



UvA-DARE (Digital Academic Repository)

Epistemic Capture Through Specialization in Post-World War II Parliamentary Debate

Ros, R.; Wevers, M.

DOI

[10.1017/chr.2025.10008](https://doi.org/10.1017/chr.2025.10008)

Publication date

2025

Document Version

Final published version

Published in

Computational Humanities Research

License

CC BY

[Link to publication](#)

Citation for published version (APA):

Ros, R., & Wevers, M. (2025). Epistemic Capture Through Specialization in Post-World War II Parliamentary Debate. *Computational Humanities Research*, 1, Article e6. <https://doi.org/10.1017/chr.2025.10008>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, P.O. Box 19185, 1000 GD Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (<https://dare.uva.nl>)

Epistemic Capture Through Specialization in Post-World War II Parliamentary Debate

Ruben Ros^{*,1} and Melvin Wevers²

¹Department of History and Art History, Utrecht University, The Netherlands

²Department of History, Universiteit van Amsterdam, Amsterdam, The Netherlands

Article-type

Keywords:

parliamentary discourse, epistemic capture, temporal networks

*Author for correspondence. Email: r.s.ros@uu.nl

Abstract

This article investigates how parliamentary debate in the Dutch House of Representatives (“Tweede Kamer”) (1945–1995) narrowed as MPs turned into domain specialists. We call this narrowing epistemic capture: a few experts progressively bound what can be said. To detect epistemic capture, we deploy a three-layer computational pipeline. Latent-Dirichlet topic modeling converts 8.2 million sentences into 250 semantic themes; Pointwise Mutual Information networks connect themes within six-month windows; Louvain clustering traces the birth, drift, and endurance of topical communities.

Capture appears on every scale. Macro-level: network modularity almost doubles after 1960 while density falls, marking compartmentalized debate. Meso-level: cabinet turnovers act as “reset switches”: topic-neighborhood similarity drops in the half-year after a new coalition forms, then anneals along partisan lines. Micro-level: enduring communities—foreign policy, agriculture, education—lock topics and MPs together for decades, yet resistance to capture is visible in distinct contentious topics.

These multiscale patterns show how twentieth-century Dutch parliamentary debate saw a rise of technical specialism that significantly constrained the breadth of political debate. Methodologically, the study demonstrates the value of structural (network) distant reading over purely lexical counts and offers a transferable workflow for measuring how democratic discourse undergoes structural transformations.

Plain Language Summary

This study examines how political debates in the Dutch parliament became increasingly specialized between 1945 and 1995, focusing on how this specialization affected democratic discussion. While specialization—where politicians become experts in specific policy areas—can lead to more informed decision-making, it may also narrow political debate and exclude broader perspectives.

Using computational methods, we analyzed thousands of parliamentary speeches to understand these changes. Our approach combined three key techniques: First, we used topic modeling to identify what politicians discussed. Second, we created networks showing how different topics connected to each other. Third, we traced how these topic networks evolved over time, revealing stable groups of interconnected topics that we call “specialist communities”. Our findings show three important developments: First, parliamentary debate gradually split into distinct specialist domains, especially between 1950–1975. Second, these specialist domains became increasingly stable and isolated from each other. Third, individual politicians increasingly focused on specific topics, forming lasting specialist communities.

While specialization enhanced technical expertise, it also led to what we call “epistemic capture”—where specialist perspectives began to dominate certain policy areas, potentially limiting broader democratic discussion. Consequently, this research offers new insights into how modern democracies balance the need for expert knowledge with maintaining open political debate.

Conceptual Framework: Specialization and Epistemic Capture

Specialization constitutes a defining characteristic of capitalist modernity, fundamentally transforming social organization through the division of labor into distinct domains of expertise. This fragmentation into specialized spheres has catalyzed unprecedented gains in efficiency and productivity across economic, administrative, and scientific domains (Durkheim 2014). Political institutions exemplify this broader transformation: decision-making processes have increasingly been delegated to specialist experts, while politicians have evolved from generalists into specialists (Vibert 2007).

CAMBRIDGE
UNIVERSITY PRESS

This is an Open Access article, distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives licence (<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 unaltered and is properly cited. The written permission of Cambridge University Press must be obtained for commercial re-use or in order to create a derivative work.

The growing complexity of policy challenges has rendered generalist approaches increasingly impractical in political decision-making. In modern democracies, politicians have become specialists; experts who are well-versed in particular policy domains. While such specialization enhances analytical depth and procedural efficiency in governance, it also introduces tensions within democratic systems—a concern articulated by political philosophers from John Dewey to Jürgen Habermas (Dewey 2012; Habermas and Shapiro 1970). There is a growing risk that citizens' experiential lifeworlds and specialists' domain-specific reasoning diverge, undermining reciprocal understanding in democratic deliberation. In fact, contemporary populist movements can be interpreted as a manifestation of this divergence, reflecting societal resistance to the growing hegemony of technocratic expertise (Scanni 2023). This phenomenon reveals a deeper theoretical contradiction: specialization, driven by imperatives of efficiency and empirical precision, stands in fundamental tension with the principles of pluralism and contingency that underpin democratic politics (Van Bouwel 2023). Specialists value skill and truth, whereas democratic politics, conceptualized as an arena of contestation and negotiation, inherently resists the drive toward optimization and definitive truth claims (Lefort 1988; Urbinati 2014). This tension threatens the deliberative inclusivity on which democratic legitimacy depends.

The fundamental challenge that specialization poses to democratic governance manifests through what we term “epistemic capture” (Zubčić and Giacomini 2024)—a phenomenon whereby domain-specific expertise generates constrained epistemological frameworks that come to dominate analytical and decision-making processes. This epistemic dominance occurs when specialized knowledge frameworks acquire privileged status as authoritative truth claims, transcending their original contextual boundaries (Haas 1992). Political scientists have long acknowledged specialists' crucial role as epistemic mediators for policymakers, translating complex realities into actionable frameworks (Buchanan *et al.* 1960). These specialist communities increasingly exercise both political and epistemic authority within their domains of expertise. The Dutch “Green Front”—a collective of agricultural specialists in parliament—exemplifies this dual hegemony, demonstrating how expert communities can simultaneously shape both policy outcomes and the underlying conceptual frameworks through which policy challenges are understood (Karel 2013).

Such intellectual encapsulation functions as both as a cognitive filter and as an institutional mechanism, constraining the scope of democratic deliberation to what is considered to be realistic or evidence-based.¹ The resultant epistemic myopia not only circumscribes the range of potential policy solutions but fundamentally restructures the parameters of democratic discourse itself. This process creates self-reinforcing feedback loops where specialized knowledge frameworks increasingly monopolize the conceptualization and articulation of policy challenges, effectively “capturing” democratic debate within technocratic boundaries (Bohman 1999).

