What Survives Legalisation?
Governance, Enforceability, and the Filtering of Stakeholder Beliefs in the EU AI Act

David Lake

s3408639

Faculty of Governance and Global Affairs, Leiden University

Thesis MSc Public Administration: International and European Governance

Supervisor: Hsini Huang

Date: January 9th, 2026

Word count: 9996

Abstract:

This thesis examines how stakeholder beliefs expressed during the European Commission's 2020 consultation on artificial intelligence relate to the content of the EU Artificial Intelligence Act adopted in 2024. Drawing on the Advocacy Coalition Framework, it analyses consultation submissions as structured expressions of competing problem definitions rather than isolated organisational positions. The analysis finds that disagreement in EU AI governance clusters around coherent belief configurations cutting across industry, civil society, and public authorities. As policy moves from agenda-setting to binding law, beliefs articulated in governance-operable terms such as regulatory design, oversight, and administrative responsibility, persist more visibly in legislative text than broad normative or strategic claims. This pattern reflects the institutional demands of EU legalisation, which privilege enforceable obligations and administratively workable procedures. The findings support institutional filtering as a plausible explanation for why extensive consultation yields selective legislative outcomes, raising broader questions about how participation translates to representation in EU technology governance.

Table of Contents

1. Introduction

The European Union's Artificial Intelligence Act represents the world's first comprehensive attempt to regulate artificial intelligence through binding horizontal legislation. Introduced in 2021 and adopted in 2024, it seeks to reconcile competing objectives: protecting fundamental rights, ensuring safety and accountability, and supporting innovation in a rapidly evolving technological field. Given AI's technical complexity and societal stakes, the Commission placed exceptional emphasis on stakeholder consultation during agenda-setting, soliciting hundreds of submissions from industry, civil society, public authorities, and expert organisations.

Despite this extensive participation, the AI Act has attracted criticism for producing regulatory outcomes that appear narrower, more technocratic, and more governance-focused than the diversity of concerns expressed during consultation would suggest. Broad ethical debates, social visions of AI, and strategic innovation narratives featured prominently in consultation responses but are less visible in the final legislative text, which centres on risk classification, compliance obligations, and administrative oversight. This gap between consultation input and legislative output raises a fundamental question about how participation translates into law in EU digital governance.

This gap motivates the thesis, which argues that it reflects systematic institutional filtering rather than consultation failure or lobbying capture. As policy moves from agenda-setting toward binding regulation, legalisation constrains which claims can be translated into enforceable obligations, producing selective persistence of beliefs compatible with EU legal form.

EU policymaking scholarship has long emphasised that consultation serves functions beyond preference aggregation, including information gathering and legitimacy construction. Within the EU policy cycle, consultation inputs are processed through impact assessments, inter-service coordination, and legal drafting procedures that progressively narrow admissible claims. This study provides empirical evidence that this filtering operates in EU AI governance, shaping which stakeholder beliefs survive the transition from consultation to legislation.

This argument builds on research showing that policy outcomes in complex regulatory domains are shaped less by actor dominance than by the structure of ideas during agenda-setting. Tallberg and colleagues demonstrate that belief congruence, rather than resource

asymmetry alone, plays a decisive role in EU policy influence. This finding motivates the analytical approach adopted here: if legislative outcomes reflect belief selection rather than actor selection, the mechanism lies in how beliefs translate into legal form, not in participation intensity.

To examine this mechanism, the thesis adopts the Advocacy Coalition Framework (ACF), which conceptualises policy conflict as competition between coalitions organised around shared belief systems. The ACF distinguishes deep normative commitments, policy core beliefs, and secondary beliefs concerning implementation and governance design. These distinctions are well suited to analysing legalisation, as EU regulatory institutions privilege beliefs that can be rendered as administratively workable obligations while filtering out others.

Accordingly, the thesis tests whether institutional filtering explains consultation–legislation gaps in EU AI governance by addressing two questions. First, how do stakeholder beliefs expressed during consultation cluster into advocacy coalitions, and do these coalitions map onto organisational identities? Second, do governance-expressible beliefs show systematically higher alignment with binding legislation than normative or strategic beliefs as policy moves toward legal adoption?

The thesis makes three contributions. First, it shows that consultation-stage conflict in EU AI governance is organised around belief structures that cut across organisational categories, challenging simple interest-group explanations. Second, it demonstrates that beliefs articulated as governance obligations align more closely with legislative outcomes than expansive normative or strategic claims, consistent with institutional filtering theory. Third, it shows that extensive participation does not guarantee representative translation into law, raising fundamental questions about the democratic function of consultation in technically complex policy domains.

The remainder of the thesis proceeds as follows. Chapter 2 situates the study within EU governance and policy process literature. Chapter 3 outlines the research design and analytical strategy. Chapters 4 and 5 present the empirical analysis of coalition structure and belief translation. Chapter 6 concludes by discussing implications for democratic representation and consultation design in EU digital governance.

