



Faculty of Security Studies, University of Belgrade
**International Journal of Contemporary
 Security Studies (IJCSS)**



Article

From Threat to Meta-Risk: Radicalisation, Agential Risk, and the Governance of Capability Formation

Vipul V. Tamhane^{1,2,*}

¹ School of Law, Northumbria University, Newcastle upon Tyne, NE1 8ST, United Kingdom;

² Department of Defence and Strategic Studies, Savitribai Phule Pune University, Pune, 411007, India.

* Correspondence: vipul.tamhane@gmail.com.

Received: 17 June 2026; Revised: 25 July 2026; Accepted: 3 August 2026; Published: 24 August 2026.

ABSTRACT

This is a conceptual, theory-building article; it does not offer an empirical demonstration that radicalisation amplifies catastrophic risk. It asks how social-level radicalisation should be conceptualised in an era in which dual-use technologies place high-consequence destructive capability within reach of sub-state actors, and where the resulting risk can still be governed. Using a documented structured qualitative synthesis across terrorism studies, existential-risk scholarship, climate-security research, deterrence theory and proliferation-financing governance, the article develops the construct of meta-risk: a persistent societal condition that raises the probability or severity of two or more causally independent hazard classes, principally by degrading the shared institutional capacities on which responses to all of them depend. It specifies necessary and sufficient conditions for the construct, a discriminating observable property separating it from threat multipliers and systemic vulnerability, and a mechanism linking population-level radicalisation to rare catastrophic capability formation without equating the two. The dependent phenomenon is bounded as global catastrophic-risk amplification rather than human extinction. The governance corollary situates financial controls as one layer within a rights-constrained, layered capability-governance architecture.

KEYWORDS

Radicalisation; agential risk; capability formation; global catastrophic risk; deterrence by denial; security governance.

1. Introduction

1.1. Statement of significance

The widening asymmetry between the destructive capability that small groups may acquire and the institutional capacity of states to detect and contain them is among the defining security problems of the present century. Existential-risk scholarship has established that a catalogue of anthropogenic hazards, principally nuclear arsenals, engineered pathogens, advanced artificial intelligence and anthropogenic climate change, now carries the theoretical potential to inflict catastrophe at civilisational scale (Bostrom, 2002; Ord, 2020). Radicalisation sits uneasily within that catalogue. Measured



as a direct cause of death, terrorism is a comparatively minor hazard: the Institute for Economics and Peace recorded 5,582 terrorism fatalities worldwide in 2025, the lowest annual total since 2007 (Institute for Economics and Peace, 2026). Read in isolation, that figure invites complacency; read as the whole measure of the threat, it mistakes an outcome indicator for a risk assessment. This article argues that the security significance of radicalisation lies less in the fatalities it produces than in its relationship to capability formation and to the institutional capacities on which responses to graver hazards depend.

1.2. Research objective and the working definition

The objective is to reconceptualise social-level radicalisation as a *meta-risk* and to derive the governance consequences of that reconceptualisation. Throughout this article, and without variation, the following definition is used:

A meta-risk is a persistent societal condition that raises the probability or severity of two or more causally independent hazard classes, principally by degrading the shared institutional capacities on which societal responses to those hazards depend.

Section 3.2 states the necessary and sufficient conditions of this construct, distinguishes it from adjacent concepts, and specifies the observable property by which the distinction can be tested. The definition is fixed at this point and is **not extended, loosened or re-specified** later in the argument.

1.3. The research gap and problem statement

Three literatures bear on this problem and rarely speak to one another. Terrorism and radicalisation studies analyse the aetiology of mobilisation and the practice of prevention, and remain internally divided over whether radicalisation denotes a coherent object of study at all (Schmid, 2013). Existential-risk studies have produced the agential-risk framework, which asks not only what technologies could produce catastrophe but *who* might wield them, yet treat terrorism largely as illustration rather than as a governance object (Torres, 2016a). Climate-security scholarship models environmental stress as a threat multiplier for conflict and extremism but seldom connects that logic to catastrophic-risk governance (Asaka, 2021). The gap is structural: no framework integrates the willing-and-able logic of agential risk with the governance question of where prevention can operate when motivation is durable and technological knowledge cannot be withdrawn. The practical stake is direct. If terrorism fatalities fall while the accessibility of high-consequence capability rises, a security posture calibrated to attack counts will be measuring the wrong quantity.

1.4. Research question

The article addresses one principal question: *if social-level radicalisation is best understood as a meta-risk, through what mechanism does it raise catastrophic risk, and where can that risk still be governed given that catastrophic motivation is unlikely to be eliminated and technical knowledge cannot be un-invented?*

1.5. Thesis statement

The article advances four bounded claims. First, radicalisation is better modelled as a second-order meta-risk than as a first-order threat: its direct contribution to catastrophic outcomes is small, while its amplifying contribution operates through capability formation and through degraded societal response capacity. Second, the population willing to pursue catastrophic harm is unlikely to be reduced to zero, though its prevalence, composition and operational opportunity are all policy-sensitive. Third, technical knowledge cannot be withdrawn once created, but accessibility, materials,

competence, integration and deployment all remain governable. Fourth, it follows that the most tractable object of governance is capability formation, of which the financial substrate is an important but non-exclusive layer within a broader, rights-constrained architecture.

1.6. Research paradigm

The study is interpretivist and theory-building. It proceeds by documented structured qualitative synthesis rather than hypothesis-testing against original data, and its outputs are a construct, a mechanism, a set of falsifiable propositions and a governance model. The full design, corpus and coding procedure are reported in Section 4.

1.7. Scope conditions and disclaimers

Four scope conditions bound the argument. First, the dependent phenomenon is *global catastrophic-risk amplification*, not human extinction; the article does not claim that radicalisation produces plausible extinction pathways for every hazard discussed, and the term extinction is used only where the cited literature uses it. Second, the framework applies most strongly where catastrophic capability is resource- and network-intensive, and least where harm is improvised and near-costless. Third, the article distinguishes throughout between population-level radicalisation and the rare catastrophic tail, and does not treat them as equivalent empirical phenomena. Fourth, catastrophic intent is treated as ideologically plural; the article expressly rejects framings that attribute it to any single religion, community or region, both because the evidence does not support them and because such framings have themselves motivated mass-casualty violence.

1.8. Structure of the article

Section 2 reviews the three literatures. Section 3 sets out the conceptual framework: a definitional ladder, the meta-risk construct with its discriminating property, and the mechanism linking radicalisation to capability formation. Section 4 reports the research design, corpus, coding procedure and falsifiable propositions. Section 5 develops the willing-and-able architecture, replacing earlier absolute formulations with bounded ones. Section 6 presents two structured case illustrations. Section 7 differentiates the four amplification pathways. Section 8 reformulates societal response capacity in measurable institutional terms. Section 9 develops the layered governance model and its rights constraints. Sections 10 and 11 discuss implications, limitations and conclusions.

