Routledge
Taylor & Francis Group

Contemporary Security Policy

ISSN: 1352-3260 (Print) 1743-8764 (Online) Journal homepage: www.tandfonline.com/journals/fcsp20

Securitization through inscription: National security documents and the making of a global threat lexicon (1987–2025)

Andrew W. Neal & Roy B. Gardner

To cite this article: Andrew W. Neal & Roy B. Gardner (15 Nov 2026): Securitization through inscription: National security documents and the making of a global threat lexicon (1987–2025), Contemporary Security Policy, DOI: 10.1080/13523260.2026.2753752

To link to this article: https://doi.org/10.1080/13523260.2026.2753752

© 2026 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group

Published online: 15 Nov 2026.

Submit your article to this journal

Article views: 253

View related articles

View Crossmark data

Full Terms & Conditions of access and use can be found at
https://www.tandfonline.com/action/journalInformation?journalCode=fcsp20

CONTEMPORARY SECURITY POLICY
https://doi.org/10.1080/13523260.2026.2753752

Routledge
Taylor & Francis Group

OPEN ACCESS Check for updates

Securitization through inscription: National security documents and the making of a global threat lexicon (1987–2025)

Andrew W. Neal and Roy B. Gardner

School of Social and Political Science, University of Edinburgh, Edinburgh, UK

ABSTRACT

Securitization theory holds that actors construct threats by drawing on an existing "grammar" of security. We argue that they also draw on an existing threat lexicon, qualifying the authorial agency the theory attributes to securitizing actors. Using computational text analysis, we examine how threat categories appear and are reused in 660 national security and defence documents published by 118 countries and territories between 1987 and 2025. Most countries overwhelmingly reuse categories already present in the documentary record, and the introduction of new categories is geopolitically concentrated and historically structured. We conceptualize strategy documents as inscription devices that stabilize threats as durable and portable objects, and states as selectors and adapters of an inherited lexicon more than as original authors. Securitizing agency lies in reproducing, modifying, or resisting categories made available by prior inscription, as well as in naming threats. Unequal early inscription may confer lasting agenda-setting power in strategic discourse.

KEYWORDS Securitization theory; national security strategies; computational text analysis; threat construction; critical security studies; corpus analysis

Between 1987 and 2025, the number of countries known to publish a public national security or defence strategy document grew from a handful to 118. One of the most important functions of these documents is to name the threats a country faces. This is significant for force posture, strategic planning, diplomacy, alliances, budget allocation, balancing other policy priorities, and legitimating all this with the public. Naming national security threats is a core exercise of sovereign prerogative, yet our analysis shows that more than three quarters of the threats named in these documents

CONTACT Andrew W. Neal andrew.neal@ed.ac.uk School of Social and Political Science, University of Edinburgh, Chrystal Macmillan Building, 15A George Square, Edinburgh, EH8 9LD, UK

© 2026 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group
This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives License (http://creativecommons.org/licenses/by-nc-nd/4.0/), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited, and is not altered, transformed, or built upon in any way. The terms on which this article has been published allow the posting of the Accepted Manuscript in a repository by the author(s) or with their consent.

had first been named by one of just 19 other countries. It looks like most countries are not exercising their sovereign prerogative wholly independently but following a strategic agenda set by a small number of others.

In this article we investigate this pattern, asking where the available threat categories come from, who put them there, and what this means for our understanding of sovereign security prerogatives, international security agenda-setting, and the role of national security and defence documents, chiefly as seen through the lens of securitization theory.

We argue that these top-level strategy documents work as "inscription devices" that turn security concerns into stable, portable threat categories (Latour, 1987, pp. 68–69). By selecting, naming, and reiterating classes of threat—such as "terrorism," "hybrid threats," "economic coercion," "pandemics," and "foreign interference"—they stabilize such categories as durable documentary objects that later authors can reuse within a transnational documentary field. We call this mechanism *prior inscription power*. The early institutionalization of the practice of publishing recurring, top-level strategy documents allowed a small number of countries to inscribe threat categories that became part of the repertoire from which later authors construct threats. This does not mean a category's first appearance in the corpus marks its original construction in political discourse more broadly. Rather, first inscription marks the point at which a constructed threat category enters this documentary field and becomes available for subsequent adoption and reiteration.

We pursue this argument using an existing corpus of 660 top-level national security and defence documents published between 1987 and 2025 by 118 countries and territories (the National Security and Defence Documents Dataset v3.0 or NSDDD), representing the known population of countries and territories publishing eligible documents (Neal & Gardner, 2026). Through computational text analysis, we identify threat topics, date and attribute their first inscriptions, and trace their subsequent reproduction. We treat states as singular documentary producers while recognizing that, in practice, their strategies are assembled by heterogeneous internal bureaucratic and political actors.

The findings reveal a sharply unequal but historically changing field of threat construction. Nineteen of the 118 countries and territories account for four fifths of all first inscriptions, while the United States alone accounts for one quarter. The 37 countries and territories that are both wealthy and democratic (fewer than one third of the total) account for three quarters. Nearly half of all countries and territories register no first inscription in the corpus. Yet the US does not continuously dominate first inscription: the US share falls sharply during a 1998–2004 transition as the field expands, while categories inscribed during the prior formative period remain in circulation. An early advantage in threat inscription is

thus deposited in a lexicon that persists after the advantage itself has diminished.

This reframes securitizing agency. In Barthes's terms, strategy texts, like other texts, are assemblages of already circulating formulations rather than expressions of singular authorship (Barthes, 1977, p. 146; Der Derian & Shapiro, 1989). States are rarely wholly original authors of the threat categories appearing in their strategies; they more often resemble Barthes's *scriptors*, writing strategically by selecting, recombining, and reiterating threat categories that are already available as internationally legible documentary objects (Barthes, 1977, p. 147). Their securitizing agency lies less in constructing threats anew or naming issues as threats than in selecting, adapting, combining, emphasizing, or omitting inherited threat categories. Once placed into circulation, a threat category becomes available, legible, and potentially expected to appear in future documents. Reuse may then be easier than departure from the lexicon, while omission may itself become a choice requiring explanation. The question is therefore not only why and under what conditions actors securitize, but how particular threat objects become writable, portable, and recognizable as threats in the first place.

This argument is a claim about power. Although first-mover advantage gives early participants more opportunities to record threats, this does not manifest equally in the inscribed lexicon. Prior inscription power arises when states' inscribed categories persist and structure the repertoire available to later strategies. This power appears to be geopolitically as well as historically structured. Global security discourse is shaped not simply by unequal power to securitize issues, but also by unequal power to establish the categories through which subsequent securitization becomes possible.

Theoretical framework: Documentary inscription and constrained authorship

From the broadened security agenda to documentary fields

Securitization through inscription in documents, or documentary securitization, builds on the widening of the security agenda that preceded securitization theory. Buzan challenged the identification of security with military threats alone by distinguishing military, political, economic, societal, and environmental sectors (Buzan, 1983, 1991). The Copenhagen School subsequently recast these as domains in which issues could be constructed as security threats through intersubjective processes of securitization (Buzan et al., 1998). This widening enabled securitization research across each sector: on armed conflict and military threat (Buzan et al., 1998), on political subversion and the contestation of state legitimacy (Vuori, 2008), on

economic dependency and energy supply (Nyman, 2018), on migration and collective identity (Huysmans, 2006), and on climate and environmental change (Trombetta, 2008). Cyber security has since been theorized as exceeding the original five sectors because its referent objects include networks, data, and networked subjects, and because it operates through distinctive modalities (Hansen & Nissenbaum, 2009). Our analysis follows this broad agenda by examining the full range of explicit threats identified in national security and defence documents rather than selecting issue areas in advance. To relate the inductively generated topics to this agenda, we use Buzan's five sectors together with the cyber category.

During the post-Cold War period in which the Copenhagen School developed its account, security policy remained comparatively opaque, making public declarations an especially visible site for securitizing speech acts (Buzan et al., 1998, p. 23; Neal, 2019; Wæver, 2011, p. 469). Since then, however, security has become progressively publicized through national security and defence strategy documents. By the end of 1997, 13 countries and territories had entered the documentary field; by 2025, 118 had done so, creating a new terrain of written security discourse. This has created a corpus large enough to examine documentary securitization on a systematic, large-N basis, rather than through selected cases that risk confirmation bias (Balzacq, 2011, p. 34).

The documentary mode we theorize here operates alongside, not instead of, speech act-based securitization. Its distinctive feature is that threats are routinely selected, classified, and stabilized in durable texts. States publish these documents for heterogeneous reasons, including alliance interoperability, transparency, security-sector reform, and the performance of statehood (African Union Commission, 2013; Knudsen, 2012; NATO, 1999; Organization of American States, 1995). Publication therefore need not express an intention to exceptionalize threats or expand coercive power.

Documents as institutional practices

We draw on three complementary traditions: Latour on documents as inscription devices (Latour, 1987, pp. 68–69); the Paris School on institutional securitization and bureaucratic production (Balzacq et al., 2016; Bigo, 2002); and Foucault on discursive formations (Foucault, 2002; Neal, 2010). Together, they explain how threats become nameable, durable, and repeatable within a transnational documentary field.

Because these documents are outputs of governmental bureaucracies, our approach builds on previous institutional, sociological, or "Paris School" securitization research (Balzacq et al., 2016; CASE Collective, 2006; Wæver, 2004). This research shows how bureaucratic dynamics, professional

authorities, and backstage processes contribute to securitization and "insecuritization" (Bigo, 2002, 2006). Threat construction may also occur through material practices, technological systems, and bureaucratic procedures, which Huysmans calls "little security nothings": routine practices that remain below the threshold of dramatic political declarations (Huysmans, 2011).

Balzacq develops this insight through the concept of a "regime of practices" (Balzacq, 2019). Securitization is not simply the product of an elite declaration and its acceptance by an audience, but "a set of interrelated practices, and the processes of their production, diffusion, and reception/translation that bring threats into being" (Balzacq, 2019, pp. 332–333). These practices combine linguistic, institutional, material, and technical elements in relatively stable forms of organization (Balzacq, 2019, p. 340). Strategy documents are products of such regimes: they translate assessments assembled across government into authorized accounts of the threats facing the state, while enabling those accounts to circulate and be translated into other strategic contexts.

Documents thus connect the linguistic focus of Copenhagen School securitization theory and the sociological focus of the Paris School. However, documents are forms of discourse unlike the speech acts issued by individual political actors because they are produced through complex sociological processes of internal contestation, negotiation, and agreement (Freeman & Maybin, 2011; Prior, 2008). They also follow genre conventions that shape what can be said and how (Neumann, 2007). Studies of US strategy document production show how inter-agency negotiations and institutional processes transform complex security assessments into formal texts (Doyle, 2007, pp. 625–626; Stolberg, 2012, pp. 71–72). Documentary securitization therefore reflects organizational practices and professional expertise rather than the intentions of a singular speaker (Bigo, 2002, 2008; Bueger & Gadinger, 2021).