2. Literature review and theoretical framework

Conflict over the Artificial Intelligence Act is often framed as a struggle between familiar camps: industry versus civil society, innovation versus regulation. That framing captures something real, but it is incomplete. Disagreement in EU AI governance does not organise itself neatly along organisational lines. It takes the form of structured conflict over policy beliefs that change as policy moves toward binding law. What matters most is not who participates in consultation, but which belief configurations remain viable under legalisation.

AI governance is not a technical exercise insulated from politics. It is shaped by persistent trade-offs between economic competitiveness, risk mitigation, and rights protection. These trade-offs resist resolution through expertise alone. In the EU context, they are intensified by an institutional setting that invites wide-ranging input early on but requires regulatory outputs to satisfy demanding legal and administrative standards. Many claims can be articulated during consultation. Far fewer can be stabilised in regulatory form.

To analyse this process, the chapter draws on three bodies of literature. First, it treats AI governance as a political and regulatory problem structured by value conflict rather than technical consensus. Second, it examines EU consultation and policymaking as selective institutional settings that condition which arguments remain plausible as policy hardens into law. Third, it introduces the Advocacy Coalition Framework (ACF) as a tool for analysing belief-driven conflict under conditions of strong legal and administrative constraint. Applying the ACF in this context is not seamless, Sabatier developed the framework for long-running policy subsystems rather than compressed regulatory processes, but its focus on belief systems allows for a systematic examination that actor-centred approaches lack.

The discussion remains closely tied to empirical application. The belief dimensions developed here draw directly from AI governance scholarship and are formulated to be observable in consultation submissions and stakeholder position papers. They are not exhaustive, and alternative classifications are possible. But they capture recurring lines of disagreement visible across the consultation record. This grounding allows later chapters to trace how belief configurations cluster during consultation and how their correspondence with policy content changes as regulation moves from agenda-setting to law.

2.1 Artificial Intelligence Governance as a Political and Regulatory Problem

AI governance poses a policy challenge because AI systems spread rapidly across sectors and operate in decision-making contexts not designed for automation. Systems used in employment screening, healthcare allocation, policing, or migration management raise more than questions of technical reliability. They generate disputes over discrimination, accountability, opacity, and democratic control (Helberger & Diakopoulos, 2023; Krarup & Horst, 2023). These are not peripheral concerns. They place AI regulation squarely in the realm of political contestation, where decisions about acceptable risk and responsibility have distributive and normative consequences.

Policymakers confront trade-offs that resist resolution through expertise alone. Promoting innovation and economic competitiveness often conflicts with protecting fundamental rights or limiting systemic risk (Krarup & Horst, 2023). In the EU, these tensions are sharpened by a legal order that combines market integration with strong rights commitments. The Commission's White Paper on Artificial Intelligence reflects this dual orientation, framing AI as both a strategic economic opportunity and a source of risk requiring regulation (European Commission, 2020). Gorwa et al. (2024) describe this as a constitutional approach to digital governance, in which regulation translates foundational legal principles into technological oversight. This framing does not resolve conflict; it displaces it. Disagreement persists over how AI should be defined as a regulatory object, where boundaries should be drawn, and which forms of expertise count as authoritative. These choices shape regulatory scope, compliance obligations, and enforcement capacity, privileging some belief systems over others (Krarup & Horst, 2023). AI governance, then, is not technical problem-solving but an ongoing political process structured by competing beliefs about innovation, risk, and institutional responsibility.

2.2 EU Policymaking, Consultation, and Institutional Selectivity

EU policymaking shapes AI governance in ways that differ from many national regulatory settings. Authority is dispersed across institutions and levels, and policy development relies heavily on expert input, impact assessment, and public consultation, particularly in technically complex areas such as AI (Hix & Høyland, 2022). Consultation serves several purposes. It supplies information, helps manage uncertainty, and lends legitimacy to regulatory initiatives. It also creates a structured arena in which competing policy beliefs are articulated and evaluated against institutional expectations.

Research on EU interest representation complicates the view of consultation as simple preference aggregation. Dür (2008) and Bunea (2013) show that consultations are formally open but substantively selective, favouring actors with organisational resources and technical expertise. Resource advantages, however, do not translate automatically into policy dominance. In policy areas governed by strong legal mandates and entrenched normative commitments, institutional constraints narrow the range of claims that remain plausible. Well-resourced actors may articulate positions more effectively, but they still encounter limits imposed by legal form, administrative capacity, and established regulatory principles.

Discursive institutionalism helps explain how this selectivity operates. Schmidt's framework highlights how institutions shape which ideas appear plausible and legitimate within policymaking, not only which actors gain access (Schmidt, 2008; 2010). It leaves open a further question: how do ideas fall away as policy develops? In the EU context, coordinative discourse among administrative and political actors requires that claims be rendered legally precise and administratively workable. Ideas that resist this translation because they remain broad or strategically vague become increasingly difficult to carry forward as policy approaches legislation.

This dynamic aligns with the operation of the EU's Better Regulation framework as a form of meta-regulation. Better Regulation structures not only how policy options are assessed, but which types of arguments remain viable. It privileges claims expressed through causal reasoning, proportionality, and administratively operable rules, while attaching less weight to broad normative or strategic claims that resist legal specification (Radaelli & Meuwese, 2009). As policy development moves toward binding law, these requirements shape which stakeholder beliefs can be expressed in forms compatible with legal certainty, enforceability, and administrative oversight.