2. Literature review

Radicalisation entered the security lexicon as an explanatory bridge between grievance and violence, but the concept has never stabilised. Sedgwick observed that it is deployed simultaneously at individual, group and mass-public levels with little agreement on where radical ends and extremist begins (Sedgwick, 2010). Neumann distinguished cognitive from behavioural radicalisation and warned that conflating them licenses the policing of belief rather than conduct (Neumann, 2013). Kundnani traced the concept's migration from descriptive category to instrument of suspicion disproportionately applied to Muslim populations (Kundnani, 2012). These critiques are load-bearing here: they establish radicalisation as a population-level social process, and they impose a normative discipline that Section 8 takes seriously. Classical terrorism scholarship established terrorism as instrumental, communicative violence oriented to political audiences (Crenshaw, 1981; Hoffman, 2006). That instrumental model strains against actors who seek not to coerce an audience but to precipitate catastrophe, which is precisely the residual category the agential-risk literature isolates.

Existential-risk studies supply the second register. Bostrom's taxonomy defined existential risks as those that would annihilate humanity or permanently curtail its potential, identifying anthropogenic, technology-enabled hazards as dominant (Bostrom, 2002). His vulnerable-world hypothesis formalised the intuition animating this article: that technological invention may yield a capability at once highly destructive and widely accessible, such that a civilisation lacking extraordinary preventive governance is devastated by default (Bostrom, 2019). Ord's synthesis ranked engineered pandemics and advanced artificial intelligence above natural and conventional hazards (Ord, 2020). Within this tradition Torres insisted that risk analysis attend to the agent as well as the tool, arguing that any account of catastrophic risk is incomplete without study of who might wish to use such technologies, and that the per-individual probability need not be high for aggregate risk to become substantial across a century (Torres, 2016b, 2018). The framework's value here is analytic rather than predictive: it separates two variables, willingness and ability, that most policy discourse conflates.

The third literature connects environmental stress to violence, and it is the most contested. Work in this area models climate change as a threat multiplier operating indirectly through migration, resource competition, economic disruption and state fragility rather than as a direct cause of extremism (Asaka, 2021; Buhaug, 2015). Mach and colleagues, synthesising expert judgement, concluded that climate variability has influenced organised armed conflict while ranking below other drivers (Mach et al., 2019). Foresight studies report expert expectation that environmental pressures will influence radicalisation across ideological types rather than within any one (Jaworska & Nosarzewski, 2024). European counter-terrorism coordination has mapped climate narratives across right-wing, left-wing and misanthropic-environmental extremisms (EU Counter-Terrorism Coordinator, 2024). Regional analyses show the mechanism operating through resource scarcity, governance weakness and insurgent persistence rather than as automatic causation (Unaam, 2025). The disputed Syrian case marks the boundary of the evidence: where Kelley and colleagues linked anthropogenic drought to migration and pre-uprising conditions, Selby and colleagues disputed both the strength and the direction of that chain (Kelley et al., 2015; Selby et al., 2017). Recent security scholarship has correspondingly differentiated climate change as threat, as risk multiplier and as a critical-infrastructure and cascading-effects problem (Mancic, 2025).

Deterrence theory supplies the fourth. Schelling's punishment-centred model presupposes an identifiable actor who values something that can be held at risk (Schelling, 1966). Snyder's distinction between deterrence by punishment and deterrence by denial opens the alternative this article develops (Snyder, 1961). Contemporary scholarship has moved further towards denial, resilience and integrated approaches suited to diffuse and sub-threshold threats, including multi-actor arrangements in which deterrent effect arises from networked capacity rather than the threat of force (Mazarr, 2018; Pavic & Berisa, 2026). Work on identity and legitimacy in nuclear bargaining further cautions against assuming that ideologically committed actors are uniformly irrational or undeterable (Hetteh, 2026).

Three gaps follow. The radicalisation literature theorises process but lacks vocabulary for cross-domain amplification. Existential-risk studies theorise the agent-tool coupling but stop short of institutional prescription. Climate-security scholarship supplies a multiplier logic disconnected from capability governance. This article's contribution is to define a construct that integrates the three, to specify its mechanism, and to derive a governance architecture from it.

3. Conceptual framework

3.1. A definitional ladder

Much confusion in this field arises from the elision of distinct phenomena under a single term. This article therefore distinguishes nine positions on a ladder running from belief to operational preparation, set out in Table 1. The critical point is that movement up the ladder is neither automatic nor common: the overwhelming majority of individuals at the lower rungs never approach the upper ones, and the base rates at each successive rung fall by orders of magnitude.

Table 1. A definitional ladder from belief to operational preparation.

Rung	Term	Defining feature	Indicative base rate
1	Radical ideas	Beliefs departing sharply from prevailing norms	Very common
2	Cognitive radicalisation	Adoption of an extremist belief system	Common
3	Behavioural radicalisation	Movement towards action consistent with those beliefs	Uncommon
4	Violent extremism	Endorsement or use of violence for ideological ends	Rare
5	Terrorism	Instrumental violence against civilians directed at an audience	Rare
6	Catastrophic intent	Willingness to inflict indiscriminate mass casualties	Very rare
7	Apocalyptic intent	Catastrophe sought as prophesied or accelerating end	Extremely rare
8	Omnicidal intent	Human extinction sought as such	Vanishingly rare
9	Operational preparation	Concrete steps to assemble mass-casualty capability	Vanishingly rare

Two consequences follow for the argument. First, claims about population-level radicalisation cannot be transferred to the catastrophic tail, and vice versa. Second, the analytically relevant transition for catastrophic risk is not from radical ideas to violence, which the terrorism literature already treats extensively, but from catastrophic intent to operational preparation, which it treats far less.

3.2. *Meta-risk: conditions and discriminating property*

The definition given in Section 1.2 is satisfied when four conditions hold jointly. N1 (non-sufficiency): the condition is not by itself sufficient to produce the catastrophic outcome. N2 (cross-domain reach): it raises the probability or severity of at least two hazard classes that are causally independent of one another. N3 (shared mediation): at least part of that amplification operates through capacities common to responses across those classes, such as coordination capacity, institutional legitimacy or information integrity, rather than solely through hazard-specific mechanisms. N4 (persistence): it is a standing condition rather than a discrete shock. These four conditions are individually necessary and jointly sufficient.

The reviewer's question of what observably distinguishes a meta-risk from a cross-domain threat multiplier or a systemic vulnerability is answered by N3, which yields a discriminating test. A threat multiplier's effect is hazard-specific: it is transmitted through mechanisms peculiar to the hazard it intensifies, and would vanish if that hazard were removed. A meta-risk's effect is mediated by capacities shared across hazards, and therefore persists when any single hazard is removed. Operationally, the two are separable by mediation analysis: if the cross-hazard amplification attributable to a candidate condition attenuates substantially once shared response-capacity variables are statistically controlled, the condition is behaving as a meta-risk; if it does not attenuate, it is behaving as a set of independent hazard-specific multipliers. Systemic vulnerability, by contrast, is a property of a network's coupling rather than a standing societal condition, and fails N4 when it describes structural configuration rather than persistent process. Table 2 records these distinctions.

Table 2. Meta-risk and adjacent concepts, with the property that separates them.