Documents are not, however, neutral repositories of prior decisions. Following Latour, we treat them as inscription devices that make heterogeneous phenomena durable, mobile, and comparable (Latour, 1987, pp. 68–69). Like laboratory inscriptions, strategy documents place diverse problems within common classificatory forms (Latour, 1987, p. 223). Cyber operations, pandemics, migration, terrorism, economic coercion, and climate change can thereby appear together as comparable objects of national security. Inscription does not merely report threats; it gives them a stable form that can travel beyond the institutional setting in which it was produced.

This approach defines the scope of our claims. National security and defence documents do not capture every threat constructed in speeches, legislation, media debate, intelligence assessments, or other documentary

traditions. Nor does the absence of a published strategy document show that a state lacks its own security discourse. The NSDDD captures a specific transnational field of public, top-level strategy documents under defined criteria (Neal & Gardner, 2026). Although this genre might be expected to privilege military security, its substantive breadth is an empirical question, and one that differs by document type. By extracting explicit threat statements without preselecting topics, our analysis assesses whether the resulting repertoire remains military or extends across the broadened security agenda. The next section explains this procedure, while the subsequent analysis measures how the active repertoire widens over time.

Strategy writers as scriptors

The concept of a threat lexicon extends securitization theory's poststructuralist account of the grammar of security. The Copenhagen School treats this grammar as a quasi-transcendental condition of securitizing speech: it precedes particular speakers and supplies the intelligible form through which an issue can be articulated as a security threat. We argue that strategy writers draw not only on this pre-existing grammar, but also on a historically sedimented repertoire of already named threat objects with which that grammar can be populated. Securitizing writers therefore write under two related conditions: a grammar that makes security claims intelligible and a lexicon that makes particular threats readily available for use (setting aside the wider contextual felicity conditions of speech act securitization).

Our core theoretical move is to reconceptualize securitizing agency in strategy-document contexts. In Foucauldian terms, we can understand strategy documents as a "surface of emergence" on which threats become visible and nameable (Foucault, 2002, p. 41). They participate in a discursive formation whose rules of formation condition the objects that can appear as threats, the concepts that make threats comparable and policy-relevant, and the authorized positions and formats through which security can be written in the name of the state.

Foucault identifies the discursive conditions under which threats become nameable; Latour explains how inscription makes them portable; and the Paris School locates this work in institutional practices rather than isolated declarations. Unlike speech acts centred on particular speakers and audiences, strategy documents stabilize a repeatable repertoire through texts that circulate and can be reinscribed across cases.

The grammar–lexicon distinction clarifies what is at stake in documentary authorship. Securitizing agency in strategy writing is more constrained than speech-act accounts imply, particularly at the level of threat-object innovation. In top-level strategies, states rarely author threats wholly *ex nihilo*; they more often operate as scriptors who select, recombine, and reinscribe

threat categories drawn from an already available and internationally legible documentary repertoire (Barthes, 1977, p. 147). This does not eliminate agency: states still select, rank, localize, adapt, and omit threats. The concept of the scriptor instead relocates agency from sovereign invention (and its nominalist Schmittian links (Wæver, 2011, p. 470)) towards the strategic use of inherited categories. Beyond carrying securitizing moves, the documentary field conditions what can count as a threat object in strategy writing.

Documentary inscription should not, however, be equated with successful Copenhagen-style securitization. A strategy may identify a threat without invoking emergency, demanding exceptional measures, or securing the assent of a discrete audience, and, compared to the relative immediacy of a political speech for example, the slower bureaucratic and periodic production pace of the genre does not invite these aspects. Here, documentary securitization refers more narrowly to the institutional inscription and reproduction of issues as security threats. Inscription can make a category available for comparison, policy coordination, resource allocation, or later securitizing moves, but in this analysis we do not infer those consequences or audience acceptance from textual appearance alone.

We therefore distinguish the theoretical process of threat construction from its observable documentary trace. In our terminology, a first inscription is a threat topic's earliest observed appearance in this corpus, dated by year; a first inscriber is a state or territory associated with that appearance in the first-inscription year; adoption denotes a later first use by another country, while reinscription denotes any subsequent use, including by a first inscriber. These terms do not imply that a state originally constructed the threat in wider political intersubjectivity, that later states borrowed directly from its document, or that the threat was successfully securitized. The next section operationalizes these distinctions.

Prior inscription power

These arguments yield a theory of prior inscription power. Recurring strategy documents create a reusable threat lexicon through practices of naming and classification. Early participants have more opportunities to inscribe categories before most states begin publishing. Temporal precedence alone constitutes a first-mover advantage. It becomes inscription power when early inscriptions acquire durable standing and circulation within a field structured by unequal institutional and geopolitical effect.

Diffusion need not depend on direct persuasion or a hegemon's conscious imposition of particular threats. Categories may circulate through alliance planning, transparency frameworks, security-sector reform, professional networks, or imitation of an established strategic genre. Once normalized, they

offer later writers recognizable ways to articulate security concerns. States can still reinterpret, combine, or reject them, but inherited categories are easier to deploy in internationally legible form than categories that have yet to be inscribed.

This shifts agency from discrete acts of construction toward patterned practices of reinscription. The key form of power in this setting is not only persuasion in the moment of a securitizing move, but positional access to inscription opportunities that populate the repertoire from which later strategy authors draw. Inscription power thus links the production, diffusion, and reception of threats within Balzacq's regime of practices (Balzacq, 2019, pp. 332–333). It also produces constrained authorship: most writers enter a field in which much of the available threat lexicon has already been written.

Our argument has three testable implications. First, if first inscription were a broadly shared capacity, credit for new threat topics would be distributed relatively evenly across countries and territories. If credit is instead concentrated, the capacity is unequally held, and the explanation must lie in differences of entry timing, documentary output, or structural position. We test this by measuring the distribution of first-inscription credit and then asking how far it tracks each of publication volume, year of entry into the field, and country characteristics such as income, regime type, and institutional membership.

Second, if these inscriptions constitute a shared lexicon, categories should travel beyond the documents and countries in which they first appear, and should return to use after periods of dormancy. A lexicon confined to its sites of origin, or one whose categories lapse permanently, would not support the claim. We test this through cross-country adoption, widening of the active repertoire, dormancy, and reactivation.

Third, if temporal precedence matters in its own right rather than as a by-product of hegemony, early participants' first inscriptions should be concentrated in the field's formative period and fall away as participation expands, while the categories they inscribed stay in circulation. US first inscription sustained across all three periods would instead suggest that hegemony, not timing, is doing the work. We test this through the timing of recurring strategy publication and the redistribution of first inscription during and after a 1998–2004 period of publication growth.

Together, these tests can support a bounded claim. If the predictions hold, they show that a few early participants, the United States above all, were able to place categories in the record before most countries were publishing at all, and that those categories stayed in use long after their inscribers' lead had gone. The tests cannot show that later writers borrowed from those documents, identify every channel through which categories travel, or establish that earlier documents cause later adoption. They can, however, establish

the pattern on which the claim rests: how unequally first inscription is distributed, how that distribution changes as participation widens, and how long inscribed categories stay in circulation.

Testing these claims requires moving beyond case-based analysis of individual securitization episodes toward systematic examination of explicit threat identification across the full documentary field. The next section describes how we identify threat topics and operationalize first inscription, adoption, concentration, and repertoire change across a multi-country, multi-decade corpus of strategy documents.

Data and methods

Identifying threats without preselecting cases

Securitization studies have predominantly relied on case studies of successful securitizations, leading to documented selection and confirmation bias (Baele & Sterck, 2015, p. 1122; Balzacq, 2011, p. 34). From the outset, Buzan, Wæver, and de Wilde suggested that studies would focus on successful cases because these reveal how security is constructed (Buzan et al., 1998, p. 39). This disciplinary prediction has largely held true, but this means researchers tend towards preselected topics, such as migration, terrorism, and health, that are already recognized as securitized, leaving the wider repertoire of threats without systematic, inductive comparison across cases (Baele & Jalea, 2022, p. 380; Balzacq et al., 2016, p. 507). Rather than imposing preselected threat topics, we extract and classify all explicit threat statements in the corpus.

A further problem is terminological delimitation. The corpus contains many apparent threat identifications that do not use the word "threat," and instead use terms such as "concern," "challenge," "risk," or "danger." When "threat" is used, it is sometimes qualified in a variety of ways, such as "serious", "major," or "significant." Any broader search must therefore decide where the language of threat construction begins and ends. Baele and Sterck construct a 226-word "security lexicon" by aggregating two existing reference glossaries, the UN Glossary of Terms and Concepts in Peace and Conflict Studies and the index of the Palgrave Handbook of Security Studies (Baele & Sterck, 2015, pp. 1127–1128, 1135–1137). While they set out clear editorial principles for this list, its categories are drawn from prior security studies scholarship rather than induced from the texts under analysis, so the method risks measuring the extent to which a corpus reproduces an already-settled disciplinary vocabulary rather than detecting threat construction on terms internal to the documents themselves. Schäfer et al. take a more empirically grounded approach, refining an 18-word list iteratively against their corpus through repeated Keyword-in-Context checks

(Schäfer et al., 2016, pp. 84–85, 91). Their vocabulary is lexically broader than the single substring we search on, but it is built for one issue area, climate change, rather than the full range of threats a strategy document might name, and their justification is empirical validation against a corpus rather than a claim to have identified the natural boundaries of threat language as such.

Our solution is conservative standardization: we extract sentences containing the substring "threat," yielding explicit threat statements that can be compared across documents and countries. This operational decision provides the methodological consistency necessary for meaningful quantitative comparison (Mouton & Marais, 1988, pp. 70–71). It likely misses securitizing language that avoids the word "threat," but it reduces fragile judgements about what should count as a threat claim and avoids the pitfalls of under-justified security lexicons. The analysis therefore concerns explicit documentary threat identification, not every formulation that might perform a securitizing function.

A third problem is categorization granularity. Is "extremist terrorism" the same as "domestic terrorism" or "Islamic terrorism"? How granular should threat categories be? There is no uniquely correct answer to this question; it is a matter of analytical purpose (Estivill-Castro, 2002; Milligan & Cooper, 1987). We therefore group the labels by similarity of wording at several settings, from strict, which keeps close variants apart, to loose, which merges them, and check that the findings hold across all of them. Granularity then becomes something the results are tested against rather than a choice fixed at the outset. This approach also enables us to revisit the familiar claim that security has "widened" (Buzan & Hansen, 2009): rather than assuming widening, we can measure how the active repertoire of threat topics changes over time and subsequently assess its relationship to the broadened security agenda.