2.3 The Advocacy Coalition Framework under Conditions of Legalisation.

The Advocacy Coalition Framework (ACF) provides a way to analyse policy conflict in complex and contested domains such as AI governance by shifting attention from organisational categories to shared belief systems. Developed by Sabatier and Jenkins-Smith (1993), the framework conceptualises policy subsystems as arenas structured by disagreement over problem definitions, causal beliefs, and appropriate regulatory responses. Actors align into advocacy coalitions based on these shared understandings rather than formal affiliation,

making the framework well suited to consultation environments where conflict cuts across sectors.

A central feature of the ACF is its distinction between belief types. Deep core beliefs capture stable normative orientations, such as views on the role of markets or the legitimacy of state intervention. Policy core beliefs concern domain-specific problem definitions and governance principles and form the primary basis for coalition alignment. Secondary beliefs relate to implementation choices and technical design and are more open to adjustment. This hierarchy matters analytically because conflict over governance architecture and regulatory design is more tractable than conflict over fundamental values, particularly under conditions of legalisation.

This study captures a compressed empirical window: consultation-stage belief articulation and the final legislative outcome. This limits the ability to observe sustained coalition interaction, belief updating, or policy-oriented learning of as emphasised in classical ACF applications (Sabatier & Jenkins-Smith, 1993). Accordingly, the framework is not employed as a causal model of coalition competition or policy change, but as an analytical device for comparing belief configurations under legalisation.

In this study, the ACF supplies the belief taxonomy that enables systematic analysis of consultation data without assuming coordinated advocacy or coalition success. Coalition identification through clustering serves a descriptive and diagnostic purpose: it reveals how beliefs cluster during consultation rather than how actors organised strategically. This belief hierarchy makes it possible to enalyse whether different belief types correspond unevenly with binding legislation.

The ACF is therefore not employed as a causal account of influence. Instead, it is used in conjunction with institutional accounts of EU policymaking, particularly work on Better Regulation and discursive constraint (Radaelli & Meuwese, 2009; Schmidt, 2010), to interpret correspondence between consultation-stage beliefs and legislative content. The analysis thus asks not which actors prevailed, but which belief types remained compatible with the legal and administrative requirements of binding EU regulation as policy moved from agenda-setting to law

2.4 Policy Beliefs in EU AI Governance

Applying a belief-centred approach to EU AI governance requires specifying which policy core beliefs organise conflict in this domain. Existing research shows that disagreement does

not scatter across countless idiosyncratic concerns. Instead, it clusters around a limited set of recurring dimensions that structure how stakeholders understand both the problems posed by AI and the appropriate regulatory response.

One central axis concerns the balance between innovation and competitiveness on the one hand, and precaution and protection on the other (Krarup & Horst, 2023). Some actors frame regulatory restraint as a condition for technological progress, arguing that uncertainty and compliance burdens risk undermining Europe's position in global AI development. Others argue that strong ex ante safeguards are necessary precisely because AI systems scale rapidly and can generate harms that are difficult to reverse. This tension is familiar, but it does not exhaust the structure of disagreement.

Beliefs about accountability, oversight, and transparency introduce additional layers. Research on algorithmic governance documents persistent disputes over responsibility gaps created by automated decision-making, with disagreement over whether human oversight should function as a minimal procedural check or as a substantive requirement for meaningful control (de Boer & Raaphorst, 2021). Transparency is similarly contested. While most stakeholders endorse it in principle, they diverge sharply over what it should entail, ranging from technical documentation and auditability to broader forms of democratic contestability and public scrutiny (Helberger & Diakopoulos, 2023).

Conflict also centres on governance capacity and institutional design. Effective regulation depends on administrative resources, coordination across authorities, and the role of standard-setting bodies, yet actors disagree on how much institutional expansion is feasible or desirable (Gorwa et al., 2020). Proportionality and regulatory burden cut across these debates, operating both as legal principles and as political arguments shaping perceptions of legitimacy and enforceability, particularly among business actors (Coen & Richardson, 2009). These concerns are not merely strategic; they reflect different assumptions about what regulation can realistically achieve under existing institutional conditions.

These dimensions allow stakeholder positions to be analysed as structured belief systems rather than as isolated or purely strategic preferences. They provide the basis for identifying advocacy coalitions during consultation and for examining how different types of beliefs—normative commitments, governance preferences, and implementation concerns—correspond unevenly with binding legislative outcomes as policy moves from agenda-setting to law.

2.5 Implications for the Analysis

Within this setting, the Advocacy Coalition Framework (ACF) enables analysis of how beliefs cluster and correspond with policy outcomes without assuming direct influence or lobbying success. Its role here is analytical rather than causal. The framework offers a theoretically grounded belief hierarchy distinguishing deep core, policy core, and secondary beliefs. These different types of beliefs can structure text analysis beyond ad hoc categorisation. It conceptualises coalitions as groupings based on shared problem definitions rather than formal organisational ties, making it well suited to a consultation context where disagreement cuts across sectors. These features allow systematic comparison between consultation-stage belief configurations and legislative content.