Concept	Defining feature	Mediation	Separating test
Threat	Direct, first-order harm	None	Removing it removes the harm
Threat multiplier	Intensifies a specific hazard	Hazard-specific	Effect vanishes if that hazard is removed
Systemic vulnerability	Property of network coupling	Structural	Describes configuration, not persistent process (fails N4)
Cascading / compound risk	Escalating secondary failures; concurrent stressors	Event sequence	Observed during shock propagation, not as a standing condition
Enabling condition	Background factor permitting harm	Permissive	Necessary for harm but not amplifying across domains
Meta-risk	Standing condition raising probability or severity across two or more independent hazard classes	Shared response capacity	Cross-hazard effect attenuates when shared capacity is controlled (N3)

Whether social-level radicalisation in fact satisfies N1-N4 is a proposition this article advances and specifies for testing, not a finding it reports.

3.3. *The mechanism: how population-level radicalisation relates to rare catastrophic capability*

The most serious objection to any meta-risk account of radicalisation is that it slides from a common social process to a vanishingly rare category of actors. The objection is well founded, and the framework meets it by denying that the relationship is one of conversion. Increases in ordinary political or social radicalisation do not turn ordinary radicals into omnicidal agents. Three distinct sub-mechanisms are proposed instead, with markedly different evidential standing.

M1, tail generation. A larger radicalised base may yield more draws from a fixed low-probability tail. This is the weakest of the three: it assumes a stable tail proportion, and the evidence for such stability across time and setting is absent. It is stated here for completeness and is expressly *not* load-bearing for the argument. The earlier claim that demographic growth necessarily increases the absolute number of catastrophic actors is withdrawn as unsupported; establishing it would require longitudinal, cross-national data that do not presently exist.

M2, substrate provision. This is the primary mechanism. Radicalised milieus supply the intermediate goods from which capability is assembled: financing channels, procurement facilitation, technical personnel, physical operating space, organisational continuity, and ideological legitimisation that lowers moral inhibition against indiscriminate harm. The rare actor with catastrophic intent does not build capability alone; the milieu determines whether intent encounters the means of its realisation. Section 6 traces this mechanism in two documented cases.

M3, response degradation. Radicalisation is proposed to degrade the societal capacities on which detection, interdiction and crisis response depend, through reduced institutional trust and cooperation, contested information environments and delegitimised institutions. This is the mechanism that satisfies N3, and Section 8 renders it measurable.

The proposed relationship is therefore substrate and degradation, not conversion. This formulation has the further advantage of being falsifiable: if capability formation in documented cases proves independent of the surrounding radicalised milieu, M2 fails.

4. Research design and methodology

4.1. Design rationale

The article is conceptual security studies; its method is structured qualitative synthesis oriented to theory-building. The object of study cannot be observed directly, sits at the tail of the probability distribution where frequentist data are sparse by definition, and is constituted by intent and capability trajectories as much as by realised events. Where an outcome is rare, latent and multi-causal, disciplined integration of heterogeneous evidence into an explicit model with testable implications is the appropriate method, and the tradition of theory development from cases in security studies supports it (George & Bennett, 2005). The synthesis is thematic and framework-based: recurring mechanisms are identified across the corpus and mapped onto the willing-able-capability scaffold (Barnett-Page & Thomas, 2009; Thomas & Harden, 2008).

The earlier version of this argument rejected process tracing categorically. That rejection was too broad and is withdrawn. While the aggregate meta-risk proposition is too diffuse for process tracing, the method is well suited to bounded episodes of capability formation, and Section 6 applies it to two such cases (Beach & Pedersen, 2019). Structured synthesis in the spirit of analytic eclecticism governs the integration of mechanisms across research traditions with different epistemological standards (Sil & Katzenstein, 2010).

4.2. Corpus construction and synthesis procedure

Transparency requires an accurate account of how the corpus was built. It was assembled *purposively*, not through an exhaustive systematic search, and no claim of PRISMA-style completeness is made. Sources were identified through targeted searching of Scopus, Web of Science, Google Scholar and JSTOR, together with the repositories of the Institute for Economics and Peace, the Financial Action Task Force, the Council of the European Union, RAND and the Combating Terrorism Center, and were extended by backward and forward citation chaining from seed works in each of the four literatures. The search covered material published between 1961, the date of the earliest deterrence-theoretic work retained, and the final search date; English-language sources only were included, which is a recognised limitation given the regional concentration of relevant scholarship.

Inclusion criteria were: direct bearing on one of the four literatures; peer-reviewed publication, or institutional provenance with documented methodology in the case of index and policy reports; and analytic rather than merely descriptive treatment. Exclusion criteria were: advocacy material without method; sources whose claims could not be traced to a primary basis; and material duplicating a source already retained. Quality was appraised on three tiers, namely peer-reviewed empirical work, peer-reviewed conceptual work, and institutional or grey literature, with the last used for descriptive and quantitative anchoring rather than for causal claims. Coding proceeded in three passes: open coding of proposed mechanisms; axial grouping into the willing, able, amplification and governance families; and a confirmatory-bias check in which each retained proposition was deliberately re-examined against the strongest opposing source located, with the Syrian drought dispute and the securitisation critique serving as the principal instances.

A supplementary evidence matrix linking every major proposition to the sources supporting or challenging it, together with the final source inventory and the coding frame, is provided as supplementary material. The final search was conducted on 15 July 2026. The process identified 87 records across Scopus, Web of Science, Google Scholar, JSTOR, and the institutional repositories of the Institute for Economics and Peace, the Financial Action Task Force, the Council of the European Union, RAND, and the Combating Terrorism Center. After removal of 12 duplicates, 75 were screened at title and abstract, of which 62 were retained for full-text assessment and 48 are cited. Exclusions comprised 13 as off-topic, in that they did not bear directly on the four literatures or on the radicalisation-capability nexus; 3 as lacking a traceable method or evidential basis beyond advocacy; and 4 as duplicative of sources already retained.

4.3. Falsifiable propositions

The article's implications are stated as five formal propositions in Table 3, each with its variables, mechanism, indicator, competing explanation and disconfirming evidence.

Table 3. Falsifiable propositions.

No.	Proposition	Variables (independent to dependent)	Observable indicator	Competing explanation	Disconfirming evidence
P1	Capability formation leaves more reliable financial than ideological signatures	Capability-formation activity to detection yield by signature type	Comparative detection rates across documented cases	Selection bias in which cases become documented	Cases in which ideological signatures consistently precede and outperform financial ones
P2	Ecosystem rather than ideology predicts operational advancement	Milieu resources to highest rung reached (Table 4)	Rung attained against an ecosystem index	Idiosyncratic group competence or leadership	Matched groups with equivalent ecosystems diverging by ideology alone
P3	Radicalisation degrades societal absorptive capacity	Radicalisation prevalence to Table 7 dimensions	Panel measures of trust, cohesion, information integrity	Reverse causation; common economic or governance shock	No association once economic and governance shocks are controlled
P4	Degraded absorptive capacity impairs response across independent hazards	Table 7 dimensions to cross-hazard response performance	Crisis-response outcomes across differing hazard types	Hazard-specific capacity differences	Cross-hazard effect disappears when hazard-specific capacity is controlled
P5	Capability governance perceived as illegitimate degrades absorptive capacity	Perceived legitimacy of controls to trust, de-risking, displacement	Survey measures plus de-risking and displacement indicators	Pre-existing distrust unrelated to controls	No change in trust or displacement following expansion of controls

4.4. Limitations

Three limitations bound the contribution. The synthesis establishes the direction and mechanism of proposed amplification, not its magnitude. The purposive corpus cannot exclude selection effects, which the confirmatory-bias pass mitigates but does not eliminate. And the framework's central construct, societal absorptive capacity, is at present better specified than measured; Section 8 proposes indicators but does not validate them.