Corpus and unit of analysis

The National Security and Defence Documents Dataset includes National Security Strategies as well as closely related top-level publications, including defence strategies, defence white papers, and national threat assessments. Its published inclusion criteria are for documents that are horizontally comprehensive, covering multiple defence and security domains, rather than vertically specialized (Neal & Gardner, 2026). The period begins with the first eligible document in 1987 and ends with the dataset's last complete year.

The 118 countries and territories constitute the known population of publishers in the 249-member ISO country code list, not a sample. States without an included document may articulate threats in speeches, legislation,

classified assessments, specialized policies, or other documentary genres. Our claims are therefore limited to this transnational field of public, top-level strategy documents.

The corpus supplies all non-English documents in English translation (131 of 660), machine-translated by the dataset compilers, so the same threat substring rule can be applied uniformly across languages. Machine translation has been validated for computational text analysis work similar to this: de Vries et al. (2018) find that word-frequency analyses of machine-translated political texts return substantively the same results as analyses of human translations. Translation nonetheless comes at some cost to lexical variation, particularly where official English versions standardize national terminology. Because the text encoder used here operates on English, working from the original-language texts would require both a multilingual encoder and a version of the dataset that retains them.

The method operates at three levels. Sentences are labelled and category strings are clustered into topics; clusters must account for at least five sentences to enter the inventory. For the subsequent analysis, these sentence-level occurrences are collapsed to binary country–topic–year mentions, indicating whether a country mentions a topic at least once in a year. Repeated mentions within a document, across several documents issued by one country in the same year, or through several labels mapping to the same topic therefore contribute only one country–topic–year observation. Counting mentions rather than presence would measure documentary volume and drafting convention as much as threat attention: a country that publishes longer or more frequent documents registers more mentions of the same topic, and repetition within a single document may reflect section structure or standard phrasing rather than added emphasis. Binary coding makes countries with very unequal publication records comparable, at the cost of any measure of how insistently a threat is pressed.

Constructing the topic inventory

We construct the threat-topic inventory in four steps: extraction, categorization, filtering, and clustering.

Step 1, extraction. A fuzzy keyword search returned 23,630 candidate sentences. Step 2, categorization. All were passed to the categorization model. Each sentence was assigned a short threat-category label such as “terrorism,” “cyber attacks,” or “territorial invasion” using Gemini 2.5 Flash through the Google Batch API on 11 January 2026. A fixed prompt asked the model to identify the threat in three to seven words or return “unclear” where none could be identified. Categories were not predetermined, and each sentence was classified without surrounding document context. No system prompt or generation parameters were supplied, so the contemporaneous API

defaults applied. The prompt, inputs, and returned labels are provided in the replication materials (Neal, 2026); no labels were manually edited.

Step 3, filtering. A case-insensitive literal substring check was then applied to each sentence's text, retaining the 18,592 sentences that contain "threat." (Because each sentence was classified in isolation, filtering the text categorisation after rather than before it does not affect the labels of the retained sentences.) Of these retained sentences, the model returned "unclear" for 5,207 (28.0 percent), leaving 13,385 labelled sentences.

Step 4, clustering. We embedded the resulting category strings using all-mpnet-base-v2 and grouped them through hierarchical agglomerative clustering with average linkage on cosine distance. Each topic was named using its medoid category string. We computed solutions at six similarity thresholds from 0.40 to 0.90 and selected 0.80 by inspecting the topic inventories for whether clusters merged threat types that remained analytically distinct. We preferred higher resolution, accepting some redundancy among near-synonymous topics, because distinctions lost through merging cannot be recovered in subsequent analysis. Selection was based on the topic inventories rather than the country rankings or temporal findings, which we test across all six thresholds.

Clusters represented by fewer than five sentences corpus-wide were then excluded, a final filter applied after clustering. At 0.80, the resulting inventory contains 421 topics covering 10,478 sentences. The five-sentence floor focuses the analysis on recurrent categories but makes it conservative towards rare and newly emerging threats. Because this floor is based on sentences rather than country–topic–year mentions, a topic cluster could in principle qualify through repetition within one document. In practice, this occurs for only one of the 421 topics. Sections A, B and I of the online appendix report the complete pipeline, threshold comparisons, and all 421 topics.

Validation

Because the topic inventory underpins every subsequent result, we validated separately the identification, labelling, and clustering stages using random samples coded by an author and an independent second coder. Full sampling procedures, coding instructions, confidence intervals, and error analyses are reported in Sections E and F of the online appendix.

In a stratified sample of 200 retained and excluded sentences, agreement between the model and the first human coder over whether a sentence identified a specific threat was 86.0 percent, with Cohen's $\kappa = 0.720$. Human–human agreement on a subsample was closely comparable at $\kappa = 0.714$; among cases on which both humans agreed, model agreement was $\kappa = 0.754$. In this sample, the model applied the exclusion criterion about as consistently as the second human coder.

In an anchored assessment of 100 retained sentences, 94 percent of model labels were judged accurate and a further 5 percent partly accurate. Because the coder saw the assigned label, this establishes that the labels are defensible summaries rather than that an independent coder would have generated identical wording. Labels were less reliable when a sentence named several threats, because the model assigned only one dominant category. Co-occurring threats are therefore undercounted.

Topic coherence was assessed through the word intrusion method (Chang et al., 2009). An intruder was identified in 39 of 40 sampled topics, compared with a chance probability of one in six. We also examined 60 sentences removed by the five-sentence floor. Most were genuine but infrequent threat identifications, and roughly one third contained material not represented in the retained inventory. The floor therefore excludes substantive rare content rather than noise alone. Exclusion rates are stable over time, making it unlikely that the floor produces the observed widening, but they are modestly higher among low-volume countries. To that extent, measured concentration may be mildly overstated.

Relating topics to the broadened security agenda

The 421 topics are generated inductively rather than assigned to predetermined security sectors. To assess whether national security and defence documents privilege military security, we separately classify the topics using Buzan's military, political, economic, societal, and environmental sectors, a supplemental cyber category following Hansen and Nissenbaum, and a category for undetermined (Buzan, 1991; Hansen & Nissenbaum, 2009). This coding does not alter the topic inventory or any first-inscription result.

Sector assignment is an interpretive judgement, so we validated it against independent human coding. A stratified sample of 92 topics was coded by hand under the same written definitions supplied to the model, blind to its assignment, by two coders working independently. Agreement between the model and the human coders (Cohen's $\kappa = 0.673$ and 0.665; $\kappa = 0.709$ pooled across 114 topics) was slightly higher than agreement between the two human coders themselves ($\kappa = 0.628$, 95 percent CI 0.497–0.760), and the intervals overlap throughout. Sector assignment is therefore reproducible to about the limit that independent human coders reach on the same material.

The disagreement is concentrated at a single boundary. Excluding topics that either coder placed in the undetermined category, human–human agreement rises to $\kappa = 0.832$ on the remaining 72 topics. The coders agree substantially about which sector a threat belongs to; what they disagree about is whether labels such as "hybrid threats" or "geopolitical threats" can be assigned a sector at all. The undetermined share is therefore the least stable quantity reported here. We do not report a cross-sectoral flag:

the two coders applied it to 5 and 19 percent of topics against the model's 38 percent, which measures a coder-specific threshold rather than a property of the topics.

Definitions and measurement

A topic is a semantically coherent cluster of threat-category labels. Its first-inscription year is the earliest year in which it appears in the corpus. A first inscriber is a country or territory mentioning it in that year. An adopter is a different country whose first mention occurs later. Reinscription refers to any subsequent use, including by an original first inscriber.

The retained inventory contains 4,041 distinct country–topic pairs: 463 first-inscriber pairs and 3,578 adopter pairs. Each pair is counted once irrespective of how many years the topic recurs in that country. Topics may have several first inscribers when they appear in multiple countries in their earliest observed year. Such countries are treated as joint first inscribers, not adopters.

When a topic has k first inscribers, each receives $1/k$ inscription credit. Sole-credit and full-credit alternatives are reported as robustness checks; designation as a first inscriber or adopter is unchanged.

These measures describe positions within the documentary corpus. First inscription identifies where a topic enters this record; threat construction denotes the wider social process through which it becomes intelligible and accepted as a security concern.

Country characteristics and periodization

Regime type, income, and institutional membership change over time. We therefore code them across each country's document-publication years rather than applying present-day classifications retrospectively. Regime type comes from Freedom House, income from the World Bank, and institutional affiliations from historical membership records. Modal ties are assigned to the chronologically earlier status; institutional membership requires membership in at least half of a country's publication years. Alternative coding rules produce substantively similar results; full rules and coverage are reported in Section G of the online appendix.

Field expansion occurred over a transition rather than at a single cut-point. We distinguish a formative phase (1987–1997), a transition (1998–2004), and an expanded phase (2005–2025). By the end of 1997, 13 countries and territories had entered the corpus; cumulative participation rose to 62 by 2004 and reached 118 by 2025. These phases follow the observed expansion of documentary participation.

Scope and limitations

The analysis maps the documentary inscription and circulation of explicit threat categories; it does not assess audience reception or downstream policy effects (Buzan et al., 1998, p. 26). The corpus is the empirical "surface of emergence" on which these categories become observable as documentary objects (Foucault, 2002, p. 87).

A key limitation is that the topic inventory is constructed from an unequal pool of threat sentences. Countries with more documents contribute more labelled sentences, potentially shaping which category variants survive filtering and how variants cluster. Sentence contribution is associated with firstinscription credit (Spearman's $\rho = 0.746$[1]), indicating a plausible pathway by which documentary throughput influences the constructed topic space.

The corpus captures only public, top-level documents and explicit uses of "threat." The extraction rule also operates unevenly across genres: National Security Strategies use this target substring of "threat" 88 percent more densely per 1,000 words ($p < 0.0001$) and contain a broader sectoral distribution than defence-oriented texts. They therefore contribute more heavily to the extracted inventory relative to their textual volume, while the numerical predominance of defence documents pulls the aggregate towards military security. Translation and one-label-per-sentence classification may further reduce lexical variation and undercount co-occurring threats.

The analysis maps recurrence rather than transmission; tracing how particular categories travel would require evidence from citations, drafting networks, alliance planning, or template circulation.