If institutional filtering operates as theorised, coalitions articulating beliefs as regulatory design, administrative responsibility, and oversight should show higher alignment with binding legislation than coalitions emphasising broad normative commitments or strategic innovation goals. The analysis is therefore not merely descriptive: it tests whether observed correspondence follows the trajectory predicted by institutional filtering theory.

The ACF is analytically useful because it distinguishes belief types that legal institutions privilege from those they filter out. Used in this way, it enables examination of why some beliefs persist through legalisation while others contract, without attributing outcomes to actor dominance or lobbying capacity. Beliefs framed as obligations, procedures, or enforceable standards are more readily comparable with legislative texts than those expressed as ethical aspirations, strategic visions, or political principles. This distinction reflects institutional requirements rather than normative judgments about claim validity.

Applied as a structuring device rather than a causal model, the ACF avoids two limitations of its original use. First, it does not require demonstrating sustained policy-oriented learning within coalitions, which would demand longitudinal evidence beyond consultation data. Second, it avoids attributing legislative outcomes to coalition success, sidestepping claims about which actors “won” the regulatory contest. Coalition identification captures how beliefs cluster during consultation, not how actors coordinated strategically.

Correspondence between coalition belief profiles and legislative content is therefore interpreted as an outcome of institutional selectivity rather than coalition influence. Beliefs that can be translated into enforceable obligations, procedural requirements, and administrative oversight structures are expected to persist as policy moves from agenda-setting to law, regardless of which actors articulated them. This interpretation draws on

scholarship on Better Regulation and discursive institutionalism (Radaelli & Meuwese, 2009; Schmidt, 2010), which highlights how legal form constrains what can be codified in binding rules. This perspective yields testable empirical predictions. If institutional filtering operates primarily through legal-form constraints, three observable trends should emerge.

First, consultation-stage conflict should be structured around competing problem definitions rather than organisational identities, as legal constraints condition participants in similar ways. Second, coalitions emphasising governance architecture, administrative oversight, and legally enforceable obligations should exhibit greater correspondence with the final text of the AI Act than coalitions focused on broad normative commitments or strategic innovation claims. Third, as the process moves from the White Paper stage to binding regulation, thematic emphasis should shift toward governance-oriented content and away from expansive normative framing.

Chapters 4 and 5 assess whether these patterns are supported by systematic belief analysis. While correspondence alone cannot establish causation, structured alignment across belief types and formulation would suggest institutional filtering rather than consultation failure or actor dominance to shape the translation of consultation inputs into EU AI legislation. Alternative explanations, such as Commission discretion, parliamentary dynamics, or stakeholder adaptation, would need to explain why governance-expressible beliefs persist in the legislative outcome regardless of how extensively other beliefs were articulated during consultation.

3. Research design and methodology

This chapter outlines the research design and methodological approach used to analyse advocacy coalitions and policy movement in European Union (EU) artificial intelligence (AI) governance. It draws on the Advocacy Coalition Framework (ACF) and focuses on the European Commission's 2020 public consultation on Artificial Intelligence and the subsequent adoption of the Artificial Intelligence Act (AI Act). The methodological task is to translate a heterogeneous corpus of stakeholder position papers into a structured, comparable representation of policy beliefs and to examine how variation in belief emphasis corresponds with binding EU legislation.

Consultation is treated as a setting in which actors articulate beliefs prior to full legalisation. The analysis examines which belief configurations correspond more closely with legislative content, treating alignment as an outcome shaped by institutional context rather than as evidence of influence. It therefore does not seek to identify policy "winners," but instead examines which types of claims remain compatible with legally enforceable regulation as policy moves from agenda-setting toward legalisation, and which are narrowed, reframed, or lose prominence.

Stakeholder position papers constitute the primary empirical material. These submissions provide direct insight into how actors articulate regulatory preferences, normative commitments, and problem definitions, but they also pose substantial methodological challenges. The texts are unstructured and vary widely in length, style, and legal sophistication. The central methodological problem is how to extract comparable belief signals from heterogeneous documents without relying on predefined survey categories or impressionistic qualitative interpretation.

Language models and embedding-based approaches were evaluated but deemed incompatible with the analytical and computational constraints of this study. Given the scale and heterogeneity of the corpus, they exceeded available computing capacity and lacked sufficient transparency for operationalising beliefs within the Advocacy Coalition Framework. A deterministic, rule-based approach was therefore adopted, prioritising replicability and theoretical coherence over expressive richness.

The rule-based approach translates the belief taxonomy into explicit coding rules, each mapped to a defined belief dimension and level. Emphasis on interpretability is necessary because coalition membership is inferred from similarity in belief emphasis rather than

observed through sustained interaction or coordinated behaviour. Where coalitions cannot be empirically validated as strategic actors, the underlying belief measures must remain transparent and tractable.