5. The agential-risk architecture: willing and able

5.1. The willing-able coupling

Catastrophe of the kind considered here requires the conjunction of a capable means with a willing agent. Existential-risk analysis long concentrated on the means, treating the agent as undifferentiated. The agential-risk turn reverses the emphasis: emerging technologies progressively satisfy the able variable for a growing number of agents, which makes the question of who would be willing increasingly consequential (Torres, 2018). Figure 1 renders the resulting architecture, incorporating the capability ladder developed in Section 5.4.




Figure 1. The willing-able architecture and the capability ladder. Social-level radicalisation is proposed to raise catastrophic risk not by converting ordinary radicals into catastrophic actors but through substrate provision (M2) and response degradation (M3). Solid arrows denote proposed causal direction; the governance panel marks points at which the ladder can be interrupted. Only rung 1 is irreversible.

The framework's diagnostic device is a thought experiment: who would press a button that ended human life? The exercise excludes the overwhelming majority of political violence. The insurgent seeking territory, the separatist seeking a state and the coercive terrorist seeking concessions all require a population whose allegiance or acquiescence is the objective; indiscriminate annihilation is self-defeating for each (Torres, 2018). What remains is a residual category for whom destruction is an end, or a prophesied good.

5.2. A taxonomy of catastrophic agents

Apocalyptic agents understand catastrophic violence as instrument or accelerant of a prophesied end. The eschatological strand within the self-styled Islamic State is the most documented recent instance (McCants, 2015; Cottee 2017). It must be stated precisely that this is a marginal and widely repudiated ideology, not a property of a faith of nearly two billion adherents, the great majority of whom reject such violence and who constitute the overwhelming share of its victims. Apocalypticism of the relevant kind is cross-traditional: Aum Shinrikyo, a syncretic Japanese movement, pursued chemical and biological weapons in explicit anticipation of an engineered Armageddon (Danzig et al., 2011; Juergensmeyer, 2017). *Omnicidal and misanthropic* agents desire human extinction as such, whether from radical anti-humanism or the conviction that the species is a planetary pathology. *Idiosyncratic* actors have private motivations irreducible to any programme (Torres, 2018). Contemporary far-right accelerationism and violent misanthropic environmental extremism belong to the same analytic family despite opposed ideological content (EU Counter-Terrorism Coordinator, 2024; Torres, 2023).

5.3. The willing variable: bounded, not fixed

The earlier formulation described the willing population as politically irreducible. That claim was too deterministic and is replaced. The defensible proposition is that this population is *unlikely to be reduced to zero* within any realistic policy horizon, while its prevalence, composition, persistence and operational opportunity all vary and are policy-sensitive. It is expected to contract where de-radicalisation and disengagement programmes function, where grievances are addressed through legitimate channels, where institutional legitimacy is high, and where community-based prevention operates with consent. It is expected to expand under conditions of state collapse and impunity, active conflict, sustained repression experienced as illegitimate, and dense ideological ecosystems that supply legitimation and contact. The policy implication is not that motivation-focused work is futile, but the opposite: such work is one of the layers in the architecture of Section 9. What does not follow is the assumption that it can drive the relevant population to zero, which is why capability-side governance remains necessary.

5.4. The able variable: a capability ladder

The earlier formulation described technological capability as irreversible. That claim conflates the existence of knowledge with the capacity to use it. Only the first rung of Table 4 is genuinely irreversible; every rung above it remains governable, and it is precisely this that makes capability governance possible. Diffusion is also domain-specific: nuclear capability remains gated by fissile-material production, which is detectable and highly controlled; biological capability is gated less by materials than by tacit skill and increasingly by synthesis screening; cyber and drone capability are the most diffused and least gated; and advanced artificial-intelligence capability is gated principally by compute, data and model access. These domains should not be collapsed into a single able variable, and the pathway analysis in Section 7 treats them separately.

Table 4. The capability ladder: what is irreversible and what remains governable.

Rung	Requirement	Governable	Principal instrument
1	Knowledge exists in the public domain	No	None available
2	Practical accessibility of protocols, designs and tools	Partly	Publication norms, synthesis screening, model and compute access
3	Acquisition of materials, precursors or components	Yes	Export control, biosecurity, procurement monitoring
4	Technical competence to execute reliably	Partly	Personnel reliability, control of specialist training access
5	Integration into a working operational system	Yes	Detection, interdiction, supply-chain controls
6	Deployment at the intended scale	Yes	Physical security, protective and consequence-management measures

5.5. Scope across actor types

Because the framework locates danger in a coupling rather than in an ideology, its reach extends across actor types while the balance of variables differs. Ideological extremists supply the clearest instances of catastrophic willingness. Lone actors matter less for the scale of their intent than for their invisibility to actor-centred interdiction. Organised criminal networks rarely harbour omniscid-al aims but are consequential as capability brokers trafficking materials, expertise and finance. State proxies occupy the most dangerous cell, coupling state-grade ability to sub-state willingness, and

the sponsorship literature documents both the transfer and its limits (Byman, 2005). Cyber actors extend the logic into a domain where the tacit-knowledge barrier is lowest and attribution hardest. The framework has high applicability where capability is resource- and network-intensive, moderate applicability where it is partially industrialised, and limited applicability to improvised, near-cost-less harm.

6. Two structured case illustrations

Two documented cases are traced through the capability-formation sequence: financing, procurement, expertise, materials, networks, detection opportunity, intervention and outcome. They are selected as a matched pair on intent and contrasting outcome, and are illustrative rather than representative.

6.1. Aum Shinrikyo: intent without capability

Aum Shinrikyo possessed what few non-state actors have possessed: explicit apocalyptic intent, substantial assets, front companies, university-trained chemists and biologists recruited from its own membership, dedicated facilities and freedom from meaningful scrutiny. It nonetheless failed to achieve mass-casualty biological capability. Its anthrax programme cultivated a vaccine strain rather than a virulent one; its dispersal engineering failed repeatedly; and its 1995 sarin release, though lethal, produced casualties far below what its investment implied (Danzig et al., 2011; Tucker, 2000). The decisive constraints were tacit knowledge, quality control and dispersal engineering, not intent or funding. Detection opportunities existed and were not taken: prior chemical incidents, procurement anomalies and a documented earlier attack were not aggregated into an interdiction. The case is important *negative* evidence. It demonstrates that catastrophic intent does not translate into catastrophic capability, that resources are not sufficient, and that the capability ladder of Table 4 has real rungs.