Unequal first inscription

As noted in the previous section, the main clustering specification (0.80 threshold) yields a retained inventory of 421 distinct threat topics (all figures use the principal 0.80 clustering specification). Across these topics, the corpus contains 4,041 distinct country–topic pairs. Of these, 3,578, or 88.5 percent, are adoptions: the country first mentions the topic only after it has appeared elsewhere in the documentary record. Threat writing in this field therefore consists overwhelmingly of selecting from categories already present in the lexicon rather than introducing categories first.

First inscription is also highly concentrated. Sixty-three of the 118 countries and territories receive some first-inscription credit; 55 receive none. The latter group includes countries that participate through adoption and four whose documents yield no retained country–topic pair. The absence of first-inscription credit therefore does not imply an absence from security discourse. Rather, most of these countries participate in a documentary field whose vocabulary was partly established before they entered it.

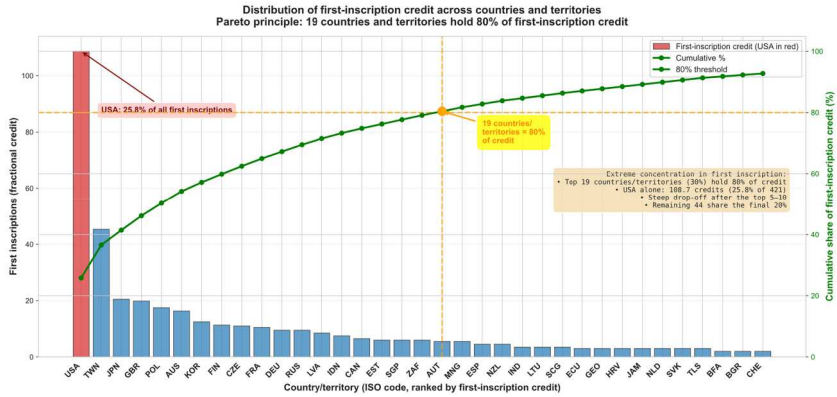


Figure 1. "Distribution of first-inscription credit across countries and territories".

The distribution of first-inscription credit shows a steep drop after the leading countries (Figure 1). The Gini coefficient across all 118 countries and territories is 0.822, indicating an extremely high level of inequality, with first inscriptions heavily concentrated among a small fraction of the population. The top five account for 50.3 percent of first-inscription credit, nineteen account for 80.4 percent, and the bottom half of first inscribers collectively account for 9.4 percent. The United States alone receives 108.7 of the 421 fractional credits, or 25.8 percent.

Publication volume and first-inscription opportunity

First-inscription concentration might be a by-product of documentary output: countries that publish more have more opportunities to register a first appearance. Output and credit are indeed associated (Figure 2), but not in the way that objection requires. Countries that write more do tend to inscribe more: rank the 118 countries and territories by words contributed

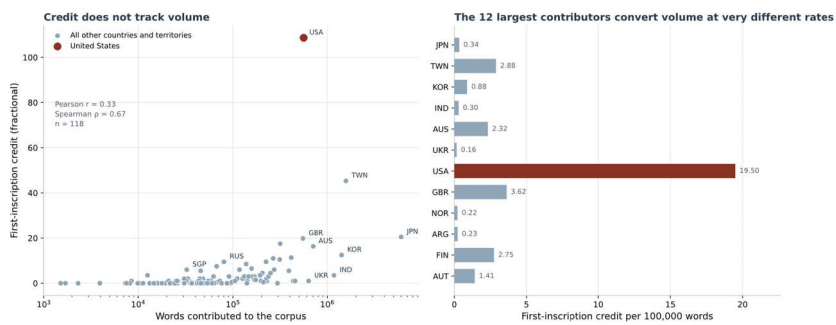


Figure 2. "First-inscription credit and corpus volume".

and by credit received, and the two orders broadly agree ($\rho = 0.67$). What does not hold is proportionality. Writing ten times as much does not earn ten times the credit, and once the comparison is made in amounts rather than ranks the relationship is weak ($r = 0.33$). Japan contributes 6.0 million words to the corpus against the United States' 0.56 million, yet receives 20.5 fractional credits to the United States' 108.7. Per 100,000 words the United States inscribes at fifty-seven times Japan's rate, 19.50 against 0.34. Among the twelve largest contributors the rate spans 0.16–19.50. Volume of any kind, whether documents, pages, or words, creates opportunities for first appearance without converting into them at anything like a constant rate. Because credit and output come apart in this way, the concentration shown in Figure 1 reflects how unevenly countries convert documentary output into first inscription, while volume of output explains relatively little of it.

Robustness and substantive composition

The concentration of first inscriptions and the leading formative US position persist across clustering thresholds (online appendix, Section B). Alternative rules for allocating credit among joint first inscribers produce similar results, as expected given that most topics have only one first inscriber. The concentration result is therefore not an artefact of a single level of topic granularity.

The inventory has a substantial military component, as expected from a corpus that includes defence strategies and white papers, but it is not confined to military security. Military topics account for 56.1 percent of the inventory and 60.6 percent of country–topic–year mentions; political for 13.8 and 10.6 percent; environmental for 5.7 and 5.1 percent; cyber for 4.5 and 6.6 percent; economic for 2.9 and 2.1 percent; and societal for 2.6 and 1.0 percent. Prominent examples include "Islamist terrorism" and "military aggression" in the military sector, "organised crime activity" and "regional geopolitical instability" in the political, "cyber attacks" in the cyber, and "climate change impacts" and "pandemic disease outbreak" in the environmental. A further 14.5 percent of topics could not be assigned a sector from their labels.

The military emphasis is concentrated by genre rather than spread across it (Figure 3). Military threats account for 50.6 percent of retained threat sentences in national security strategies, against 72.5 percent in defence documents and 71.2 percent in defence white papers. The difference is taken up by the political (14.7 percent in strategies), cyber (9.0 percent) and environmental (7.7 percent) sectors rather than distributed evenly across the remainder. The aggregate military skew therefore reflects the mix of genres represented in the corpus and differences in their substantive

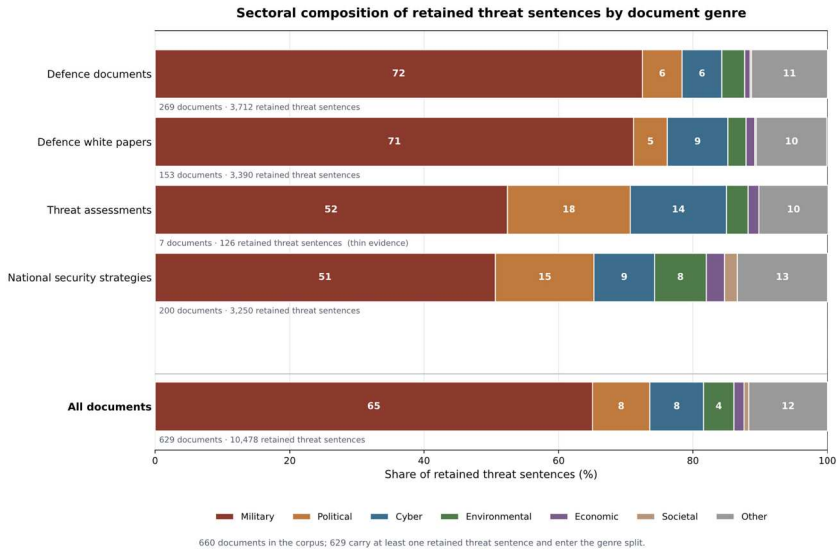


Figure 3. "Sectoral composition of retained threat sentences by document genre".

composition; an extraction-density effect discussed below also shapes their relative contribution.

This comparison should be read alongside unequal extraction density across genres: National Security Strategies yield more threat sentences per 1,000 words, so their broader agenda carries more weight in the inventory than their document count alone suggests, while the larger number of defence-oriented documents pulls the aggregate the other way (see *Scope and Limitations*).

From formative dominance to dispersed first inscription

The full-period distribution conceals a pronounced historical transition. The United States' leading position is concentrated in the formative phase of the documentary field and declines as participation expands. We distinguish three periods: a formative field (1987–1997), a transition (1998–2004), and an expanded field (2005–2025). This periodization follows cumulative entry into strategy-document publication (Figure 4).

By the end of the formative period, 13 countries and territories had entered the corpus. The United States receives 90.8 of 175 first-inscription credits, or 51.9 percent. During the transition, cumulative participation rises to 62 and the US share falls to 15.5 of 139 credits, or 11.2 percent. In the expanded period, it falls to 2.3 of 107 credits, or 2.2 percent. The transition is itself highly productive, accounting for nearly one third of the 421-topic inventory. Field expansion therefore disperses first inscription rather than bringing documentary novelty to an end (Figure 5). The country with the largest share also

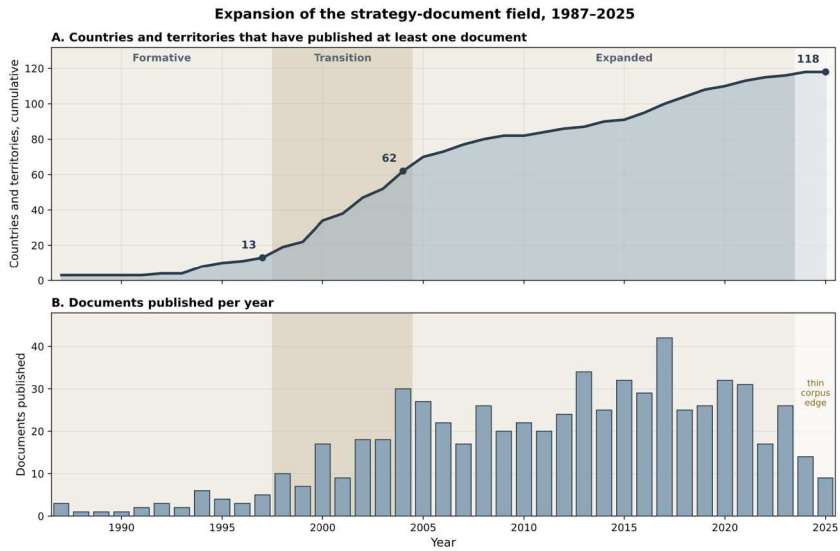


Figure 4. "Expansion of the strategy-document field, 1987–2025".

changes: the United States ranks first in the formative and transition periods, while Japan ranks first in the expanded field. Japan's position is stable across clustering thresholds from 0.50 to 0.90, but its 19.2 percent share at the principal specification reflects a dispersed distribution rather than a new concentration comparable to the formative US pattern.