Specialization and the Study of Parliament

Viewing specialization through the lens of “epistemic capture” challenges conventional understandings of parliamentary development in political science and historical research. First, political scientists often consider specialization an unavoidable adaptation to the growing complexity of policymaking, positioning the parliamentary committee as central to this development (Shaw 1998). Historians, inspired by Max Weber's work on bureaucracy and parliament, have observed how the consolidated power of bureaucratic and executive institutions has prompted procedural changes within legislative bodies (Koš 2018). Twentieth-century democracies experienced a process of “governmentalization”: the concentration of legislative prerogatives and agenda-setting powers in the hands of executive institutions such as presidents or, in the Dutch case, the Council of Ministers (Palonen 2015; Koole 2018). Specialization can be seen as parliament's reflex, both triggered by and accommodating this process.

Rather than treating specialization as an inevitable adaptation to policymaking demands, our concept of epistemic capture highlights its potential tension with political contestation: specialization is not merely a neutral instrument for improving politics, but a development that shapes and can constrain parliamentary debate. In this sense, it offers a way to trace the “informal” impact of long-term institutional processes such as governmentalization (Te Velde and Voermans 2019). Building on this insight, our multi-level analytical framework highlights the subtle, informal epistemic dimensions of parliamentary evolution that extend beyond structures like committees. It suggests that the very framing and understanding of policy issues within parliament may become dominated by specific expert perspectives, narrowing the scope of political debate. Consequently, understanding parliament requires attention not only to formal architecture, but also to the less visible, knowledge-based dynamics that shape how issues are conceptualized.

Second, viewing specialization through the lens of epistemic capture challenges prevailing approaches to the study of (historical) parliamentary discourse. In political and parliamentary history, research has traditionally relied on hermeneutic methods, centered on close textual analysis, to examine how politicians debate. More recently, computational approaches have gained prominence in this field (Van Eijnatten and Huijnen 2021; Guldi 2019). By applying algorithms and statistical techniques to large-scale corpora of historical texts, these methods can uncover patterns in historical discourses that were previously difficult to detect. These approaches treat linguistic patterns as indicators of shared mental models, social representations, or rhetorical styles, thereby enabling systematic analysis of parliamentary discourse across both temporal and contextual dimensions (Marjanen 2023; Bolla *et al.* 2020; Moretti 2013).

While these computational frameworks offer novel analytical capabilities, traditional distant reading approaches often reduce multifaceted intellectual phenomena to discrete lexical units, treating keywords as direct indices of ideational

content. The resultant analytical framework focuses predominantly on frequency distributions, collocational patterns, and vector space representations of individual keywords. It may, while methodologically robust, inadequately capture the nuanced dynamics of (historical) idea systems. Though these approaches enhance research efficiency, their contribution to fundamental historiographical understanding remains limited by their inherent methodological assumptions.

Recently, scholars have begun employing more sophisticated methods to transcend the study of individual concepts and ideas. By prioritizing theory-driven modeling over exploratory lexical approaches, they have tackled questions and concepts in the history of ideas that were long regarded as too complex or large to empirically study. For instance, Ryan Heuser tests several long-standing hypotheses regarding general conceptual transformations during the “Sattelzeit” (saddle time), a transitional era between early modern and modern history (Heuser 2023). Recent work by Vicinanza et al. (Vicinanza, Goldberg, and Srivastava 2023) shows how ideas tend to emerge in peripheral spaces, illustrating how various studies have recently used information-theoretic measures to quantify the novelty and resonance of intellectual developments. These methodological advances depart from conventional “distant reading” approaches by emphasizing structural dynamics and processes (Poulsen and DeDeo 2023). Moving beyond simple distributional analysis of lexical patterns, they enable scholars to construct more ambitious theoretical frameworks for understanding historical discourse. Our study builds upon this methodological foundation, conceptualizing specialization as both a discursive process and an epistemic transformation that can be systematically analyzed through computational means.

The Dutch House of Representatives: A Case Study in Specialization

The Dutch House of Representatives presents a compelling empirical case for examining parliamentary specialization in the twentieth century. This institution underwent a clear shift in its conceptualization of specialized knowledge: from an initial skepticism of specialization as an intellectual constraint to its eventual embrace as an essential mechanism for effective governance. This epistemic reorientation marked a decisive break from nineteenth-century parliamentary culture, where specialized knowledge was often conflated with intellectual parochialism (Hoetink 2018). The interwar period catalyzed this transformation through an emerging cognitive tension: politicians increasingly observed a widening epistemological gap between the expanding state bureaucracies and the generalist orientation of parliamentary representatives. This perceived knowledge asymmetry precipitated the emergence of domain-specific parliamentary expertise, particularly in strategically critical areas such as foreign policy and agricultural governance—domains characterized by high technical complexity and specialized knowledge requirements. The post-war period witnessed the institutional crystallization of this epistemological transformation. Following broader Western European patterns, the establishment of permanent commit-

tees created formal spaces for specialized technical discourse, altering the structure of parliamentary deliberation (Siefken 2021; Aylett 2019). The Netherlands’ distinctive multiparty system, coupled with the desire of the main parties to depoliticize contentious issues, further accelerated this specialization trajectory (Lijphart 1968). This development marked a decisive break from the previous organizational paradigm of randomly allocated “departments”—an institutional arrangement that reflected and reinforced generalist epistemological assumptions. The 1966 Rules of Order, implemented after a decade of experimentation with permanent committees, dismantled the departments system (Hoetink 2018), institutionalizing specialized knowledge domains as the primary organizing principle of parliamentary work. Since then, the Dutch House of Representatives has grown into a legislative institution that is marked by high levels of specialization (Andeweg 1997).

Situated within this historical context and informed by the broader theoretical implications of specialization in democratic institutions, our study aims to answer the question of how specialization engendered epistemic capture in parliamentary debates in the post-war period (1945–1995). To systematically investigate this question, we developed a novel computational framework that combines three complementary methodological approaches: topic modeling for semantic content extraction, network analysis for relational structure identification, and dynamic community detection for temporal pattern recognition. This integrated analytical approach allows us to trace the evolution of specialized knowledge domains within the Dutch House of Representatives across five decades, illuminating both the mechanisms of specialization and their implications for parliamentary discourse.

By applying temporal network analysis to topic linkages in parliamentary debates, we aim to advance understanding of how epistemic structures shape politics. This approach allows us to trace the complex interplay between specialized knowledge formation and democratic deliberation. Moreover, our study contributes to the broader scholarly discussions regarding the inherent tensions between technical expertise, pluralistic discourse, and political contingency in democratic systems, offering empirically grounded insights into contemporary challenges facing democratic institutions.