6.2. The Islamic State's chemical programme: ecosystem-enabled advancement

The Islamic State presents the contrasting outcome. Territorial control, diversified revenue, captured industrial precursors, access to technically trained personnel and functioning procurement networks enabled indigenous production of sulphur mustard and its repeated battlefield use in Iraq and Syria. Yet the same record documents low technical sophistication and limited lethality relative to state programmes (Strack, 2017). The instructive contrast with Aum is not that intent differed in kind but that the surrounding ecosystem differed: territory, revenue and personnel converted intent into sustained, if crude, operational output. This is the clearest available illustration of mechanism M2.

6.3. Comparative analysis

Table 5. Comparative case analysis of capability formation.

Dimension	Aum Shinrikyo	Islamic State (chemical programme)
Catastrophic intent	Present and explicit	Present and explicit
Financing	Substantial internal assets and front companies	Diversified territorial revenue
Procurement	Open commercial purchase, largely unimpeded	Captured stocks and regional procurement networks

Expertise	University-trained chemists and biologists recruited internally	Technically trained personnel within controlled territory
Materials	Acquired; biological agent strain ineffective	Acquired industrial precursors
Networks and space	Dedicated facilities, minimal scrutiny	Territorial control and organisational continuity
Detection opportunity	Multiple missed signals, including a prior attack	Limited external visibility during territorial phase
Intervention	Post hoc, after mass-casualty attack	Military degradation of territory and personnel
Highest rung reached (Table 4)	Rung 5 in chemistry; failed at rung 4 in biology	Rung 6 at low technical sophistication
Outcome	Lethal but far below invested potential	Repeated crude battlefield use, limited lethality

Three inferences follow, each bounded. First, across both cases the binding constraint sat on the ability side, which supports the article's governance emphasis. Second, the ecosystem, not the ideology, best explains the difference in operational advancement, which supports M2 and is consistent with M3 given the missed detection opportunities in the Aum case. Third, the earlier assertion that the willing did not change across these cases is withdrawn: their ideologies, resources, structures, technical personnel, geography and objectives differed materially, and the comparison holds only on the narrow dimension of catastrophic intent being present in both.

7. Differentiated amplification pathways

The four pathways require separate causal treatment: they involve different actors, technologies, time horizons and mediating institutions, and they differ substantially in evidential support. Table 6 sets them out on common attributes.

Table 6. Differentiated amplification pathways.

Attribute	Nuclear	Biological	Artificial intelligence	Climate
Initiating actor	State decision-makers reacting to sub-state action	Sub-state actor or insider	Sub-state actor using available systems	Sub-state actor within fragile governance
Relevant form of radicalisation	Any producing mass-casualty attack in a rivalry	Catastrophic or apocalyptic intent	Instrumental extremism of any type	Grievance-based mobilisation
Capability required	None by the initiator; state arsenals supply effect	High: agent, competence, dispersal	Low to moderate; systems widely available	None; operates through conditions
Capability-formation pathway	Not applicable; escalation is the mechanism	Materials, tacit skill, integration	Access to models, tools and compute	Not applicable
Potential consequence	Catastrophic to civilisational	Catastrophic if engineered agents mature	Amplifying rather than terminal	Indirect; conflict and displacement
Mediating institutions	Attribution, command and control, escalation doctrine	Biosecurity, synthesis screening, health surveillance	Model governance, safety systems, platform control	Governance capacity, adaptation, humanitarian systems

Evidential support	Historical near-misses; no case of sub-state-triggered nuclear use	Consistent record of non-state failure	Documented misuse; prospective for higher classes	Contested and indirect
Alternative explanation	Escalation driven by state rivalry independent of sub-state action	State programmes remain the principal risk	Harms attributable to platforms, not radicalisation	Governance failure and poverty as primary drivers
Uncertainty	Moderate	High	High	Very high

7.1. Nuclear

The proposition is not that a mass-casualty attack within a nuclear rivalry automatically produces escalation; the historical record does not support that, and several such attacks have not produced it. The claim is narrower and conditional. Escalation risk from sub-state action rises when four conditions co-occur: attribution to a state adversary is rapid, confident and politically consequential; the attack occurs within an active rivalry where retaliatory doctrine is pre-committed; command-and-control or early-warning systems are implicated, degrading decision time; and domestic political conditions raise the cost of restraint. Where attribution is slow, contested or directed at a non-state actor without a state patron, escalation is markedly less likely. The relevant vulnerabilities are therefore attributional and organisational rather than ideological, and the literature on the limits of organisational safety in nuclear systems supplies the mechanism (Sagan, 1993). A related caution follows from work on identity and legitimacy in nuclear bargaining: ideologically committed actors are not uniformly beyond deterrent or bargaining logic, and treating them so degrades analysis (Hetteh, 2026).

7.2. Biological

Five distinct phenomena are frequently conflated and are separated here: naturally occurring outbreaks; laboratory accidents; deliberate misuse of existing pathogens; synthetic-biology-enabled weaponisation; and genuinely catastrophe-capable engineered agents. Only the last two are candidates for civilisational-scale consequence, and they are the least documented in the non-state record. The empirical record of non-state biological attempts is one of consistent failure at the competence and integration rungs (Carus, 2000; Tucker, 2000). The pathway's significance is therefore prospective, resting on the proposition that synthesis and design tools are lowering the tacit-knowledge barrier that historically defeated such attempts, and on the vulnerable-world logic that follows if they do (Bostrom, 2019).

7.3. Artificial intelligence

Six risk classes are conflated in much commentary and are separated here: AI-assisted propaganda and recruitment; lowered access barriers to harmful knowledge; automation of cyber operations; manipulation or circumvention of safety systems; autonomous weapons; and misaligned advanced systems. Only the first four are plausibly connected to radicalised non-state actors, and this article makes claims about those only. Misaligned advanced artificial intelligence is a distinct risk class involving different actors, mechanisms and governance instruments; it is *not* an agential-risk problem of the kind analysed here, and the framework should not be read as implying otherwise. That AI governance has become a codified policy object in its own right reflects this separation (National Institute of Standards and Technology, 2023).

7.4. Climate

This is the weakest pathway and is qualified accordingly. The cited literature supports an indirect, context-dependent relationship among climate stress, governance capacity, conflict and radicalisation; it does not establish that radicalisation materially amplifies climate change itself, and no such claim is made. The pathway asserted here is narrower and operates in one direction: climate stress may expand the conditions under which recruitment succeeds and state capacity degrades, principally through resource scarcity, displacement and governance failure in already fragile settings (Jaworska & Nosarzewski, 2024; Unaam, 2025). Four other candidate mechanisms, political obstruction of climate governance, attacks on critical infrastructure, environmentally motivated extremism, and cascading damage caused by conflict, are analytically distinct and are not aggregated into a single claim here. The contested status of climate-conflict causation is treated as a constraint on the whole pathway rather than as an objection confined to a later subsection (Mach et al., 2019; Selby et al., 2017).