Early strategy publication and inscription opportunity

Document-type evidence identifies a plausible mechanism for this temporal pattern. In the sparsely populated formative field, the United States invents

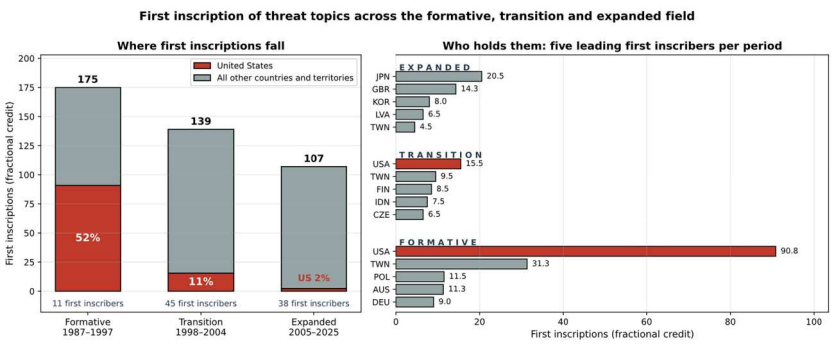


Figure 5. "First inscription across the formative (1987–1997), transition (1998–2004), and expanded (2005–2025) field".

the National Security Strategy as a distinctive genre: a recurring broad-government strategy coordinated from the centre of government rather than produced by the defence department. Other early participants publish mainly defence strategies and white papers (Figure 6). The US concentration in first inscription therefore coincides with an unusual combination of genre formation, early entry, recurring publication, and geopolitical authority.

The claim is not that the NSS label itself causes later adoption. Taiwan and Australia also enter early, but principally through defence-oriented formats, while later first inscriptions occur across several document types. The distinctive US position combines recurring publication in a sparsely populated field with a genre that uses explicit "threat" terminology more densely than defence-oriented documents. Both increase the likelihood that the United States is recorded as the first documentary site of a category under our extraction rule. Genre-specific extraction density is therefore part of the measured inscription advantage as well as a methodological limitation: the early US pattern cannot be attributed to temporal precedence or geopolitical authority alone (see *Scope and Limitations*).

These conditions establish an opportunity for prior inscription power, not the power itself. That depends on whether categories acquire authority, circulate beyond their first site, and persist through adoption and reactivation. Prior inscription power therefore does not require direct copying or contemporaneous persuasion; it lies in the durable standing acquired by categories first placed into an emerging repertoire.

Field expansion and alternative explanations

The documentary field expands through heterogeneous institutional channels. NATO accession encouraged the publication of national security concepts aligned with alliance planning (NATO, 1999); OAS initiatives

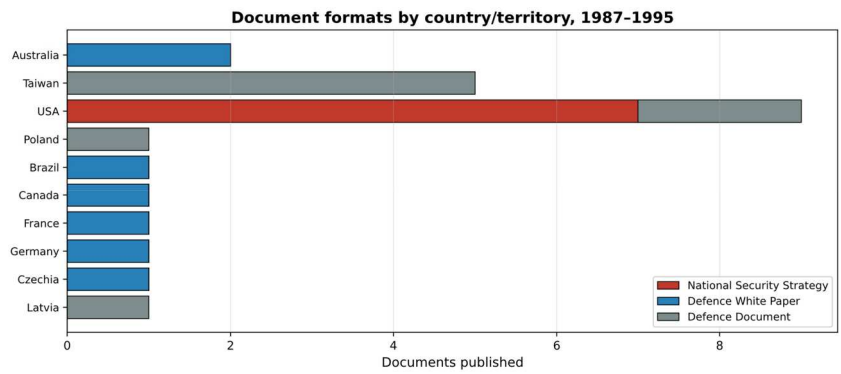


Figure 6. "Document formats by country/territory, 1987–1995".

promoted transparency and confidence-building (Organization of American States, 1995); AU frameworks encouraged security-sector strategies in post-conflict settings (African Union Commission, 2013); and Pacific regional frameworks created space for distinctive priorities such as climate security (Pacific Islands Forum, 2018). These initiatives served different purposes but collectively expanded public strategy production around the turn of the millennium. Beyond these groupings, many other countries and territories have adopted the publication practice; the range of reasons why needs further study.

Material events plainly also shape threat identification. The end of the Cold War, terrorism, cyber capabilities, hybrid warfare, aerial drones, and artificial intelligence all alter what governments confront and write about. Events alone, however, do not explain which country becomes the first documentary site of a category. New concerns continue to arise while first appearances disperse across an expanding field. The evidence therefore supports a bounded opportunity claim: early and recurring participation, document genre, and the density of explicit threat language help explain historically concentrated first inscription, but the analysis does not establish that earlier documents cause later adoption or that material developments are irrelevant.

Unequal redistribution and inscription opportunity

Once the field expands, first inscriptions are redistributed rather than equalized. Descriptively, first-inscriber status is concentrated among high-income countries, but it does not rise monotonically across income categories. It is held by 76.2 percent of high-income countries, compared with 32.1 percent of upper-middle-income, 45.5 percent of lower-middle-income, and 41.7 percent of low-income countries (Figure 7). The relevant divide is therefore between high-income countries and all others, not a general income gradient ($\chi^2 = 12.81$, $p = 0.0003$). First-

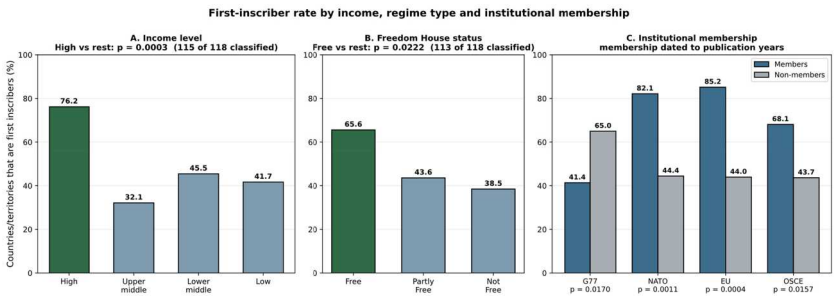


Figure 7. "First-inscriber rate by income, regime type and institutional membership (publication-year coding)".

inscriber status is also more common among countries classified as Free during their publication years: 65.6 percent, compared with 43.6 percent of Partly Free and 38.5 percent of Not Free countries ($\chi^2 = 5.23$, $p = 0.0222$).

NATO, EU, and OSCE members are likewise more likely to receive first-inscription credit, while G77 members are less likely to do so in unadjusted comparisons (41.4 against 65.0 percent). These memberships are nested rather than merely correlated: every NATO and every EU country-year in the corpus is also an OSCE country-year, so their separate effects are not identified and should not be interpreted as independent institutional contributions. These associations instead locate first inscription within an interconnected Euro-Atlantic structure of wealth, democratic governance, and institutional membership; they do not establish that any one characteristic causes documentary novelty.

The country-level tests ask whether a country is ever a first inscriber, which conflates entering the field with naming something new once inside. We therefore repeat the analysis at the level of the publishing year. Across the 585 country-years in which a country published at least one eligible document, we model whether it recorded a first inscription that year, controlling for the calendar year, the years since the country entered the field, the volume of text it published that year, and whether a national security strategy was among its documents, with standard errors clustered by country. Each structural position is fitted in a separate model (Figure 8). Income, regime type, and NATO, EU and OSCE membership do not predict the rate: every interval crosses one. G77 membership does, roughly halving the odds of an inscribing year (odds ratio 0.43, 95

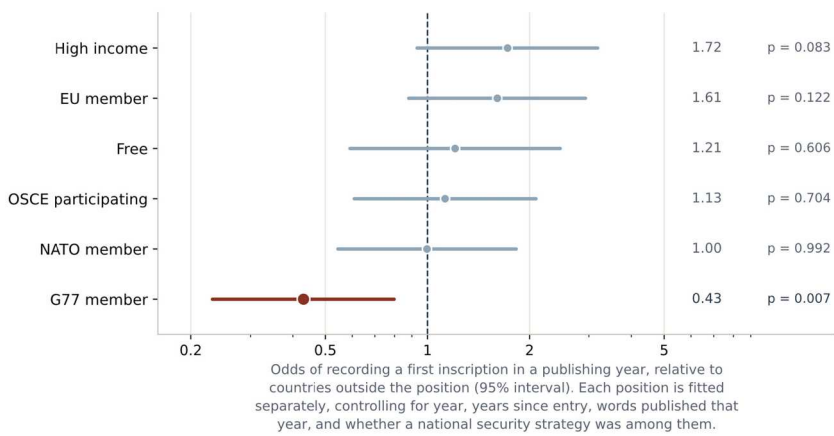


Figure 8. "Structural position and the rate of first inscription among countries already publishing".

percent CI 0.23–0.80, $p = 0.007$). In observable terms, 56 percent of G77 publishing country-years yield a first inscription during the transition against 69 percent of non-member years, and 10 percent against 20 percent in the expanded period.

The difference lies in whether a country has an inscribing year at all, not in how much it inscribes when it does. A model of the credit earned returns no difference (rate ratio 0.73, $p = 0.19$). G77 members have fewer productive years rather than less productive ones.

Entry timing and throughput

These associations should be interpreted in light of unequal opportunity to register first appearances, especially earlier entry into the strategy-document field and higher rates of ongoing publication within the corpus. G77 members typically enter later than non-members (median first-document year 2007 versus 2002) and publish fewer documents (median documents 2 versus 7), yielding substantially fewer opportunities to register a first appearance. Earlier entrants consequently accumulate significantly more first-inscription credit (Figure 9).

Entry timing and documentary output account for the G77 gap in whether a country is ever recorded as a first inscriber. They do not account for the rate at which G77 members inscribe once they are publishing, which remains about half that of non-members (Figure 8).

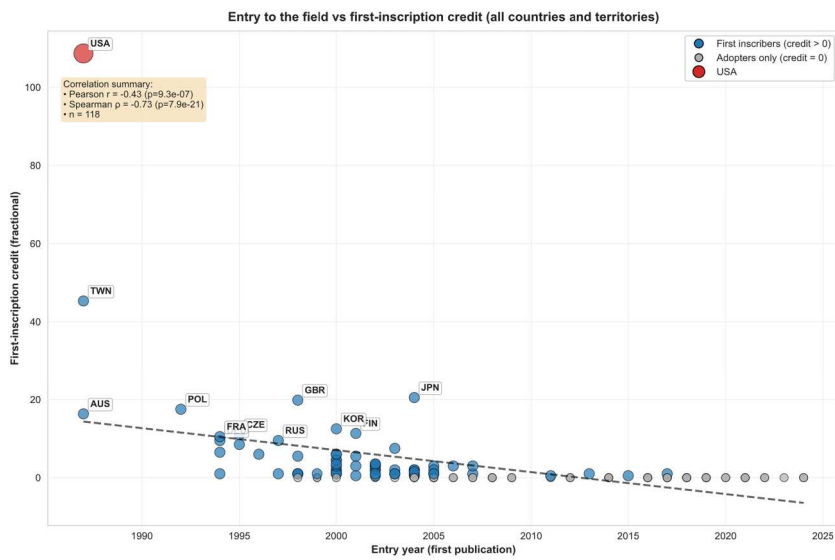


Figure 9. "Entry year and first-inscription credit".