Data: Parliamentary Proceedings

This study draws upon a comprehensive corpus of digitized parliamentary proceedings of the Dutch House of Representatives from 1946 to 1995. The proceedings consist of speeches held in the House that are transcribed, edited, and published by the Stenographical Service. In recent years, the proceedings have been digitized and linked with metadata on speakers, parties, and dates. The quality of the Optical Character Recognition improved considerably in post-war proceedings, ensuring robust data quality for computational analysis (Van Lange 2023). To prepare the data for analysis, we performed several preprocessing steps. First, we removed stop words to eliminate common but uninformative words. Next, we filtered out speeches shorter than ten words, as these are unlikely to

contain substantial content. We constrained our analysis to semantically rich parts of speech—nouns, verbs, adjectives, and adverbs—because these most effectively capture conceptual content and thematic structure.

In addition to text preprocessing, we also excluded several types of speeches. First, we restrict our analysis to plenary debates. With the omission of committee meetings (also present in our data), we prevent our models from reflecting only committee language and ensure they capture specialization in plenary debates. Second, we excluded speeches by the House chair, which often contain procedural and formulaic language that could potentially skew our topic models and linkage analyses.

The resulting corpus comprises 52,396,073 tokens and 495,053 unique types. The data size varies unevenly over the period; after 1967, the size of the parliamentary text gradually expands. Whereas an average year in the 1950s accounted 1.25% of the tokens, this share rose to 4% in 1979. This quantitative shift aligns with broader institutional transformations in parliamentary practice.

Methods

Our methodological approach consists of four main steps: 1) pre-processing of parliamentary proceedings, 2) topic modeling to identify thematic content, 3) network analysis to map relationships between topics, and 4) dynamic community detection to identify clusters of specialized domains (see Figure 1). This multi-step approach allows us to trace the evolution of specialized knowledge domains in the Dutch House of Representatives across five decades.

Topic Models

To investigate specialization in parliamentary debates, we used topic modeling to analyze the semantic content of speeches. We employed Latent Dirichlet Allocation (LDA) using the Mallet toolkit to generate low-dimensional representations of speech content (Blei, Ng, and Jordan 2003; McCallum 2002).²

To prepare the data for the model, we segmented each speech into 50-word chunks. This approach accounts for the general decline in speech lengths over the period we studied and its potential effect on topic diversity.³ After training the model, we averaged the topic distributions for each member within a single session to facilitate our subsequent network analysis.

Our topic model captured a mix of distinct policy areas as well as procedural and rhetorical categories (e.g., filing motions or citing newspapers). Since our focus was on thematic specialization in policy domains, we excluded the procedural and rhetorical topics. We first ranked topics based on their document entropy—a measure of how uniformly a topic is distributed across documents. Higher entropy scores indicated more general topics, which were more likely to be non-thematic. After this initial ranking, we manually reviewed and filtered out any topics that were clearly procedural or rhetorical. This ensured that our final set of topics represented substantive content.

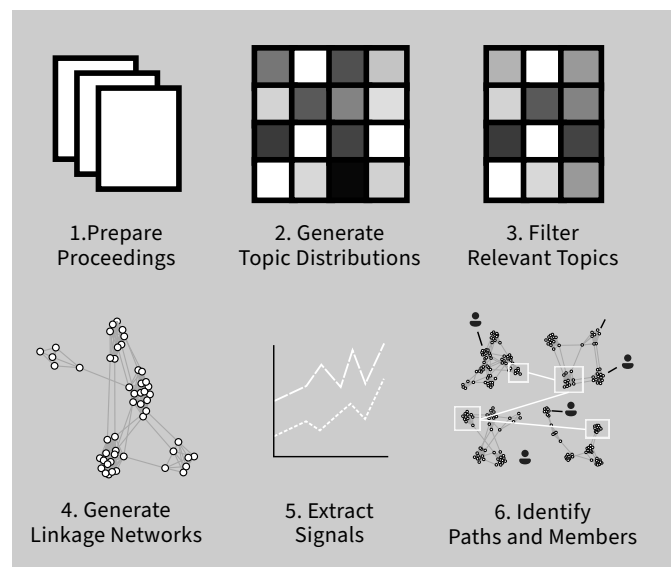



Figure 1. *Top:* example topic-linkage network. Nodes are detected by Louvain community; edge thickness is proportional to Pointwise Mutual Information. *Bottom:* methodological workflow from corpus preprocessing to dynamic community detection.

Filtering complicated the use of standard metrics, such as coherence scores (a common measure of topic quality), to determine the optimal number of topics. We ultimately selected 250 topics, of which 116 were identified and removed as non-thematic. This number provided the optimal trade-off between thematic granularity and interpretability: fewer topics risked conflating distinct policy areas, while more topics tended to yield overly specific representations, often limited to single events or pieces of legislation. After excluding the non-thematic topics, we renormalized topic distributions to ensure comparability across parliamentary sessions.⁴ This approach enhances the resolution of inter-thematic connections in our network analysis, and our primary analytical metrics remained stable under this protocol, supporting its validity.

Networks and Communities

To measure the topological relationships between parliamentary speech topics, we used Pointwise Mutual Information (PMI)-based topic linkage scores, extending the methodological framework developed by Perry and DeDeo Perry and DeDeo (2021) (see Appendix Linkage function). PMI provides a probabilistic measure of topic co-occurrence that exceeds random expectation, where elevated scores indicate stronger thematic associations. We construct network structures by calculating pairwise PMI scores between topics, establishing

edges where PMI values exceed zero, thereby capturing substantive thematic relationships (see Figure 1 as an example).

This network-analytic approach enables a systematic investigation of temporal evolution in thematic linkage architectures. To mitigate PMI's known sensitivity to low-frequency events, we implement a joint frequency-based smoothing function, preventing spurious centrality of rare topics and enhancing the reliability of topical connection representations (Pantel and Lin 2002). Our temporal granularity is calibrated to six-month intervals, aligning with the periodicity of parliamentary activity—specifically accounting for the distinct characteristics of budgetary debates (August–December) versus general legislative debates (January–July).

Following network construction, we applied the Louvain community detection algorithm to identify clusters of related topics (Blondel et al. 2008). This algorithm optimizes modularity by quantifying the differential density of intra-community versus inter-community links, revealing emergent specialized epistemic domains within parliamentary discourse. The identified communities represent distinct spheres of specialized knowledge and discourse.

Through systematic parameter optimization, we evaluated multiple resolution parameters that govern the size and number of detected communities. A resolution parameter of 3 yielded optimal results, producing communities that demonstrate both thematic coherence and historical validity when manually evaluated against Dutch parliamentary historiography. To analyze the properties and evolution of these topic networks and communities, we rely on several key metrics.