This framing also shapes how correspondence with legislative outcomes is assessed. In the absence of longitudinal evidence of belief adjustment or policymaker response, the analysis cannot evaluate learning dynamics or coalition influence. It can, however, document which belief configurations articulated during consultation correspond more closely with the final AI Act. Explaining these results therefore requires reference to the institutional constraints and selectivity mechanisms developed in Chapter 2, rather than to the internal causal logic of coalition competition.

3.1 Data Sources and Corpus Construction

The analysis relies on three sources: consultation metadata from the European Commission's 2020 public consultation on the White Paper on Artificial Intelligence, stakeholder position papers submitted during that consultation, and the texts of the 2020 White Paper and the Artificial Intelligence Act adopted in 2024.

Consultation metadata are used to identify actors and link organisational attributes to position-paper submissions. Structured questionnaire responses included in the consultation are not used to operationalise beliefs. Reliance on predefined response categories would restrict the analysis to distinctions embedded in the consultation design rather than those derived from theory and textual expression.

Of 440 submitted position papers, 326 were retained after language and content filtering. Exclusions are based on language and minimum content thresholds (see Appendix A.2, Tables A.2.1 and A.2.2 for the full preprocessing pipeline and exclusion criteria). Exclusions occurred before belief extraction. The reduced corpus size ensures that retained documents are comparable in language and substantive content, and metadata linkage, limiting the risk of systematic bias at later analytical stages.

Policy movement is examined using two EU policy documents: the Commission's 2020 White Paper on Artificial Intelligence and the final AI Act. These texts differ substantially in institutional role and genre. Comparability is therefore achieved not by equating absolute text volume, but by assessing relative thematic emphasis within each document using an identical belief extraction framework.

3.2 Operationalising Policy Beliefs in Text

Within the Advocacy Coalition Framework, belief systems are organised hierarchically into deep core beliefs, policy core beliefs, and secondary aspects. Because policy core beliefs structure coalition formation and regulatory conflict, the analysis concentrates on identifying and measuring policy core beliefs as expressed in stakeholder texts. Deep core and secondary beliefs are retained primarily to support interpretation rather than measurement.

Belief dimensions are derived from the AI governance literature reviewed in Chapter 2 and correspond to recurring lines of disagreement in EU AI governance. These include innovation and competitiveness, fundamental rights and social justice, accountability and human oversight, transparency and explainability, governance capacity and regulatory design, proportionality and regulatory burden, and sector-specific risk considerations. The dimensions were refined through close reading of a subset of position papers and cross-checked against existing scholarship. They are not normatively neutral. Instead, they reflect how actors themselves articulate concerns, avoiding the imposition of an external classificatory scheme.

Belief extraction relies on a deterministic, rule-based, sentence-level classification framework. Documents are converted from PDF to text and pre-processed through correction of OCR artefacts, removal of boilerplate material, and formatting normalisation. Texts are segmented into sentences, with only sentences exceeding a minimum length threshold retained. Each retained sentence is evaluated against a manually specified library of regular-expression rules mapped to belief themes and belief levels. Rules are weighted to reflect prescriptive strength, with explicit policy demands receiving higher weights than descriptive statements. When multiple rules apply, the highest-weight match is retained. Sentences that match no rules are excluded.

This approach imposes structure on the text rather than uncovering latent beliefs. Alternative coding choices are therefore possible. For example, a sentence arguing that documentation requirements undermine innovation without improving accountability could plausibly be classified as innovation-related or accountability-related. In this analysis, such cases are coded under proportionality, reflecting an evaluation of regulatory means against expected benefits. This choice is contestable and reflects an explicit interpretive judgment rather than a claim of semantic necessity.

Across the corpus, belief extraction captures a minority of sentences per document. On average, approximately nine per cent of sentences are classified as belief-bearing, with substantial variation across submissions. Appendix A.5 (Tables A.5.1 and A.5.2) reports

summaries and examples of uncoded material. This is intentional. The approach targets explicit normative and policy-relevant claims rather than background description or technical exposition. Restricting analysis to belief-bearing sentences also enables tractable comparison without excessive computational burden. Uniform segmentation and filtering procedures ensure comparability across documents of different length and structure. Appendix A.1 documents the condensed rule library, while Appendix A.3 (Table A.3.1) provides illustrative examples of coded sentences and their associated rule triggers.

Belief profiles approximate expressed emphases rather than exhaustive belief systems. Implicit, highly technical and legally nuanced claims may be underrepresented, a trade-off accepted for transparency and traceability.

3.3 Aggregation, Coalition Identification, and Validation

Belief-bearing sentences are aggregated to construct actor-level belief profiles. Sentence weights are summed by belief theme and normalised to account for differences in document length, producing belief vectors that capture relative thematic emphasis across actors.

Advocacy coalitions are identified using k-means clustering applied to actor-level belief vectors. Belief dimensions are standardised prior to clustering to ensure equal weighting across themes. The number of clusters is specified ex ante. A six-cluster solution is selected as a compromise between analytical differentiation and interpretability, based on internal validity diagnostics and substantive coherence. A five-cluster solution yields a nearly identical structure, indicating that coalition boundaries are present but not sharply delineated.