A widening repertoire with long memory

Concentration and temporal precedence matter only if inscribed categories persist. We therefore turn from first appearances to changes in the active repertoire and the reactivation of previously used topics.

The active repertoire widens

We measure widening using the number of active topics per year: the count of distinct topics mentioned by at least one country in a given year. This is intentionally non-cumulative. A topic contributes to the annual count only if it is mentioned in that year, allowing the active repertoire to expand, contract, or fluctuate rather than assuming monotonic growth.

At the main clustering specification of 0.80, and across alternative thresholds, annual topic activity is substantially higher in the 2010s–2020s than in the late 1980s–1990s (Figure 10). This indicates that the documentary field sustains a wider set of threat topics in annual circulation once it matures, even though the number of active topics varies from year to year with publication cycles and shifts in strategic attention.

Widening describes the growing range of threat categories in annual documentary circulation; it does not establish whether real-world dangers have increased. How material developments become constructed as threats remains a central question for securitization theory.

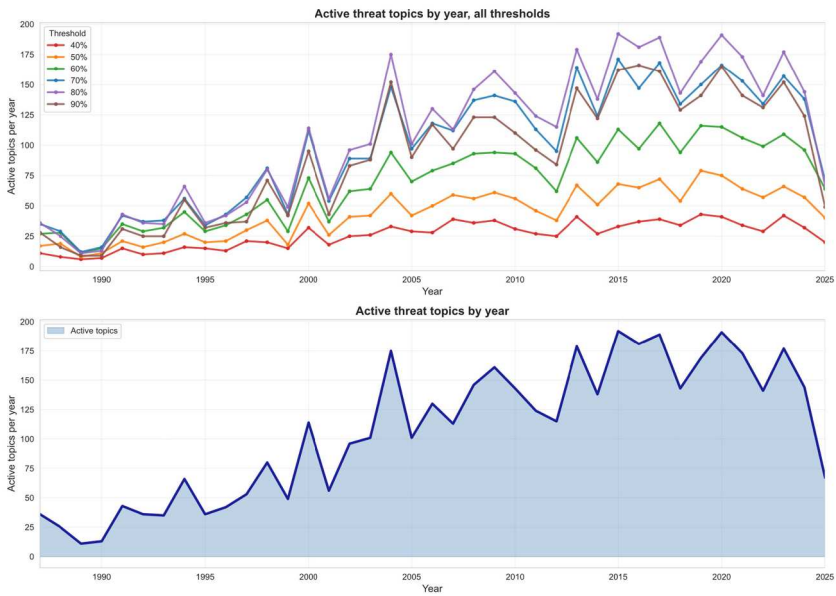


Figure 10. "Active threat topics by year".

Widening beyond early US inscription

Because the United States is the dominant historical site of first inscription in the documentary record, the widening observed might be driven largely by the continued circulation of topics it first inscribed. A related possibility is that widening merely reflects the accumulation of categories already present during the formative period. We assess these possibilities by recalculating the active repertoire under three restrictions: excluding the 113 topics first inscribed by the United States, alone or jointly; considering only the 139 topics first inscribed during the transition (1998–2004); and considering only the 107 topics first inscribed in the expanded field (2005–2025).

Widening persists under all three restrictions (Figure 11). Excluding topics first inscribed by the US lowers the level of annual activity but leaves the temporal expansion intact. Topics first inscribed during the transition form a widening active repertoire of their own, while topics first inscribed in the expanded field likewise accumulate and circulate after 2005. Widening is therefore neither solely a legacy of early US inscription nor an artefact of categories already present in the formative field.

The decline across all four series after 2023 coincides with thinning corpus coverage at the collection boundary: the corpus contains only fourteen documents in 2024 and nine in 2025, compared with around thirty in a typical recent year. It should therefore not be interpreted as evidence that the threat repertoire is contracting.

Persistence and reactivation

Widening through accumulation and reactivation depends on whether threat categories persist in the documentary repertoire, either through

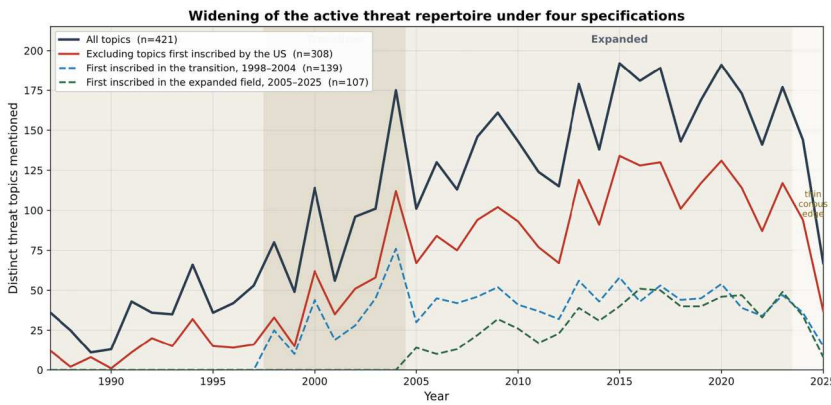


Figure 11. "Widening of the active threat repertoire under four specifications".

continuous mention or as latent categories that can be revived. We therefore trace the status of every topic from its first inscription onwards, allowing topics to move repeatedly between active, dormant and inactive states.

A topic is active when mentioned; dormant for one to five years or six to ten years since its last mention; and inactive when more than ten years have passed. Topics do not enter the calculation before their first inscription. These classifications refer to observed mentions, not to the formal currency or continuing policy influence of the documents. A strategy and its threat categories may remain operative beyond its publication year, often until superseded, but its practical authority may also change with a new government, altered circumstances, institutional revision, or other political developments. Because the corpus does not provide a comparable expiry date for every document, we use annual mentions of topics as the consistent observable measure. "Dormant" and "inactive" therefore mean absent from new documents in those years, not absent from policy or political discourse generally. This time-varying classification models the repertoire as a dynamic archive in which topics accumulate, recede, and re-enter documentary circulation as strategic attention shifts (Figure 12).

Three patterns stand out. First, the repertoire expands cumulatively: the number of topics ever used grows from 36 in 1987 to 421 by 2025. Second, only a minority of the accumulated repertoire is active in any given year: across the period, the mean active-to-ever-inscribed ratio is 37.9 percent, indicating that widening operates largely through layering rather than through the continuous use of every category. Third, persistence is typically episodic rather than continuous. We observe 1,066 reactivations, defined as a new mention after at least two full years without one. The median gap between mentions is four years and the longest 27,

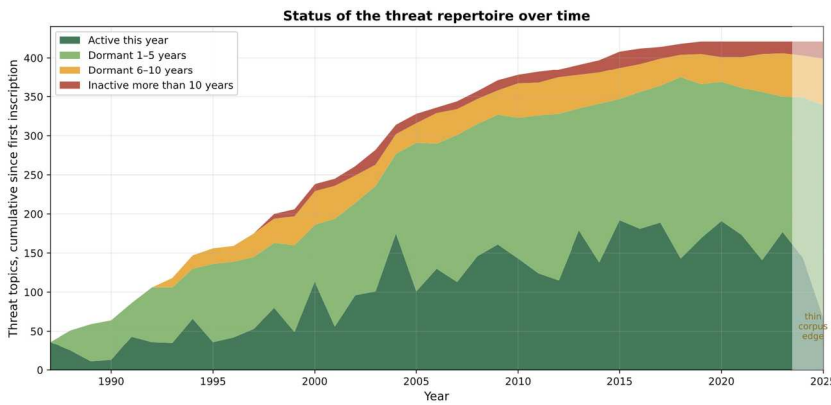


Figure 12. "Status of the threat repertoire over time: active, dormant, and inactive for more than ten years, stacked area".

corresponding to three and 26 full years of absence. Of these reactivations, 102, across 95 topics, occurred more than ten years after the topic's previous mention, while 712 occurred at least ten years after its first inscription, whether or not the topic had been mentioned in the interim.

Together, these results indicate a repertoire with long but episodic memory: categories recede from annual use but can return long after their first inscription and after extended periods of inactivity. This allows early inscription to have effects beyond the conditions under which it occurred: a category need not remain continuously prominent to retain a place in the documentary repertoire. Widening therefore occurs through accumulation, dormancy, and reactivation rather than straightforward replacement.

The apparent rise in dormancy and inactivity after 2023 coincides with thinning corpus coverage at the collection boundary and should not be interpreted as evidence of repertoire contraction or disappearance, since only 82 of the 421 topics had gone unmentioned for six years or more by 2025.

The findings thus establish the three empirical components of prior inscription power. First appearances are unequally distributed; their distribution changes as the field expands; and the resulting categories persist through layering and reactivation. The next section considers what this pattern implies for securitizing agency and documentary authorship.

Discussion

Documentary securitization and constrained authorship

These findings contribute to debates about securitization by specifying a mode of security construction operating through what Balzacq terms "regimes of practices" (Balzacq, 2019, pp. 332–333). Rather than displacing the Copenhagen School's attention to speech acts, the analysis clarifies how security is constructed differently when threats are articulated in documentary form within an increasingly public, internationally comparable field.

A Foucauldian framing shifts attention from discrete performative moments to the conditions under which certain threat objects become sayable, repeatable, and governable. Strategy documents can be treated as part of a discursive formation whose "rules" condition which objects become legible as threats, which vocabularies are available for their description, and which authoritative formats are available for reiteration across cases (Foucault, 2002; Neal, 2010). Latour adds that documentary inscription stabilizes those objects, while Barthes clarifies the constrained authorship of those who write with them.

Of 4,041 country–topic pairs, 88.5 percent are adoptions of categories already present in the documentary record, and nearly half of the countries

and territories receive no first-inscription credit. Securitizing agency in this documentary setting therefore lies in selecting from an inherited lexicon more than in authoring novel threat objects.

Prior inscription as temporal and structural power

Our central explanatory move is to treat top-level national security and defence strategy documents as inscription devices that stabilize threat categories into durable, reusable documentary objects. The United States' early dominance in documentary first inscriptions is associated with early institutionalization: it stabilized a recurring central government driven National Security Strategy template early, when few other states were producing comparable publications. Its leading full-period position is overwhelmingly historical: the United States receives 51.9 percent of first-inscription credit in the formative period, 11.2 percent during the transition, and 2.2 percent in the expanded field.