1. **Network Topology Metrics.** Modularity and network density, and node degree constitute our primary network-structural measures. Modularity quantifies the strength of community partitioning within the network—where elevated values indicate dense intra-community connections coupled with sparse inter-community links, suggesting well-defined specialized domains (Newman 2006). Network density provides a global measure of topological cohesion, formally defined as the ratio between observed edges and the maximum possible number of edges (Mitchell 1969). Complementing this global measure, node degree quantifies local connectivity by measuring the number of direct connections maintained by individual topics. Together, these metrics enable identification of both system-wide connectivity patterns and locally influential topics that may serve as structural anchors in parliamentary discourse networks. The combination of global density and local degree centrality metrics reveals both the overall cohesion of parliamentary debate and the emergence of focal topics that disproportionately influence thematic organization.
2. **Temporal Clustering Metrics** We employ Normalized Mutual Information (NMI) to assess temporal stability in community structures. This information-theoretic measure quantifies the preservation of community membership across time periods, generating normalized scores between 0 and 1. Higher scores indicate stronger temporal persistence in specialized domain structures, while lower scores

suggest structural reorganization (McDaid, Greene, and Hurley 2013).

3. **Set Similarity Metrics.** To analyze the evolutionary dynamics of community structures, we implement the overlap coefficient as our primary set-theoretic similarity measure (Szymkiewicz 1934). This metric, calculated as the ratio of intersection size to minimum set size, enables quantification of node neighborhood persistence and community evolution. The overlap coefficient proves particularly suitable for tracking specialized domain evolution, as it effectively captures nested relationships where smaller specialized communities may be subsumed within larger thematic structures. The overlap coefficient demonstrates particular methodological advantages over traditional Jaccard similarity for quantifying epistemic capture phenomena. Whereas Jaccard similarity, normalized by set union size, potentially obscures complete subset relationships, the overlap coefficient's normalization by minimum set size preserves critical information about nested structural relationships. This property proves especially valuable in longitudinal analyses where community sizes exhibit temporal variation, enabling robust detection of specialized domain absorption and emergence patterns. We generate temporal stability profiles by computing pairwise overlap coefficients between node neighborhoods in adjacent time periods (T and $T - 1$), then averaging these coefficients to construct a composite measure of structural persistence. This sequential comparison approach, while methodologically straightforward, effectively captures both gradual drift and abrupt reorganizations in specialized knowledge domains. The resulting temporal stability mappings enable identification of both persistent structural features and periods of significant epistemic reorganization.
4. **Dynamic Community Detection.** To systematically trace the temporal evolution of specialized communities, we implement dynamic community detection using the Community Discovery Library (CDlib). The Temporal Network Clustering algorithm enables tracking of community formation, dissolution, and transformation across sequential time periods (Rossetti, Milli, and Cazabet 2019). We operationalize community similarity through the previously established overlap coefficient, measuring temporal coherence through topic composition consistency. The algorithm links communities across adjacent time steps based on node overlap, providing insights into the stability and dynamics of specialized knowledge domains in the Dutch House of Representatives. The resulting temporal community paths illuminate multiple aspects of parliamentary specialization: the persistence of established policy domains, the emergence of novel specialized communities, and the structural reorganization of existing knowledge domains in response to evolving political priorities and institutional changes.

Results

Our analytical approach moves beyond single-metric measurements of specialization, implementing instead a multi-level network analysis framework that captures the complex dynamics of parliamentary specialization. This methodological strategy integrates multiple analytical scales: global network metrics illuminate long-term structural transformations, while granular analyses of community structures, topical relationships, and edge patterns reveal localized dynamics and contextual variations in specialization processes.

Our analysis distinguishes three distinct yet interrelated analytical levels: macro-level institutional patterns, meso-level political dynamics, and micro-level individual specialization processes. Using metrics at each level, we integrate different explanations, allowing us to differentiate between various causal forces and contextual factors that shape specialization.

Macro-Level Analysis: The Emergence of Specialized Communities

Our macro-level analysis examines the structural evolution of specialization within the House of Representatives through network topology metrics. Our primary indicator, network modularity, quantifies the degree of discourse compartmentalization, where increasing values signal progressive differentiation into distinct specialist domains. The temporal analysis reveals two distinct phases: a notable increase in modularity between 1950 and 1975, followed by stabilization persisting until the 1990s. This modularity trajectory exhibits an inverse relationship with network density (the overall “connectedness” of a network). This divergence between rising modularity and declining density (Figure 2) indicates a fundamental transformation in parliamentary discourse: while topics increasingly cluster into specialized communities (increasing modularity), they simultaneously become more isolated from other topics (decreasing density), suggesting progressive compartmentalization of expertise at the expense of broader policy discussions. We use Cliff’s delta (see Appendix), a robust measure for ordinal data, to test the effect size of this development (Meissel and Yao 2024). The overall results indicate a statistically significant and large ($\delta = 0.59, [-1, 1]$) increase in modularity in the 1950s and 1960s.

While modularity and density metrics demonstrate the emergence of distinct community structures, we examined the temporal persistence of these specialized domains. If communities differ in population over time, but the overall level of clusteredness (i.e. modularity) remains stable, this does not indicate specialization. This would indicate fluctuating topic clustering rather than enduring specialization. To quantify this structural stabilization, we employ Normalized Mutual Information (NMI) to measure inter-temporal community stability (Danon *et al.* 2005). This information-theoretic measure provides a standardized metric of clustering consistency across successive time periods.

Analysis of consecutive six-month periods from 1944 to 1995 reveals a systematic increase in NMI scores, rising from approximately 0.5 to 0.6 (Figure 3). This upward trajectory

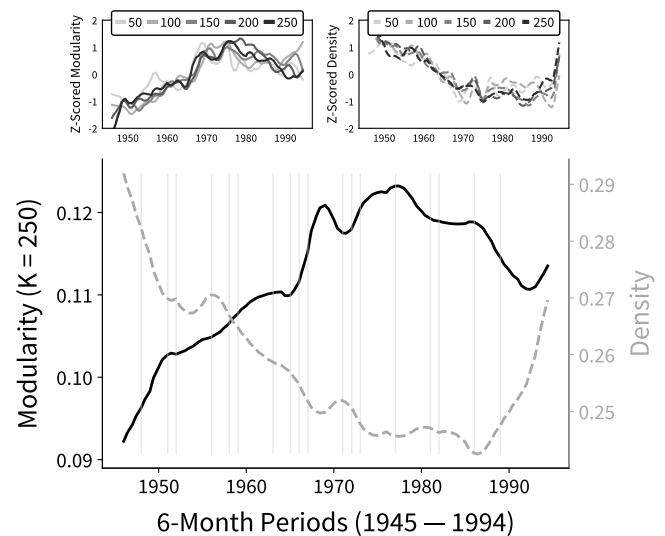


Figure 2. Modularity (solid black) and density (dashed grey) of six-month topic-linkage networks. The main panel shows the signals based on the topic model with 250 topics. Vertical dashed lines mark cabinet changes. The top panels show the z-score normalized modularity (left) and density (right) signals for topic models with different numbers of topics.

indicates progressive crystallization of specialized communities, suggesting not merely the formation but the institutional entrenchment of specialized discourse patterns. The temporal stability of these topic clusters demonstrates the consolidation of distinct epistemic domains within parliamentary debate.