Coalitions are interpreted by examining average belief profiles, internal dispersion, and actor composition. Labels are assigned according to dominant belief emphases rather than organisational affiliation. Diagnostic measures indicate higher similarity within clusters than between clusters, suggesting that the resulting groupings capture substantively coherent patterns in belief articulation.

The identified coalitions are treated as analytical groupings rather than as coordinated or self-conscious political alliances. Clustering serves as a heuristic for revealing recurrent configurations of belief emphasis in the data, not for asserting fixed coalition membership or strategic coordination.

As a validity check, random subsets of coded sentences were manually reviewed to assess correspondence between assigned belief themes and substantive content.

3.4 Policy Movement and Coalition Alignment

The final stage of the analysis extends belief extraction from consultation inputs to policy outputs. The White Paper and the AI Act are processed using the same belief extraction framework, producing document-level belief profiles based on extracted and weighted belief-bearing sentences.

Policy movement is operationalised as change in relative thematic emphasis between the White Paper and the AI Act Coalition–policy alignment uses cosine similarity to assess thematic correspondence, not causal influence. Differences in correspondence across belief dimensions and coalitions indicate which belief emphases align more closely with binding regulation as policy moves from consultation toward legalisation, and which lose relative prominence under requirements associated with enforceability, legal certainty, and administrative governance capacity. Figure 1 summarises the analytical workflow from corpus construction and sentence-level belief extraction to coalition identification and policy-alignment assessment.