Yet timing alone is insufficient to explain this pattern. Other countries entered the field early without acquiring a comparable share of first inscriptions. This occurred near the post-Cold War height of US hegemony and alongside the expansion of NATO and other US-centred security institutions, giving US strategy documents exceptional geopolitical authority and potential circulatory reach.

This shifts the account of securitization away from contemporaneous persuasion under felicity conditions. Here, power operates through the capacity to convert early inscription into durable and circulating categories, not necessarily through direct persuasion, audience acceptance, or downstream policy effects. Once inscribed, topics need not be continuously promoted to remain consequential as part of an inherited lexicon. The frequent reactivation of dormant categories shows why this historically accumulated advantage can persist.

This distinguishes prior inscription power from first-mover advantage. First-mover advantage concerns the number of opportunities to appear first; inscription power concerns the unequal capacity to turn those opportunities into categories that acquire durable standing within a transnational repertoire. The US position in the formative field joined both. Its advantage was temporal, but it was also embedded in wider hierarchies of institutional resources and geopolitical authority.

The findings give empirical force to the grammar–lexicon distinction developed above. Securitizing actors draw not only on pre-existing rules for speaking security, but also on a historically sedimented repertoire of threat objects. This is where the distinction between authors and scriptors becomes consequential: actors author particular strategies, but rarely originate the categories with which they write. Prior inscription power is

therefore power over the historically constituted contents available to securitizing practice, not only over individual securitizing moves.

This interpretation is deliberately bounded. The evidence does not demonstrate that later writers directly copied US documents or were persuaded by their authors. Nor does first appearance prove original construction beyond this corpus. What the analysis establishes directly is the conjunction of early recurring publication, genre-specific opportunities for explicit threat identification, highly concentrated first inscription, widespread later adoption, and long documentary persistence. The early US advantage measured here therefore combines historical position with the genre-dependent measurement opportunity noted under Scope and limitations. This does not negate prior inscription power, but it means that the observed US concentration should not be attributed solely to timing, hegemony, or institutional reach.

What the historical transition rules out

The near-total decline in US first inscription across the turn-of-the-millennium transition is what distinguishes the early-institutionalization account from continuous dominance. Early inscription predicts historically front-loaded novelty contributions, followed by dispersion once the field expands and the opportunity structure for first appearances widens. That prediction matches the documentary record. The United States remains pivotal in security politics, but its distinctive position as a site of first inscription in this documentary record is overwhelmingly a feature of the formative period.

The same transition supports a field-expansion understanding of diffusion. Uptake in this setting is plausibly driven by the growth of documentary participation, and by the multiplication of internationally legible strategy documents across countries, rather than by immediate imitation alone. It also shows that widening is not simply a legacy of formative US inscription: the repertoire continues to widen when topics first inscribed by the United States are excluded and among topics first inscribed during the transition and expanded phases.

Inequality of opportunity

Field expansion does not produce levelling. During the transition and expanded phases, first inscriptions are redistributed across a wider set of countries, but first-inscriber status remains unequally distributed and is strongly conditioned by differential opportunity to register first appearances, especially earlier entry into the field and higher documentary throughput. For that reason, institutional and bloc-level patterns should be interpreted in light of opportunity for first inscription, not as direct measures of intrinsic innovativeness.

Descriptively, first-inscriber status is more common among high-income states and democracies, and is concentrated among countries embedded in NATO, EU, and OSCE-linked documentary infrastructures. These patterns partly reflect unequal opportunities for first inscription: some countries entered the documentary field earlier and published more frequently. Their concentration among wealthy democracies and Euro-Atlantic institutions also suggests that first inscription is embedded in wider geopolitical hierarchies. Those hierarchies operate on entry and output rather than on productivity once inside. Among countries already publishing, neither income, nor regime type, nor NATO, EU or OSCE membership predicts whether a given year yields a first inscription; only G77 membership does, and this is net of entry timing and publication volume, which the model holds constant.

G77 members generally entered the documentary field later and published fewer documents, and those differences account for the gap in whether they are ever recorded as first inscribers. They do not account for the gap in rate, which remains once timing and output are controlled for. Among countries already publishing, G77 membership roughly halves the odds that a given year yields a first inscription, holding constant the year, the time since entry, the volume published and the genre. The inequality therefore lies both in whether countries enter early and publish often, and in what they find left to name when they do.

G77 members and other later or lower-volume participants still adopt, adapt, and circulate the lexicon, but they enter a field in which many categories have already been inscribed. Their lower first-inscription credit reflects their temporal position within the documentary field as well as the concerns they articulate. What matters is therefore which parts of the field possess the temporal and structural capacity to place categories into the record and support their circulation.

Persistence and the long tail of inscription

The repertoire dynamics matter because they determine whether unequal first inscription has lasting consequences. The global strategic threat lexicon widens over time, and once topics appear, they tend not to disappear permanently; instead, they move through dormancy and reactivation. This creates a field in which threat categories accumulate by layering rather than being routinely replaced.

Persistence amplifies the consequences of unequal inscription. If threat categories quickly disappeared permanently, unequal first appearances would matter less. The reactivation counts reported in the previous section show that categories can return after extended inactivity. Inequality in who is recorded as contributing to the lexicon is not only a snapshot distributional fact; it can become an accumulated property of the field. Through

documentary memory, prior inscription power can therefore outlast the opportunity advantage under which a category first entered the record.

Methodological contribution and limits

The paper combines a post-structuralist account of threat inscription with corpus-wide measurement. Both target the author-centrism of much securitization research, which presumes that threat construction originates in the choices of identifiable securitizing agents. By decentring the agent and foregrounding the documentary field as a condition of possibility for threat-writing, the analysis aligns with Foucauldian attention to discursive formations (Foucault, 2002). By operationalizing this through corpus-wide computational analysis, it reveals patterns that studies of preselected cases cannot observe.

The combination is not paradoxical. Quantitative methods do not carry a fixed epistemological commitment and can be used to investigate questions central to critical and constructivist research (Barkin & Sjoberg, 2017). Here, corpus-wide comparison serves a poststructuralist concern with the conditions under which threat objects become sayable and repeatable, rather than treating texts as the transparent expressions of individual intention (Foucault, 2002). Moreover, these methods do not merely record a self-evident existing reality; by making some patterns visible and comparable, they also shape what becomes available for analysis (Latour, 1987, pp. 65–70). In this case, they reveal concentration, historical redistribution, adoption, widening, and reactivation that topic-preselected designs cannot establish. This approach also addresses a recognized limitation in securitization research: the selection bias inherent in case-based designs, which tend toward pre-selected topics already recognized as securitized (Baele & Sterck, 2015, p. 1122; Balzacq, 2011, p. 34).

Several scope conditions delimit what the evidence supports. First, the analysis concerns publicly available top-level strategy documents and explicit "threat" formulations; it does not measure audience acceptance, classified assessments, or policy implementation. Second, topic construction depends on clustering model-generated category labels, and the labelled sentence pool is unequal across countries and document genres. Documentary throughput and genre-specific use of explicit threat language may therefore affect both the constructed topic space and the distribution of opportunities for first inscription. Third, the analysis maps first inscription and adoption but does not trace transmission mechanisms. Combining these maps with evidence on drafting infrastructures, alliance planning cycles, donor-supported strategy writing, template circulation, and textual borrowing would strengthen causal claims about how threat categories become internationally legible and reusable. The findings therefore support a bounded theory of

prior inscription power, not a causal claim that earlier documents directly produce later texts.

Conclusion: Reframing the securitization question

This paper asked how documentary securitization operates at scale when measured as first inscription and adoption. Using 660 top-level strategic security and defence documents produced by 118 countries and territories from 1987 to 2025, we extracted explicit threat statements, clustered them into comparable topics, and mapped first inscriptions, adoption, persistence, and widening across time.

Three empirical conclusions follow. First, documentary threat construction overwhelmingly draws on an existing lexicon, and first inscriptions are highly unequal. Second, US dominance is historical rather than continuous: its share of first-inscription credit falls steeply across the three periods, and first inscriptions disperse as the field expands without ceasing to be structured. Third, the global threat repertoire widens through layering and has long, episodic memory: categories can return after extended periods of inactivity.

These findings support an account of agenda-setting in strategic discourse in which power operates not only through immediate persuasion but through the unequal capacity to inscribe categories early, render them internationally legible, and sustain their circulation. Inequalities in who first inscribes a category can therefore outlast the conditions that produced them, even as first inscription becomes more widely distributed in the mature field.

The implications extend to how we formulate the normative stakes of securitization. Wæver's original ethical-political question was: "Why do you call this a security issue?" (Wæver, 2011, p. 468). This question presumes authorial agency: a choice by the securitizing actor that requires justification. Our findings suggest a second question alongside the first for strategy-document contexts: "*Why do you not call this a security issue?*"

In a documentary field where the threat lexicon is substantially inherited, where most strategy-writing involves selecting from and recombining an already-inscribed repertoire, the burden of justification may shift. When a threat category is already widely inscribed and institutionally expected, non-adoption may become the choice requiring justification. The actor who declines to include a widely recognized threat is the one who has some explaining to do. The normative question must therefore address not only why actors securitize, but also how they reproduce, modify, omit, or resist categories that are already legible within the documentary field.

This reframing captures the documentary sense in which securitizing agency is more often reproductive than authorial. Most states write security by reinscribing an inherited threat lexicon, not by freely inventing new threat

objects. Securitization theory has always recognized that agency operates under felicity conditions, but the documentary field constitutes a further layer of constraint, shaping which threat objects are available for inscription before any particular securitizing move is attempted. It structures which categories are already available, recognizable, and potentially expected before any particular strategy is written.

This does not mean that all securitization is compelled or that agency disappears. Strategy authors make choices: about emphasis, about which inherited topics to include or omit, about how to frame and adapt existing categories. But the space of available threat objects is substantially prepopulated. There is no formal bar on introducing novel threat categories, yet the path-dependent structure of the field makes reinscription far easier than invention: the inherited lexicon is already written, formatted, and internationally legible. Introducing a genuinely novel category requires not only naming it but establishing its recognizability within documentary conventions.

Understanding this requires attending to the documentary conditions under which threats become writable. These conditions are historically constituted through early institutionalization of recurring strategy publication, persist through the long memory of an archive-like repertoire, and are structured by unequal opportunities for inscription and circulation.

Note

1. Spearman's rho (ρ) measures correlation on a scale from -1 to $+1$. A value of 0.746 indicates a strong positive relationship, meaning the two variables tend to increase together.

Acknowledgements

We thank Sally Gardner and Joe Murphy for their work on text coding and validation. We are also grateful to Francesco Ragazzi, Matthias Leese and Emily Gilbert for their comments on an earlier version of this paper. Any remaining errors are our own.