These empirical patterns align with the documented institutionalization of the committee system during the 1950s and 1960s, when specialized committees emerged as primary parliamentary actors. However, our analysis reveals that this institutional transformation extended beyond procedural changes to fundamentally reshape parliamentary discourse: as specialized domains crystallized, debate patterns exhibited greater analytical depth within domains but reduced thematic breadth across domains.

Surprisingly, specialization does not persist throughout the entire period under study. In the late 1970s, the modularity signal exhibits a modest decline, while density increases from the late 1980s onward. We argue that a general focus on economic issues, framed through a neoliberal lens, fostered greater associations between topics, thereby challenging stable clusters. By the end of the 1970s, topics related to economic crisis, budget shortages, and austerity came to dominate parliamentary debate. Notably, the topic with the strongest growth in degree centrality concerns business subsidies. This centrality stems from the so-called RSV affair, in which financial mismanagement and scandals surrounding government-backed industrial subsidies attracted widespread attention between 1983 and 1985 (Kroeze and Keulen 2020). The affair stimulated a neoliberal critique asserting that state intervention had overreached, casting government economic involvement in a critical light. Consequently, this topic exemplifies how various policy areas became subject to a broad neoliberal critique, which blurred the boundaries between previously compartmentalized spe-

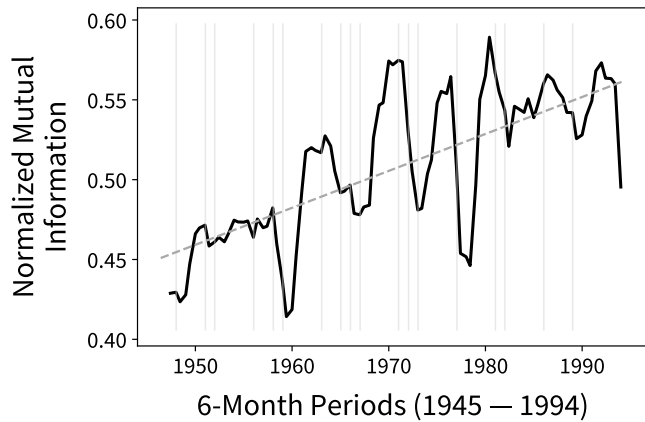


Figure 3. Clustering stability, expressed as Normalized Mutual Information (NMI) between successive linkage networks. A linear regression line is plotted to illustrate the overall trend, and vertical dashed lines mark cabinet changes.

cialist clusters. This broader ideological framing increased associations between topics. These patterns indicate that specialization does not follow a linear trajectory but is shaped by political developments that challenge specialist frameworks.

Meso-Level Analysis: Specialist Communities between Politics and Procedures

Macro-level analyses confirm the gradual specialization of the House of Representatives, especially in the decades between 1950 and 1975. However, they reveal little about the driving factors and local dynamics that are clearly present in the previous figures. To investigate specialization beyond linear increase, we focus on two meso-level methods.

First, we examine the development of overlap in node neighborhoods over time. Using the overlap coefficient, we study how much of the topic neighborhood is preserved from one time point to the next. We refer to this metric as “Topic Neighborhood Similarity” (TNS). The lower panel in Figure 4 shows the average overlap ratio as it develops over time.⁵ The upward trend reflects an increasing overlap in node neighborhoods, confirming our macro-level finding that parliament witnesses the emergence of stable specialist clusters. However, the figure also shows more granular patterns of specialization development. These patterns appear as systematic perturbations in neighborhood stability coinciding with cabinet transitions: each change in government precipitates an initial destabilization of topic relationships, followed by gradual re-stabilization during the cabinet’s tenure. We observe cycles of epistemic “reset,” followed by periods of progressive network crystallization. These patterns challenge deterministic models of specialization: rather than exhibiting immunity to political changes, specialist communities demonstrate dynamic responsiveness to shifts in parliamentary composition and political leadership. This sensitivity to political transitions suggests that specialized knowledge structures, while increasingly stable over time, remain susceptible to reconfiguration through political contingencies.

Beyond identifying meso-level “resets”, our approach also




Figure 4. Topic Neighborhood Similarity (TNS). TNS measures the similarity between the neighborhood of a topic (node) in a period and its N preceding period as measured with the overall coefficient. The bottom figure shows the average similarity for all nodes over time. Vertical dashed lines indicate cabinet changes. The top figures show the similarity scores for three specific nodes (topics): a topic with a rising neighborhood similarity (environmental management), a topic with a persistently high similarity (international conflict) and a topic with persistently low similarity (public broadcasting). For each of these topics, the average similarity is indicated with the horizontal line.

assesses which domains actually specialize. By tracing temporal linkages between communities across successive periods, we identify persistent “paths” of specialized discourse—coherent sequences of thematically related communities that persist over time. These community paths, composed of interconnected topic clusters, provide granular evidence of specialization’s structural manifestation.

The identification of paths of connected communities allows us to assess which political themes are most specialized. Our analysis highlights three domains with particularly pronounced specialization: foreign policy, education, and agriculture. These areas demonstrate the highest degrees of temporal continuity and thematic coherence in their respective community paths, suggesting successful epistemic capture within these policy domains.⁶

Temporal analysis of community paths (see Appendix Community Paths) reveals distinct periods of accelerated specialization, characterized by the simultaneous emergence of multiple specialist domains and increased community integration into existing paths. Three critical junctures emerge with particular salience: The late 1960s witnessed the crystallization of multiple specialist domains, incorporating emergent policy areas including public housing, municipal governance, wage policy, and social welfare into persistent community structures. The second peak of epistemic capture appears around 1980. This period witnesses the emergence of specialist paths centered on constitutional reform, nuclear waste management, and business subsidies—issues that exemplify the increasingly technical nature of political discourse. A third wave of specialization crystallizes circa 1990, marked by the formation of distinct specialist communities around social welfare policy, institutionalized consultation processes, and public transportation

infrastructure. These emergent paths reflect both the evolution of policy priorities and the progressive technocratization of political debate. Our analysis of node neighborhoods and paths of connected communities indicates that the macro-level trend inhabits meso-level fluctuations, apparent in epistemic “resets” that follow a change in cabinet and distinct periods of relatively intense capture. These methods thus complement the macro-level trend of gradual specialization with an overview of thematic foci and temporal oscillations.

Micro-Level Specialization: Linking Epistemic Structures to Political Reality

The combined macro- and meso-level analyses of specialization as epistemic capture can be further enriched by a micro-level analysis of specific topics and actors. This level reveals how broad structural trends materialize in concrete domains of debate and political practice. First, our earlier examination of neighborhood overlap allows us to trace the specialization trajectories of individual topics. These trajectories illustrate the interaction between institutional (macro) specialization and political (meso) dynamics (illustrated in the upper panels of Figure 4). The “public broadcasting” topic shows consistently low neighborhood similarity, reflecting its entanglement in multiple controversies and policy domains—an epistemically volatile field resistant to stable capture.