8. Societal absorptive capacity

8.1. Retiring the immunological metaphor

An earlier formulation of this argument described radicalisation as eroding civilisational immunity and invoked the language of civilisational cancer. That framing is withdrawn. The objection is not merely that the metaphor is imprecise, though it is; it is that biological metaphors of infection and immunity carry a serious normative hazard in radicalisation research, since they position disfavoured populations as pathogens or contamination and thereby license exactly the pathologising of dissent that the critical literature warns against (Kundnani, 2012; Sedgwick, 2010). Given that this article's own scope conditions reject essentialising framings, retaining such a metaphor would be inconsistent.

The construct is therefore reformulated in institutional terms as **societal absorptive capacity**: the set of measurable institutional and informational capacities that determine whether a society absorbs or amplifies a shock. This is closer to established work on resilience and social capital, and it has the advantage of being operationalisable rather than evocative (Bourbeau, 2015; Putnam, 2000).

8.2. Measurable dimensions

Table 7. Societal absorptive capacity: dimensions and candidate indicators.

Dimension	What it captures	Candidate indicator
Institutional trust	Confidence in public institutions to act competently and fairly	Standardised trust items in population surveys
Interpersonal trust and social cohesion	Willingness to cooperate across group boundaries	Generalised trust and intergroup contact measures
Legitimacy	Perceived rightfulness of authority and procedure	Procedural-justice and compliance measures
Information integrity	Existence of a shared, reliable evidential basis	Disinformation prevalence and source-trust measures
Coordination capacity	Ability to act collectively across institutions	Inter-agency exercise and response performance
Crisis-governance capacity	Institutional ability to manage acute shocks	Documented crisis-response outcome assessment
Resilience of essential systems	Continuity of critical services under stress	Service-continuity and recovery-time indicators

Recent security-studies work demonstrates that several of these dimensions, including institutional trust, perceived threat and dimensions of digital and personal security, can be measured through structured instruments in defined populations, which offers a practical route to operationalisation (Cvetkovic et al., 2025).

8.3. Status of the claim

The proposition that radicalisation degrades these dimensions, and that degraded dimensions impair response to independent hazards, is advanced here as a *testable proposition and not as an established finding*. Figure 2 renders it as such. The connection among radicalisation, contested information environments, institutional distrust and weakened crisis response is plausible and partially supported at the level of individual links, but the compound claim has not been demonstrated, and Propositions 3 and 4 in Table 3 specify what would confirm or disconfirm it.




Figure 2. The societal absorptive capacity proposition. Dashed arrows denote proposed, not demonstrated, degradation pathways; the dotted arrow denotes proposed feedback whereby governance perceived as illegitimate may itself raise radicalisation. The compound claim is specified for testing in Propositions 3 and 4 of Table 3.

9. From deterrence to layered capability governance

9.1. Conceptual precision: deterrence, denial and prevention

Not everything proposed under the heading of capability governance is deterrence. Deterrence proper requires that the adversary perceive the obstacle and adjust behaviour in anticipation; where a measure simply prevents capability formation irrespective of the adversary's beliefs, it is prevention, interdiction or regulation, and should be described as such (Snyder, 1961). The distinction matters practically. Publicised financial controls that raise expected detection probability may deter; unpublicised interdiction of a procurement channel does not deter but does prevent. Deterrence by denial occupies the middle ground: it works on feasibility rather than on the cost of consequences, and it requires perception only to the extent that the adversary must believe success unlikely (Mazarr, 2018). Against an actor who does not fear death, punishment is largely inert while denial and prevention retain force. Contemporary work on multi-actor and networked deterrence indicates that deterrent effect can arise from the density of institutional capacity rather than from the threat of force, which is consistent with the architecture proposed here (Pavic & Berisa, 2026).

9.2. A layered architecture

The earlier version of this argument presented a false trilemma: actors undeterrable, technology irreversible, therefore financial ecosystems the only tractable point of intervention. That inference does not hold, and it is replaced by a layered model. Capability formation can be governed at many points, and the financial layer is one among several. Table 8 sets out the layers, their instruments and their limits.

Table 8. Layers of capability governance, instruments and limits.

Layer	Principal instruments	Principal limit
Motivation	Prevention, disengagement, grievance redress, legitimacy-building	Cannot reduce the willing population to zero
Knowledge and access	Publication norms, synthesis screening, model and compute access governance	Cannot withdraw knowledge already public
Material	Export control, biosecurity, personnel reliability, procurement monitoring	Weak against improvised and dual-use commodity pathways
Financial	Counter-terrorism financing, counter-proliferation financing, transaction monitoring	Weak signatures for low-cost and insider pathways; rights costs
Detection	Intelligence, supply-chain monitoring, cyber defence	Attribution difficulty; displacement into unobserved channels
Resilience	Institutional capacity, crisis response, consequence management	Acts after formation rather than preventing it




Figure 3. Layered capability governance. The financial layer is one of six and is not sufficient alone. All layers operate within the rights constraint shown, which Section 9.6 treats as a condition of effectiveness rather than an external limitation.

The financial layer retains a distinctive role for three reasons: capability formation is resource-dependent and therefore leaves financial traces; the financial system is already an instrumented, choke-pointed and intelligence-rich domain; and financial measures target pathways rather than persons, which suits a threat characterised by distributed agency. But distinctive is not exclusive, and the architecture fails if any single layer is treated as sufficient.

9.3. Counter-terrorism financing and counter-proliferation financing

These are distinct regimes and are conflated at analytical cost. Counter-terrorism financing addresses the funding of violence by designated persons and entities, typically involving small sums, domestic predicate offences and listing-based sanctions. Counter-proliferation financing addresses the financing of prohibited weapons programmes and their procurement, typically involving larger transactions, complex trade structures, state-linked actors and targeted financial sanctions with a different legal basis and evidentiary standard (Financial Action Task Force, 2023). Capability formation of the kind analysed here sits closer to the second regime than the first, which is a further reason not to treat ordinary terrorist-financing controls as sufficient. The earlier citation of civil-war economics literature in support of proliferation-financing claims was inapposite and has been removed.

9.4. The limits of financial signatures

The claim that capability formation reliably leaves financial signatures requires qualification. Low-cost biological, cyber, insider and improvised pathways may leave little or no conspicuous financial trace. Even where traces exist, they may be obscured by cash transactions, informal value transfer systems, virtual assets, trade-based mechanisms, front companies, opaque beneficial ownership, transaction structuring and procurement through legitimate suppliers whose activity appears unremarkable. False positives impose their own costs, considered next. The defensible claim is therefore narrower than the one earlier advanced: financial signatures are informative for resource-intensive, network-dependent capability formation, and weak for improvised or insider pathways.

9.5. Rights, proportionality and safeguards

Capability-focused financial governance carries real costs. Documented effects of expansive financial-crime controls include de-risking and wholesale withdrawal of services from disfavoured sectors, discriminatory profiling, exclusion of individuals and organisations from the banking system, obstruction of lawful scientific research and humanitarian activity, privacy intrusion, and difficulty in obtaining correction or remedy for erroneous designation (Financial Action Task Force, 2023; Waldron, 2003). These are not incidental. They fall disproportionately on precisely the communities the radicalisation literature identifies as already over-policed, and they therefore interact directly with the mechanisms in Section 8.