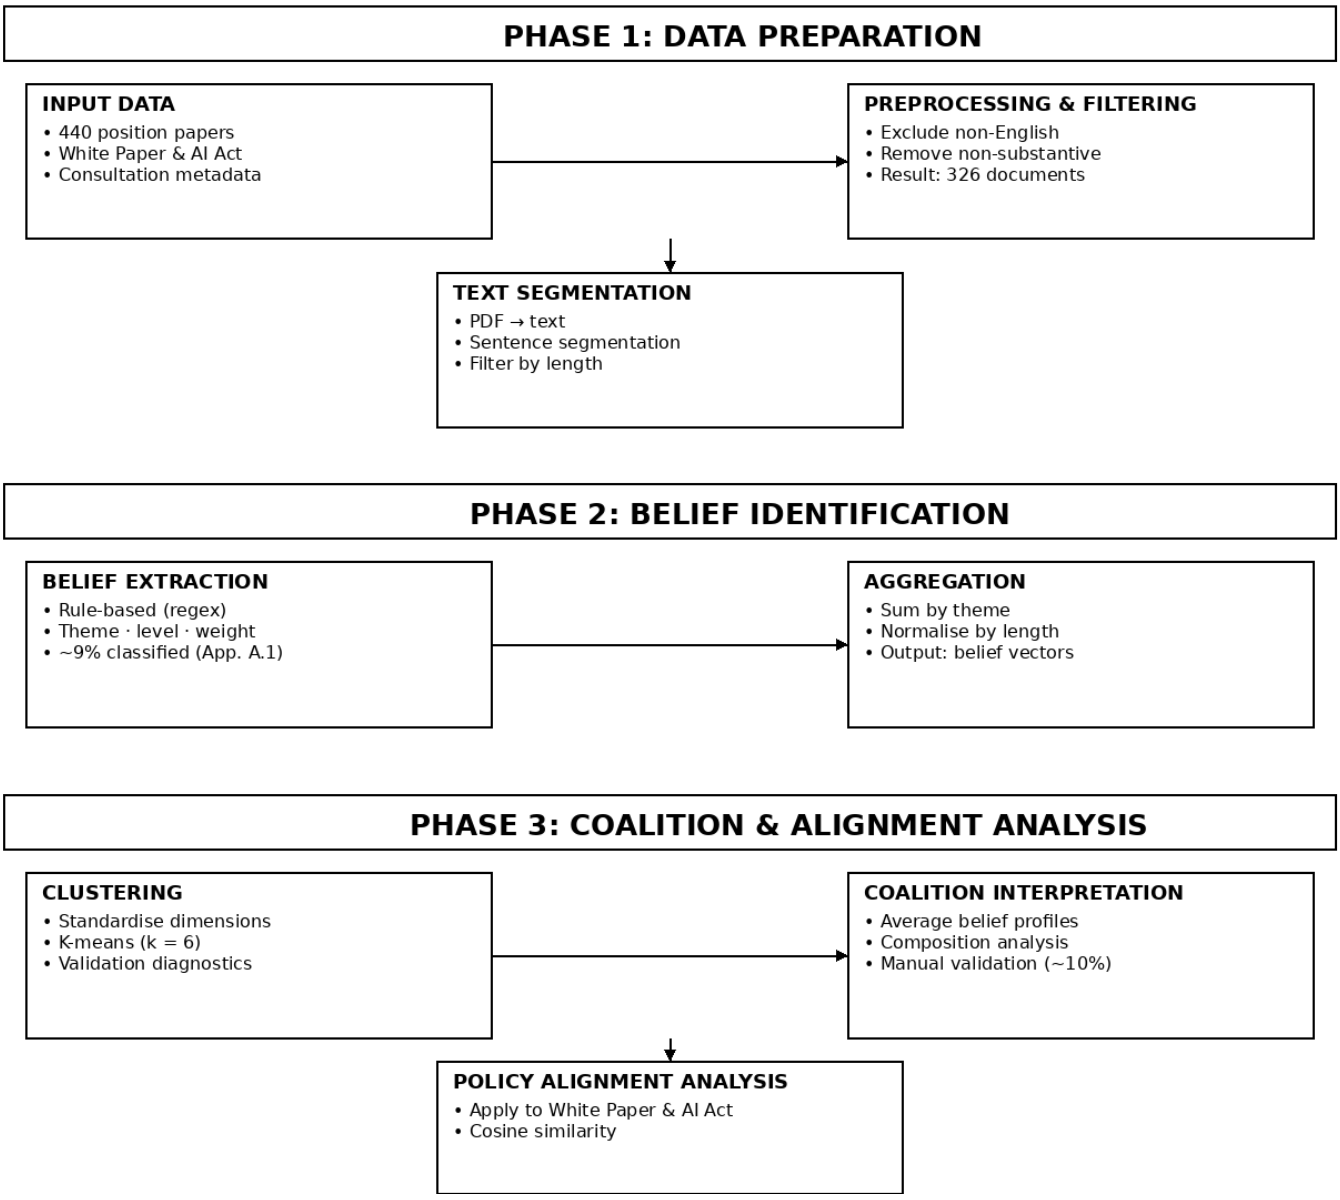

Figure 1. Methodological flowchart: text processing and analytical steps

4. Coalition analysis

Advocacy coalitions are used here as analytical labels for belief-based clusters identified from similarity in consultation position-taking. They do not denote durable alliances, coordination, or collective action. Clustering captures shared problem definitions and regulatory preferences at a single consultation stage, not strategic alignment, learning, or belief change over time.

4.1 Identifying Advocacy Coalitions

Applying k-means clustering to actor-level belief profiles yields a six-coalition solution. As described in Chapter 3, this specification reflects a balance between analytical differentiation and interpretability, informed by internal validity diagnostics. Appendix A.4 reports the full set of diagnostics, including within-cluster sum of squares, silhouette scores, and the Davies–Bouldin index, as well as stability checks across random initialisations. Together they indicate that a six-cluster solution can capture meaningful variation in belief configurations without fragmenting the data into weakly differentiated or analytically redundant groupings.

Although some overlap appears at the margins, the resulting clusters show clear separation in belief space. Multidimensional scaling of actor-level belief profiles indicates higher similarity within coalitions than between coalitions, supporting interpretation of the clusters as belief-based groupings rather than artefacts of the clustering procedure. Coalition sizes vary substantially. The Technical governance-oriented coalition is the largest, followed by the Rights-oriented, Innovation and competitiveness-oriented, General or mixed governance-oriented, Industry-oriented, and Medical and sectoral safety-oriented coalitions. This distribution suggests that certain belief configurations were more prevalent or more consistently articulated during the consultation process.

Coalitions also differ in their degree of internal consolidation. The Rights-oriented, Technical governance-oriented, Innovation and competitiveness-oriented, and Medical and sectoral safety-oriented coalitions exhibit relatively tight clustering, indicating more coherent belief configurations. By contrast, the General or mixed governance-oriented and Industry-oriented coalitions display greater internal dispersion and partial overlap with adjacent coalitions. This suggests that some coalitions correspond to clearly delineated belief packages, while others capture broader spaces combining feasibility-oriented and mixed governance

claims. Figure 2 plots actors in belief space and illustrates the separation of the six advocacy coalitions identified by k-means clustering; the axes represent a distance-preserving projection of high-dimensional belief vectors and carry no substantive interpretation.

Figure 2. Clustering results. Figure shows advocacy coalitions in belief space.