By contrast, international conflict exhibits persistently high neighborhood similarity. Despite periodic perturbations triggered by international events, the topic’s stability reflects the epistemic consolidation of foreign policy discourse: interpretative frameworks were maintained by established specialists who shaped debate across shifting geopolitical contexts.

A third case, environmental management, demonstrates gradual stabilization. Beginning in the late 1960s, this topic persistently clustered with related issues such as environmental pollution and nuclear waste management, indicating a trend toward epistemic stabilization. This process of specialization was formally recognized in 1971, when a permanent committee for environmental hygiene was established, following proposals initiated by the progressive liberal party D66.

Taken together, these trajectories show that epistemic capture is not an autonomous, linear process. Instead, it unfolds as a dynamic interplay between institutional change, political events, and ideological struggles. Topics may crystallize into specialized domains, but they can also undergo destabilization when contested. Macro-level capture thus manifests through contingent, politically shaped micro-level dynamics.

Identifying Specialists

In addition to topic-level specialization, our data enables the identification of specialist politicians. This analysis bridges structural patterns with political agency by examining how epistemic domains are embodied in the activity of individual parliamentarians.

We quantify politician–topic relationships through conditional probability distributions $P(\text{topic}, \text{member})$ for each period. After z -score normalization, we classify specialists as those

whose engagement with a community exceeds one standard deviation above the mean. Linking these results to community paths allows us to trace durable associations between politicians and domains of expertise.

These quantified politician–community associations validate our network-based analysis of specialization by demonstrating concrete manifestations in parliamentary practice. Known specialists, such as Anton Lucas (Catholic People’s Party), Jur Mellema (Christian Historical Union), and Ad Hermes (Christian Democratic Appeal) appear as strongly bound to respectively finance, agriculture, and education. Confirming their status as experts, these politicians demonstrate strong and persistent links to particular paths of specialist communities. Besides profiling known specialists, this micro-level analysis also allows us to track lesser known specialists and specialisms. In Table 1 we present examples of speakers with high overlap scores between communities, illustrating the method’s capacity to address new questions regarding politicians’ careers, choices, and areas of interest. For instance, Piet Elfferich consistently engages with municipal affairs throughout his parliamentary tenure—a focus that diverges from his formal specialist role on the permanent agriculture committee and reflects his prior experience as a decade-long alderman in the city of Delft.

By linking epistemic structures to political behavior, this micro-level analysis opens new avenues for research. Future work could combine this approach with roll-call voting data or committee memberships to deepen our understanding of how specialization shapes not only discourse but also institutional decision-making.

Table 1. Examples of politicians with strong and persistent links to community paths. The table shows the five speakers with the highest mean overlap between their associated communities and paths of communities. The topic most often found in the community path is displayed along with the average overlap score between the speaker’s communities and the path communities.

ID	Speaker	Most Frequent Topic	Av. Overlap
00321	Van den Doel (PVDA)	public housing (7)	0.73
02976	Ybema (D66)	budgetary space 1 (5)	0.65
00375	Elfferich (ARP)	municipal restructuring (7)	0.55
00964	Van Nierop (ARP)	social benefits (9)	0.53
01876	Staneke (DS70)	budget cuts (9)	0.49

Analysis of micro-level dynamics, such as the persistence and dissolution of specialist paths and the crystallization of expert communities, reveals specialization as a complex, multi-faceted process. The temporal patterns demonstrate epistemic capture operating through multiple mechanisms: the sustained persistence of specialized communities and the rapid emergence of new specialist domains at critical historical junctures. Our network-based methodology thus illuminates both the systematic and contingent aspects of epistemic capture, revealing how structural transformations in parliamentary discourse emerge through the interaction of institutional frameworks, political dynamics, and individual specialization trajectories.

Conclusion

Our analysis of Dutch parliamentary proceedings (1945–1995) highlights the deep entanglement between epistemic capture and parliamentary specialization. Using an integrated framework that combines computational and network-based methods, we show how specialization unfolds differently across macro-, meso-, and micro-levels, each exposing a distinct dimension of epistemic transformation in parliamentary discourse.

At the macro-level, we find three concurrent trends: declining network density, increasing modularity, and greater clustering stability. Most evident in the 1950s and 1960s, these developments reflect the rise of topical specialization alongside institutional reforms such as the establishment of permanent committees. Crucially, our focus on plenary debates reveals that this shift extended beyond procedural innovation, signifying a deeper reorganization of parliamentary discourse itself.

At the meso-level, we uncover how political transitions act as moments of epistemic reorganization. Cabinet changes consistently produce sharp declines in neighborhood stability, followed by gradual reconsolidation during each cabinet's tenure. These fluctuations suggest that shifts in political leadership destabilize established communities of expertise before new epistemic alignments take shape.

The fact that these patterns emerge in topic linkage networks—structures independent of formal parliamentary design—underscores that political transitions trigger substantive epistemological reconfigurations rather than merely procedural adjustments. Leadership change thus reshapes not only who governs but also how knowledge is structured and debated within the legislature.

Our multi-level analysis advances theoretical and methodological understanding across three distinct scholarly domains: political science, the study of historical discourse, and computational humanities. In political science, we show how specialization and epistemic capture reshape knowledge production and decision-making in democratic institutions, revealing the complex interplay between expertise, institutional structures, and deliberation. Our findings empirically substantiate tensions long theorized but rarely demonstrated: the trade-off between democratic pluralism and the technocratic imperatives of expertise and efficiency. Put succinctly, as specialized knowledge deepens, the scope of democratic deliberation narrows.

For the study of historical political discourse, our analysis demonstrates that specialization transforms not only the content of debate but also the mechanisms of idea formation and circulation. The contrast between ideas forged in pluralistic, ideologically diverse debates and those emerging within specialist communities is central to understanding twentieth-century intellectual development.

Methodologically, we show how computational approaches can move beyond conventional lexical analysis to capture structural transformations in intellectual history. The integration of topic modeling with network analysis opens new perspectives on the institutional and social contexts that shape knowledge

production.

In the computational humanities, our use of network analysis provides an alternative to the conventional divide between distant and close reading. By capturing both systematic patterns and contextual variation across multiple scales, network-based approaches enable a more integrated analysis of regularities and contingencies in historical discourse. While our network-based approach yields significant insights, future research could productively complement these findings through advanced computational methodologies. Agent-based modeling (ABM) remains a valuable direction, with simulations incorporating detailed representations of political actors, policy frameworks, and institutional constraints offering potential to illuminate the micro-foundations of specialization processes. Also, the emergence of transformer architectures and large language models presents compelling new avenues for this research agenda. These technologies could enhance ABM approaches by providing more sophisticated natural language processing capabilities for analyzing policy documents, legislative texts, and political communications at scale. Additionally, LLMs could serve as computational agents themselves, simulating diverse stakeholder perspectives and policy positions. The ability of transformers to capture complex contextual relationships and generate human-like reasoning patterns could prove particularly valuable for modeling the subtle dynamics of epistemic capture, potentially revealing how information asymmetries and specialized knowledge create feedback loops that reinforce institutional specialization. Such AI-augmented methodologies could bridge the gap between our macro-level network findings and the micro-level decision-making processes that generate the observed patterns of specialization in democratic institutions.