The safeguards required are correspondingly specific: necessity and proportionality assessment before deployment; independent or judicial authorisation; purpose limitation and data minimisation; auditability and effectiveness review; non-discrimination testing; accessible remedy and correction of erroneous designations; and explicit protection for legitimate scientific, humanitarian and journalistic activity.

9.6. When capability governance becomes a meta-risk

This point belongs in the model rather than in a closing caveat. Capability-governance measures degrade the same dimensions of absorptive capacity that Section 8 identifies when they are perceived as illegitimate, discriminatory or unaccountable: they reduce institutional trust, supply

grievance narratives that ideological entrepreneurs exploit, and displace activity into unobservable channels, reducing the very visibility that justified them. By the criteria in Section 3.2, a governance regime that persistently degraded shared response capacities across hazard domains would itself satisfy N1 to N4 and would constitute a meta-risk. The architecture is therefore self-limiting by design: proportionality is not an external ethical constraint upon it but a condition of its own effectiveness.

10. Discussion

The framework's contribution is threefold. Conceptually, it supplies a defined construct with necessary and sufficient conditions and a discriminating test, addressing the proliferation of overlapping vocabulary in this field. Mechanistically, it specifies how a common social process relates to a rare catastrophic tail without equating the two, locating the relationship in substrate provision and response degradation rather than conversion. Prescriptively, it reorients governance from actor-centred interdiction towards layered, rights-constrained capability governance in which financial controls are one instrument among several.

Three recommendations follow for practice. Assessment should track capability-formation indicators alongside incident counts, since falling fatalities and rising investigations are consistent with many underlying trends and demonstrate nothing about catastrophic capability on their own. Institutional convergence between counter-terrorism, counter-proliferation and financial-intelligence functions should take the capability-formation network rather than the discrete incident as the shared unit of analysis. And proportionality review should be built into capability-governance instruments at design stage, for the reasons in Section 9.6.

The limitations are substantial and are stated plainly. The framework is a proposal for testing, not a validated model. Its central construct awaits measurement. Its climate pathway is weakly supported and is presented as such. Its case illustrations are two, purposively selected, and cannot establish generality. And its governance model, while layered, is derived analytically rather than evaluated empirically against programme outcomes.

11. Conclusions

This article asked how social-level radicalisation should be conceptualised in an era of diffusing high-consequence capability, and where the resulting risk can still be governed. It argued that radicalisation is better modelled as a meta-risk, defined as a persistent societal condition raising the probability or severity of two or more causally independent hazard classes principally by degrading shared institutional capacities, and it specified the conditions, the discriminating property and the mechanism that give that construct content. It replaced earlier deterministic formulations with bounded ones: the willing population is unlikely to reach zero but is policy-sensitive; technical knowledge cannot be withdrawn but accessibility, materials, competence, integration and deployment remain governable. It bounded the dependent phenomenon as global catastrophic-risk amplification rather than extinction. And it derived a layered, rights-constrained governance architecture in which the financial substrate of capability formation is an important but non-exclusive layer.

Four research priorities follow. The absorptive-capacity construct requires measurement and validation against the indicators proposed in Table 7. The propositions in Table 3 require testing, beginning with the comparative reliability of financial and ideological signatures in documented capability-formation cases. The case method applied in Section 6 should be extended to a larger and more systematically selected set, including cases of successful interdiction. And the conditions under which capability-governance regimes themselves degrade absorptive capacity require dedicated empirical study, since that question determines whether the architecture proposed here is self-limiting in practice as well as in principle.

Author Contributions: Vipul V. Tamhane is the sole author and was responsible for conceptualisation, methodology, formal analysis, investigation, and writing (original draft, review and editing). The author has read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable. This study did not involve human participants or animals.

Informed Consent Statement: Not applicable.

Data Availability Statement: No primary dataset was collected. The study did, however, generate a structured document, screening and inclusion decisions, the tabular thematic coding frame, and a proposition-to-evidence matrix.

Generative AI Statement: Generative artificial-intelligence tools were used to assist with conversion of concept to figures, structuring and language editing. They were not used to generate substantive theoretical claims. Every citation was independently verified by the author against the original publication, and each theoretical connection was checked by the author against the cited sources to ensure that no unsupported association was introduced. The author reviewed and edited all output and takes full responsibility for the content of the publication.

Acknowledgments: The author thanks the reviewers for detailed and constructive comments, which materially improved the structural precision, transparency and normative treatment of this article.

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

12. References

1. Asaka, J. O. (2021). Climate change-terrorism nexus? A preliminary review/analysis of the literature. *Perspectives on Terrorism*, 15(1), 81-92. ISSN: 2334-3745
2. Barnett-Page, E., & Thomas, J. (2009). Methods for the synthesis of qualitative research: A critical review. *BMC Medical Research Methodology*, 9, 59. <https://doi.org/10.1186/1471-2288-9-59>
3. Beach, D., & Pedersen, R. B. (2019). *Process-tracing methods: Foundations and guidelines* (2nd ed.). Ann Arbor: University of Michigan Press. <https://doi.org/10.3998/mpub.10072208>
4. Bostrom, N. (2002). Existential risks: Analyzing human extinction scenarios and related hazards. *Journal of Evolution and Technology*, 9(1).
5. Bostrom, N. (2019). The vulnerable world hypothesis. *Global Policy*, 10(4), 455-476. <https://doi.org/10.1111/1758-5899.12718>
6. Bourbeau, P. (2015). Resilience and international politics: Premises, debates, agenda. *International Studies Review*, 17(3), 374-395. <https://doi.org/10.1111/misr.12226>
7. Buhaug, H. (2015). Climate-conflict research: Some reflections on the way forward. *WIREs Climate Change*, 6(3), 269-275. <https://doi.org/10.1002/wcc.336>
8. Byman, D. (2005). *Deadly connections: States that sponsor terrorism*. Cambridge: Cambridge University Press. <https://doi.org/10.1017/CBO9780511790843>
9. Carus, W. S. (2000). The Rajneeshees (1984). In J. B. Tucker (Ed.), *Toxic terror: Assessing terrorist use of chemical and biological weapons* (pp. 115-137). Cambridge, MA: MIT Press.
10. Crenshaw, M. (1981). The causes of terrorism. *Comparative Politics*, 13(4), 379-399. <https://doi.org/10.2307/421717>
11. Cvetkovic, V. M., Lipovac, M., Renner, R., Stanarevic, S., & Raonic, Z. (2025). A predictive framework for understanding multidimensional security perceptions among students in Serbia: The role of institutional, socio-economic, and demographic determinants of sustainability. *Sustainability*, 17(11), 5030. <https://doi.org/10.3390/su17115030>