AI disclosure statement

Generative AI tools were used both as reported analytical instruments and as aids to coding, validation design, and manuscript preparation. Gemini 2.5 Flash (gemini-2.5-flash, Google), accessed through the Google Batch API on 11 January 2026, assigned short threat-category labels to the extracted sentences. Its role, prompt, inputs, outputs, limitations, and human validation are reported in the Methods and replication materials (Neal, 2026). GPT-5.4 (gpt-5.4-2026-03-05, OpenAI) was used to classify the resulting topics by security sector; this analysis was conducted using a fixed codebook and structured output and was subject to independent human validation.

Claude models, including Sonnet and Opus 4–5 (Anthropic, accessed via Claude.ai and Claude Code between December 2025 and August 2026), assisted with Python and Jupyter code development, debugging, refactoring, validation and clustering protocols, robustness checks, documentation, and figure code. ChatGPT models, including GPT-5.2 and GPT-5.6 Sol (OpenAI, accessed between January and August 2026), assisted with debugging, methodological review, language editing, restructuring, and drafting alternative formulations. AI-generated suggestions were not accepted without author review, testing, and revision. All figures and quantitative results were produced from the study data by author-run Python code rather than generative image tools.

The authors determined the research questions, theoretical framework, study design, analytical decisions, interpretation, and final wording. They verified the originality and accuracy of the manuscript, checked references and empirical outputs, and take full responsibility for the integrity of the entire work. No AI tool is credited as an author. The authors reviewed the applicable terms of use and confirm that the tools were used in ways suitable for publication. Prompts, recorded model identifiers and configurations, inputs, saved outputs, validation materials, and relevant analysis code for AI tools used as analytical instruments are included in the replication package (Neal, 2026). Original and AI-assisted manuscript versions have been retained and can be shared with the journal on request.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

This work was supported by BA/Leverhulme Small Research Grants: [Grant Number SRG23\231686]; Foreign, Commonwealth and Development Office: [Grant Number 300708-144]; Scottish Council on Global Affairs.

ORCID

Andrew W. Neal http://orcid.org/0000-0002-7386-9255
Roy B. Gardner http://orcid.org/0000-0002-1336-1476

Reference list

African Union Commission. (2013). *African Union policy framework on security sector reform*. https://www.peaceau.org/uploads/au-policy-framework-on-security-sector-reform-ae-ssr.pdf

Baele, S. J., & Jalea, D. (2022). Twenty-five years of securitization theory: A corpus-based review. *Political Studies Review*, *21*(2), 376–389. https://doi.org/10.1177/14789299211069499

Baele, S. J., & Sterck, O. C. (2015). Diagnosing the securitisation of immigration at the EU level: A new method for stronger empirical claims. *Political Studies*, *63*(5), 1120–1139. https://doi.org/10.1111/1467-9248.12147

Balzacq, T. (2011). *Securitization theory: How security problems emerge and dissolve*. Routledge. https://doi.org/10.4324/9780203868508

Balzacq, T. (2019). Securitization theory: Past, present, and future. *Polity*, *51*(2), 331–348. https://doi.org/10.1086/701884

Balzacq, T., Léonard, S., & Ruzicka, J. (2016). 'Securitization' revisited: Theory and cases. *International Relations*, *30*(4), 494–531. https://doi.org/10.1177/0047117815596590

Barkin, J. S., & Sjoberg, L. (2017). *Interpretive quantification: Methodological explorations for critical and constructivist IR*. University of Michigan Press. https://doi.org/10.3998/mpub.7361329

Barthes, R. (1977). The death of the author. In S. Heath (Ed.), *Image, music, text* (pp. 142–148). Fontana.

Bigo, D. (2002). Security and immigration: Toward a critique of the governmentality of unease. *Alternatives: Global, Local, Political*, *27*(1), 63–92. https://doi.org/10.1177/03043754020270s105

Bigo, D. (2006). *Illiberal practices of liberal regimes: The (in)security games*. L'Harmattan.

Bigo, D. (2008). Globalized (in)security: The field and the ban-opticon. In D. Bigo & A. Tsoukala (Eds.), *Terror, insecurity and liberty: Illiberal practices of liberal regimes after 9/11* (pp. 10–48). Routledge. https://doi.org/10.4324/9780203926765.ch2

Bueger, C., & Gadinger, F. (2021). Making grand strategy in practice. In T. Balzacq & R. R. Krebs (Eds.), *The Oxford handbook of grand strategy* (pp. 142–158). Oxford University Press. https://doi.org/10.1093/oxfordhb/9780198840299.013.8

Buzan, B. (1983). *People, states and fear: The national security problem in international relations*. Wheatsheaf.

Buzan, B. (1991). *People, states and fear: An agenda for international security studies in the post-Cold War era* (2nd ed.). Wheatsheaf.

Buzan, B., & Hansen, L. (2009). *The evolution of international security studies*. Cambridge University Press. https://doi.org/10.1017/CBO9780511817762

Buzan, B., Wæver, O., & de Wilde, J. (1998). *Security: A new framework for analysis*. Lynne Rienner. https://doi.org/10.1515/9781685853808

CASE Collective. (2006). Critical approaches to security in Europe: A networked manifesto. *Security Dialogue*, *37*(4), 443–487. https://doi.org/10.1177/0967010606073085

Chang, J., Gerrish, S., Wang, C., Boyd-Graber, J., & Blei, D. (2009). Reading tea leaves: How humans interpret topic models. *Advances in Neural Information Processing Systems*, *22*, 288–296.

Der Derian, J., & Shapiro, M. J. (1989). *International/intertextual relations: Postmodern readings of world politics*. Lexington Books.

de Vries, E., Schoonvelde, M., & Schumacher, G. (2018). No longer lost in translation: Evidence that Google Translate works for comparative bag-of-words text applications. *Political Analysis*, *26*(4), 417–430. https://doi.org/10.1017/pan.2018.26

Doyle, R. B. (2007). The U.S. national security strategy: Policy, process, problems. *Public Administration Review*, *67*(4), 624–629. https://doi.org/10.1111/j.1540-6210.2007.00748.x

Estivill-Castro, V. (2002). Why so many clustering algorithms: A position paper. *ACM SIGKDD Explorations Newsletter*, *4*(1), 65–75. https://doi.org/10.1145/568574.568575

Foucault, M. (2002). *The archaeology of knowledge* (A. M. S. Smith, Trans.). Routledge. https://doi.org/10.4324/9780203604168

Freeman, R., & Maybin, J. (2011). Documents, practices and policy. *Evidence & Policy*, *7*(2), 155–170. https://doi.org/10.1332/174426411X579207

Hansen, L., & Nissenbaum, H. (2009). Digital disaster, cyber security, and the Copenhagen School. *International Studies Quarterly*, *53*(4), 1155–1175. https://doi.org/10.1111/j.1468-2478.2009.00572.x

Huysmans, J. (2006). *The politics of insecurity: Fear, migration and asylum in the EU*. Routledge. https://doi.org/10.4324/9780203008690

Huysmans, J. (2011). What is in an act? On security speech acts and little security nothings. *Security Dialogue*, *42*(4–5), 371–384. https://doi.org/10.1177/0967010611418713

Knudsen, B. B. (2012). Developing a national security policy/strategy: A roadmap. *Sicherheit und Frieden (S+F) / Security and Peace*, *30*(3), 135–140. https://doi.org/10.5771/0175-274x-2012-3-135

Latour, B. (1987). *Science in action: How to follow scientists and engineers through society*. Harvard University Press.

Milligan, G. W., & Cooper, M. C. (1987). Methodology review: Clustering methods. *Applied Psychological Measurement*, *11*(4), 329–354. https://doi.org/10.1177/014662168701100401

Mouton, J., & Marais, H. C. (1988). *Basic concepts in the methodology of the social sciences*. Human Sciences Research Council.

NATO (1999). Official text: Membership action plan (MAP), 24-apr.–1999 'Membership action plan (MAP) approved by the heads of state and government participating in the meeting of the North Atlantic Council'.

Neal, A. W. (2010). *Exceptionalism and the politics of counter-terrorism: Liberty, security and the war on terror*. Routledge. https://doi.org/10.4324/9780203867587

Neal, A. W. (2019). *Security as politics: Beyond the state of exception*. Edinburgh University Press. https://doi.org/10.3366/edinburgh/9781474450928.001.0001

Neal, A. W. (2026). Replication package for 'Securitization through inscription: National security documents and the making of a global threat lexicon (1987–2025)' [Dataset]. University of Edinburgh. https://doi.org/10.7488/ds/8197

Neal, A. W., & Gardner, R. B. (2026). National security and defence documents dataset (1987–2025) v3.0 [Dataset]. University of Edinburgh. https://doi.org/10.7488/ds/8076

Neumann, I. B. (2007). "A speech that the entire ministry may stand for," or: Why diplomats never produce anything new. *International Political Sociology*, *1*(2), 183–200. https://doi.org/10.1111/j.1749-5687.2007.00012.x

Nyman, J. (2018). *The energy security paradox: Rethinking energy (in)security in the United States and China*. Oxford University Press.

Organization of American States. (1995). *Confidence and security building measures in the Americas*. https://2009-2017.state.gov/documents/organization/120043.pdf

Pacific Islands Forum. (2018). *Boe Declaration on Regional Security*. https://forumsec.org/publications/boe-declaration-regional-security

Prior, L. (2008). Repositioning documents in social research. *Sociology*, *42*(5), 821–836. https://doi.org/10.1177/0038038508094564

Schäfer, M. S., Scheffran, J., & Penniket, L. (2016). Securitization of media reporting on climate change? A cross-national analysis in nine countries. *Security Dialogue*, *47*(1), 76–96. https://doi.org/10.1177/0967010615600915

Stolberg, A. G. (2012). *How nation-states craft national security strategy documents*. Strategic Studies Institute, US Army War College.

Trombetta, M. J. (2008). Environmental security and climate change: Analysing the discourse. *Cambridge Review of International Affairs*, *21*(4), 585–602. https://doi.org/10.1080/09557570802452920

Vuori, J. A. (2008). Illocutionary logic and strands of securitization: Applying the theory of securitization to the study of non-democratic political orders. *European Journal of International Relations*, *14*(1), 65–99. https://doi.org/10.1177/1354066107087767

Wæver, O. (2004, March 17–20). Aberystwyth, Paris, Copenhagen: New "schools" in security theory and their origins between core and periphery. 45th Annual Convention of the International Studies Association, Montreal, Canada.

Wæver, O. (2011). Politics, security, theory. *Security Dialogue*, *42*(4–5), 465–480. https://doi.org/10.1177/0967010611418718