Acknowledgments We are grateful for the valuable comments and constructive feedback provided during the Computational Humanities Research 2024 conference.

Funding Statement This research was supported by the Leiden Centre for Digital Humanities.

Competing Interests MW is also an associate editor at the Computational Humanities Research Journal.

Data Availability Statement The topic models used in this article are published on Zenodo: 10.5281/zenodo.16754330. Code can be found on: <https://github.com/rubenros1795/linkage>

Ethical Standards The research meets all ethical guidelines, including adherence to the legal requirements of the study country.

Author Contributions A1: Conceptualization, Data Curation, Methodology, Software, Visualization, Writing; A2: Conceptualization, Methodology, Writing, Supervision. All authors approved the final submitted draft.

Notes

- 1 The higher linguistic complexity found by Vodová and Voda Vodová and Voda (2024) in specialist committee meetings can be taken as a reflection of this more technocratic language of specialists.
- 2 Latent Dirichlet Allocation (LDA) is a probabilistic model that assumes documents are mixtures of topics, where a topic is a probability distribution over words. We chose LDA over more recent methods like Top2Vec and BERTopic because our goal was to produce stable and interpretable document-topic distributions suitable for a longitudinal network analysis. Newer embedding-based models often create highly dense representations that can obscure the boundaries between topics. While other models like Structural Topic Modeling (STM) and Supervised Topic Models (sLDA) offer useful extensions, they require incorporating metadata or pre-specifying outcomes, which contradicts our aim of inductively tracing emergent patterns over time.
- 3 Diachronic entropy levels reported in Figure 5 show variation in topic diversity, but this variation seems not to affect our macro-level modularity, density and stability scores in Figures 2 and 3.
- 4 Renormalization was performed by dividing each topic distribution by its sum.
- 5 We calculate the individual TNS score per topic by comparing the topic's neighborhood to four preceding neighborhoods. We then average the topic-level TNS scores to calculate an overall indication of neighborhood stability.
- 6 We calculated this as the number of periods the same topic occurs in a chain, which yields a topic-level score that indicates the persistence of a topic in a path.

References

- Andeweg, Rudy B. 1997. Role specialisation or role switching? dutch mps between electorate and executive. *The Journal of Legislative Studies* 3 (1): 110–127.
- Aylett, Philip. 2019. Crossman and beyond: house of commons select committees in the 1960s. *Parliamentary History* 38 (3): 408–430.
- Blei, David M., Andrew Y. Ng, and Michael I. Jordan. 2003. Latent dirichlet allocation. *Journal of machine Learning research* 3 (1): 993–1022.
- Blondel, Vincent D., Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefebvre. 2008. Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment* 2008 (10): P10008.
- Bohman, James. 1999. Democracy as inquiry, inquiry as democratic: pragmatism, social science, and the cognitive division of labor. *American Journal of Political Science* 43 (2): 590–607.
- Bolla, Peter de, Ewan Jones, Paul Nulty, Gabriel Recchia, and John Regan. 2020. The idea of liberty, 1600–1800: a distributional concept analysis. *Journal of the History of Ideas* 81 (3): 381–406.
- Buchanan, William, Heinz Eulau, LeRoy C. Ferguson, and John C. Wahlke. 1960. The Legislator as Specialist. *The Western Political Quarterly* 13 (3): 636–651.
- Danon, Leon, Albert Díaz-Guilera, Jordi Duch, and Alex Arenas. 2005. Comparing community structure identification. *Journal of Statistical Mechanics: Theory and Experiment*, no. 9, 1–10.
- Dewey, John. 2012. *The public and its problems: an essay in political inquiry*. University Park: Penn State University Press.
- Durkheim, Emile. 2014. *The division of labor in society*. New York: Simon / Schuster.
- Guldi, Jo. 2019. Parliament's debates about infrastructure: an exercise in using dynamic topic models to synthesize historical change. *Technology and Culture* 60 (1): 1–33.
- Haas, Peter M. 1992. Introduction: epistemic communities and international policy coordination knowledge, power, and international policy coordination. *International Organization* 46 (1): 1–36.
- Habermas, Jürgen, and Jeremy J. Shapiro. 1970. The scientization of politics and public opinion. In *Toward a rational society: student protest, science, and politics*, 62–80. Boston: Beacon Press.
- Heuser, Ryan. 2023. Computing Koselleck: Modelling Semantic Revolutions, 1720–1960. In *Explorations in the Digital History of Ideas*, edited by Peter De Bolla, 256–285. Cambridge University Press.
- Hoetink, Carla. 2018. *Macht der gewoonte, regels en rituelen in de tweede kamer na 1945*. Nijmegen: Vantilt.
- Karel, Erwin H. 2013. Boer en politiek, opkomst en teloorgang van het groene front. *Historia Agriculturae* 44:31–56.
- Koole, Ruud. 2018. Gouvernentalisering: De veranderende verhouding tussen regering en parlement in Nederland. *Tijdschrift voor Constitutioneel Recht* 2018 (5): 317–341.
- Koß, Michael. 2018. *Parliaments in Time: The Evolution of Legislative Democracy in Western Europe, 1866–2015*. Oxford University Press.
- Kroeze, Ronald, and Sjoerd Keulen. 2020. “More exciting than Watergate, more real than Dynasty”. Transparency's rise: The Dutch RSV-inquiry in the context of the 1980s. In *History of Transparency in Politics and Society*, edited by Jens Ivo Engels and Frédéric Monier, 137–154. Göttingen: V&R Unipress.
- Lefort, Claude. 1988. *Democracy and political theory*. Minneapolis: University of Minnesota Press.
- Lijphart, Arend. 1968. *Verzuiling, pacificatie en kentering in de Nederlandse politiek*. Haarlem: Becht.
- Marjanen, Jani. 2023. Quantitative conceptual history: on agency, reception, and interpretation. *Contributions to the History of Concepts* 18 (1): 46–67.
- McCallum, Andrew Kachites. 2002. *Mallet: a machine learning for language toolkit*. <http://mallet.20cs.20umass.20edu>.
- McDaid, Aaron F., Derek Greene, and Neil Hurley. 2013. *Normalized mutual information to evaluate overlapping community finding algorithms*. ArXiv preprint. arXiv: arXiv:1110.2515.
- Meissel, Kane, and Esther S Yao. 2024. Using cliff's delta as a non-parametric effect size measure: an accessible web app and r tutorial. *Practical Assessment, Research, and Evaluation* 29 (1).
- Mitchell, James Clyde. 1969. The concept and use of social networks. In *Social networks in urban situations*, 1–50. Manchester: Manchester University Press.
- Moretti, Franco. 2013. *Distant reading*. London: Verso Books.
- Newman, Mark E. J. 2006. Modularity and community structure in networks. *Proceedings of the National Academy of Sciences* 103 (23): 8577–8582.
- Palonen, Kari. 2015. Parliamentarism Challenged. In *Challenges to parliamentary politics: rhetoric, representation and reform*, edited by Suvi Soininen and Tuula Vaarakallio, 17–36. Baden-Baden: Nomos Verlagsgesellschaft.
- Pantel, Patrick, and Dekang Lin. 2002. Discovering word senses from text. In *Proceedings of the eighth acm sigkdd international conference on knowledge discovery and data mining*, 613–619. Association for Computing Machinery.
- Perry, Chloe, and Simon DeDeo. 2021. *The cognitive science of extremist ideologies online*. ArXiv preprint. arXiv: arXiv:2110.00626.
- Poulsen, Victor Møller, and Simon DeDeo. 2023. Inferring cultural landscapes with the inverse Ising model. *Entropy* 25 (2).
- Rossetti, Giulio, Letizia Milli, and Remy Cazabet. 2019. Cdlib: a python library to extract, compare and evaluate communities from complex networks. *Applied Network Science* 4 (52).