12. Danzig, R., Sageman, M., Leighton, T., Hough, L., Yuki, H., Kotani, R., & Hosford, Z. M. (2011). *Aum Shinrikyo: Insights into how terrorists develop biological and chemical weapons*. Washington, DC: Center for a New American Security. <https://www.cnas.org/publications/reports/aum-shinrikyo-insights-into-how-terrorists-develop-biological-and-chemical-weapons>
13. EU Counter-Terrorism Coordinator. (2024). The role of climate change and environmental concerns in violent extremist and terrorist radicalisation in the EU (Council document 5982/24). Brussels: Council of the European Union.
14. Financial Action Task Force. (2023). International standards on combating money laundering and the financing of terrorism and proliferation: The FATF recommendations. Paris: FATF. <https://www.fatf-gafi.org/en/publications/Fatfrecommendations/Fatf-recommendations.html>
15. George, A. L., & Bennett, A. (2005). *Case studies and theory development in the social sciences*. Cambridge, MA: MIT Press.
16. 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
17. Hoffman, B. (2006). *Inside terrorism* (Rev. ed.). New York: Columbia University Press.
18. Institute for Economics and Peace. (2026). Global terrorism index 2026: Measuring the impact of terrorism. Sydney: IEP. <https://www.economicsandpeace.org/global-terrorism-index/>
19. Jaworska, J., & Nosarzewski, K. (2024). Future nexus between climate change and terrorism: Will environmental issues influence radicalization regardless of ideology? In R. J. Chasdi & Y. Sharan (Eds.), *Issues of terrorism in the post-coronavirus era*. Cham: Springer. https://doi.org/10.1007/978-3-031-68542-2_8
20. Juergensmeyer, M. (2017). *Terror in the mind of God: The global rise of religious violence* (4th ed.). Berkeley: University of California Press. <https://doi.org/10.1525/9780520965164>
21. Kelley, C. P., Mohtadi, S., Cane, M. A., Seager, R., & Kushnir, Y. (2015). Climate change in the Fertile Crescent and implications of the recent Syrian drought. *Proceedings of the National Academy of Sciences*, 112(11), 3241-3246. <https://doi.org/10.1073/pnas.1421533112>
22. Kundnani, A. (2012). Radicalisation: The journey of a concept. *Race & Class*, 54(2), 3-25. <https://doi.org/10.1177/0306396812454984>
23. Mach, K. J., Kraan, C. M., Adger, W. N., Buhaug, H., Burke, M., Fearon, J. D., ... von Uexkull, N. (2019). Climate as a risk factor for armed conflict. *Nature*, 571, 193-197. <https://doi.org/10.1038/s41586-019-1300-6>
24. Mancic, T. (2025). Climate change as a security challenge, risk and threat of the 21st century and its consequences on critical infrastructure. *International Journal of Contemporary Security Studies*, 1(1), 191-204. https://doi.org/10.18485/fb_ijcss.2025.1.1.14
25. Mazarr, M. J. (2018). *Understanding deterrence*. Santa Monica, CA: RAND Corporation. <https://www.rand.org/pubs/perspectives/PE295.html>
26. McCants, W. (2015). *The ISIS apocalypse: The history, strategy, and doomsday vision of the Islamic State*. New York: St. Martin's Press. ISBN-13: 978-1250080905
27. National Institute of Standards and Technology. (2023). Artificial intelligence risk management framework (AI RMF 1.0) (NIST AI 100-1). Gaithersburg, MD: U.S. Department of Commerce. <https://doi.org/10.6028/NIST.AI.100-1>
28. Neumann, P. R. (2013). The trouble with radicalization. *International Affairs*, 89(4), 873-893. <https://doi.org/10.1111/1468-2346.12049>
29. Ord, T. (2020). *The precipice: Existential risk and the future of humanity*. London: Bloomsbury.
30. Pavic, A. M., & Berisa, 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

31. Putnam, R. D. (2000). *Bowling alone: The collapse and revival of American community*. New York: Simon & Schuster. <https://doi.org/10.1145/358916.36199>
32. Sagan, S. D. (1993). *The limits of safety: Organizations, accidents, and nuclear weapons*. Princeton, NJ: Princeton University Press. <https://doi.org/10.1515/9780691213064>
33. Schmid, A. P. (2013). *Radicalisation, de-radicalisation, counter-radicalisation: A conceptual discussion and literature review (ICCT Research Paper)*. The Hague: International Centre for Counter-Terrorism. <https://doi.org/10.19165/2013.1.02>
34. Sedgwick, M. (2010). The concept of radicalization as a source of confusion. *Terrorism and Political Violence*, 22(4), 479-494. <https://doi.org/10.1080/09546553.2010.491009>
35. Selby, J., Dahi, O. S., Frohlich, C., & Hulme, M. (2017). Climate change and the Syrian civil war revisited. *Political Geography*, 60, 232-244. <https://doi.org/10.1016/j.polgeo.2017.05.007>
36. Sil, R., & Katzenstein, P. J. (2010). Analytic eclecticism in the study of world politics: Reconfiguring problems and mechanisms across research traditions. *Perspectives on Politics*, 8(2), 411-431. <https://doi.org/10.1017/S1537592710001179>
37. Snyder, G. H. (1961). *Deterrence and defense: Toward a theory of national security*. Princeton, NJ: Princeton University Press. ISBN: 9780691652092
38. Strack, C. (2017). The evolution of the Islamic State's chemical weapons efforts. *CTC Sentinel*, 10(9), 19-23.
39. Thomas, J., & Harden, A. (2008). Methods for the thematic synthesis of qualitative research in systematic reviews. *BMC Medical Research Methodology*, 8, 45. <https://doi.org/10.1186/1471-2288-8-45>
40. Torres, P. (2018). Who would destroy the world? Omnicidal agents and related phenomena. *Aggression and Violent Behavior*, 39, 129-138. <https://doi.org/10.1016/j.avb.2018.02.002>
41. Torres, P. (2023). Maniacs, misanthropes, and omnicidal terrorists: Reassessing the agential risk framework. In *Proceedings of the 2023 Stanford Existential Risks Conference*. Stanford, CA: Stanford Existential Risks Initiative. <https://doi.org/10.25740/gv769bs1452>
42. Torres, P. (2016a). Agential risks: A comprehensive introduction. *Journal of Evolution and Technology*, 26(2), 31-47. <https://doi.org/10.55613/j eet.v26i2.58>
43. Torres, P. (2016b). Introducing the subfield of agential riskology. *Institute for Ethics and Emerging Technologies*. <https://ieet.org/index.php/IEET2/more/torres20160226>
44. Tucker, J. B. (Ed.). (2000). *Toxic terror: Assessing terrorist use of chemical and biological weapons*. Cambridge, MA: MIT Press.
45. Unaam, A. O. (2025). Fueling the insurgency: Climate-resource scarcity and the persistence of terrorism in Northern Nigeria. *International Journal of Contemporary Security Studies*, 1(2), 137-144. https://doi.org/10.18485/fb_ ijcss.2025.1.2.9
46. Waldron, J. (2003). Security and liberty: The image of balance. *Journal of Political Philosophy*, 11(2), 191-210. <https://doi.org/10.1111/1467-9760.00174>
47. Cottee, S. (2017). "What ISIS Really Wants" Revisited: Religion Matters in Jihadist Violence, but How? *Studies in Conflict & Terrorism*, 40(6), 439-454. <https://doi.org/10.1080/1057610X.2016.1221258>