- Scanni, Francesco Maria. 2023. Opposites but similar? technocracy and populism in contemporary european democracies. *Administration & Society* 55 (5): 1007–1029.
- Shaw, Malcolm. 1998. Parliamentary committees: a global perspective. *The Journal of Legislative Studies* 4 (1): 225–251.
- Siefken, Sven T. 2021. No paradise of policy-making: the role of parliamentary committees in the german bundestag. In *Parliamentary committees in the policy process*. London: Routledge.
- Szymkiewicz, Dezydery. 1934. Une contribution statistique à la géographie floristique. *Acta Societatis Botanicorum Poloniae* 11 (33): 249–265.
- Te Velde, Henk, and Wim Voermans. 2019. Autoriteit en onderhandelen. Over regenten en polderbestuurders. In *Van driestromenland tot delta? Beschouwingen over ontwikkelingen in de Nederlandse politiek*, edited by Tom Louwerse, Rudy Andeweg, Joop van Holsteyn, and Josje den Ridder, 159–177. Leiden: Leiden University Press. Accessed December 1, 2023.
- Urbinati, Nadia. 2014. *Democracy disfigured: opinion, truth, and the people*. Harvard University Press.
- Van Bouwel, Jeroen. 2023. Strengthening the epistemic case against epistocracy and for democracy. *Social Epistemology* 37 (1): 110–126.
- Van Eijnatten, Joris, and Pim Huijnen. 2021. Something happened to the future: reconstructing temporalities in dutch parliamentary debate, 1814–2018. *Contributions to the History of Concepts* 16 (2): 52–82.
- Van Lange, Milan. 2023. *Emotional imprints of war: a computer-assisted analysis of emotions in dutch parliamentary debates, 1945–1989*. Bielefeld: Bielefeld University Press.
- Vibert, Frank. 2007. *The rise of the unelected: democracy and the new separation of powers*. Cambridge: Cambridge University Press.
- Vicinanza, Paul, Amir Goldberg, and Sameer B. Srivastava. 2023. A deep-learning model of prescient ideas demonstrates that they emerge from the periphery. *PNAS Nexus* 2, no. 1 (January 1, 2023).
- Vodová, Petra, and Petr Voda. 2024. Practice makes ‘perfect’? The effect of committee specialization on the complexity of parliamentary speeches. *Parliamentary Affairs* (August 8, 2024): 1–25.
- Zubčić, Marko-Luka, and Gabriele Giacomini. 2024. Beyond ‘infodemic’: complexity, knowledge and populism in covid-19 crisis governance. *Social Epistemology* 0 (0): 1–21.

Topic Models

The table in this appendix contains the labels and top terms for a rhetorical, procedural, and “thematic” topic. A full list of topic labels can be found in the GitHub repository.

Label	Top Terms
Rhet/Appreciatio	appreciation, great statesman, good, wise, word, work, faction, on which to speak, hope, year, special, heart, agreement, start, new, gladly, trust, thank
Proc/Meeting Reports	report, provisional, answer, memorandum, member, remark, bill, government, reason, opinion, consider, gladly, minister, general, explanation, different, point, relation
European Community	European, country, politics, cooperation, international, economic, community, integration, common, national, territory, relation, foreign, large, unit, union, development, treaty, government

Entropy of Topic Distributions

To investigate whether our macro-level signals are statistical artifacts deriving from the overall diversification of parliamentary language, we report the average entropy of topic distributions over time (Figure 5). The higher the entropy, the greater the diversity of the system. The trend appears volatile, with a slight decline between 1950 and 1985 followed by a growth in entropy. This trend does not confirm the idea of an overall rise in topic diversity.




Figure 5. Average Shannon Entropy of topic distributions over time.

Effect Sizes

Cliff’s delta is a non-parametric effect size measure that quantifies the degree of difference between two groups by estimating the probability that a randomly selected value from one group will be higher than one from the other. It is robust to non-normal distributions and unequal variances, making it well suited for ordinal or skewed data (Meissel and Yao 2024). Values range from -1 to 1 , with 0 indicating no difference and larger absolute values reflecting stronger effects. In this analysis, we use Cliff’s delta to estimate the effect size of changes in modularity over time. Since modularity is not normally distributed, Cliff’s delta provides a more appropriate alternative to parametric measures such as Cohen’s d .

Linkage function

Given a matrix $\theta \in \mathbb{R}^{N \times K}$, where N is the number of documents and K is the number of topics, the function computes mutual information between topics.

- θ_{ij} : Document–topic mixture value for document i and topic j .

Steps:

1. Joint Probability Calculation:

$$P(i, j) = \frac{\sum_d \theta_{di} \cdot \theta_{dj}}{\sum_{i,j} \sum_d \theta_{di} \cdot \theta_{dj}}$$

2. Marginal Probability Calculation:

$$P(i) = \frac{\sum_d \theta_{di}}{\sum_i \sum_d \theta_{di}}$$

3. Mutual Information:

$$R_{ij} = \log_2 \left(\frac{P(i,j)}{P(i) \cdot P(j)} \right)$$

4. Smoothing Function:

$$\text{weight}(i,j) = \left(\frac{P(i,j)}{P(i,j) + 1} \right) \left(\frac{\min(P(i), P(j))}{\min(P(i), P(j)) + 1} \right)$$

5. Smoothed Mutual Information:

$$\text{SMI}_{ij} = \text{weight}(i,j) \times R_{ij}$$

Community Paths

Community paths are consecutive links (paths) between similar communities in topic linkage networks. In Figure ??, we show the connections between topic communities in the period between January 1970 and January 1975. Lines indicate connections between communities in different periods. Communities without any matches in next periods have no outgoing lines. Paths related to foreign policy (blue), education (red), and agriculture (yellow) are highlighted.