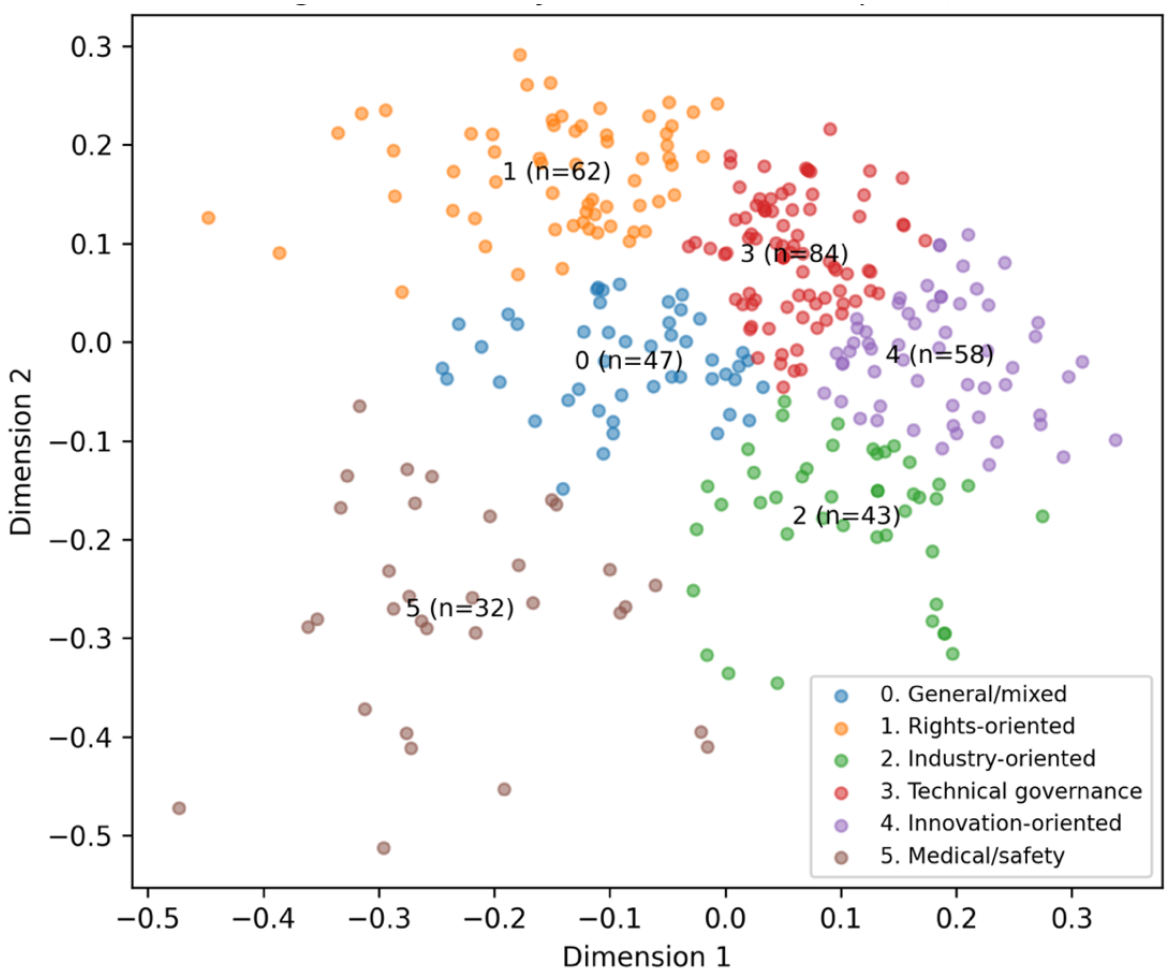

4.2 Coalition Composition

Stakeholder beliefs articulated during the consultation cluster into belief-based groupings that do not map neatly onto organisational categories such as industry, civil society, or public authorities. Although certain coalitions are dominated by particular actor types, each includes actors whose organisational identity alone would not predict membership. This accords with the Advocacy Coalition Framework's expectation that coordination is structured primarily by shared beliefs rather than sectoral affiliation.

As shown in Figure 3, the Rights-oriented coalition is composed mainly of civil society organisations and public-interest groups, with NGOs accounting for approximately 43.5 per cent of documents. Academic institutions and business organisations each contribute

roughly 14.5 per cent, while public authorities are marginally represented. This composition corresponds with the coalition's emphasis on fundamental rights protection and accountability as central regulatory concerns.

The Industry-oriented coalition, by contrast, is dominated by large firms and business associations, which together account for roughly 70 per cent of membership. Companies contribute approximately 41.9 per cent of documents and business associations 27.9 per cent. This composition aligns with a belief configuration centred on proportionality, legal certainty, and workable compliance pathways, reflecting a focus on operational feasibility and uncertainty reduction.

The Technical governance-oriented coalition is the largest and most cross-sectoral. It includes companies, academic and research institutions, business associations, NGOs, public authorities, and governance intermediaries, with companies and academic institutions most prominent. The presence of standardisation bodies and regulatory intermediaries indicates that governance design and implementability constitute a shared problem definition cutting across normative and commercial divides.

The Medical and sectoral safety-oriented coalition is smaller and more specialised, reflecting shared exposure to safety-critical applications and existing regulatory regimes. It includes professional medical associations, regulated-sector firms in healthcare technologies, and patient and public-health organisations. Alignment within this coalition is driven primarily by sector-specific safety concerns and regulatory compatibility rather than by general positions on innovation or regulation.

The Innovation and competitiveness-oriented coalition brings together start-ups, innovation networks, business associations, and selected public bodies concerned with industrial policy. Arguments within this coalition emphasise ecosystem development, competitiveness, and enabling instruments such as sandboxes and testbeds rather than regulatory burden reduction alone.

Finally, the General or mixed governance-oriented coalition comprises actors with more moderate or internally diverse belief profiles. Its composition is relatively balanced across NGOs, companies, business associations, academic institutions, and public bodies, alongside a substantial share of cases with incomplete metadata. This coalition captures integrative positions that combine public-value sensitivity with feasibility and implementation considerations without strong alignment along a single belief axis.

Overall, coalition composition reinforces a core implication of the Advocacy Coalition Framework. Organisational characteristics shape participation patterns, but belief alignment

does not reduce to organisational identity. Instead, consultation-stage discourse reflects multiple belief configurations that cut across sectoral boundaries.

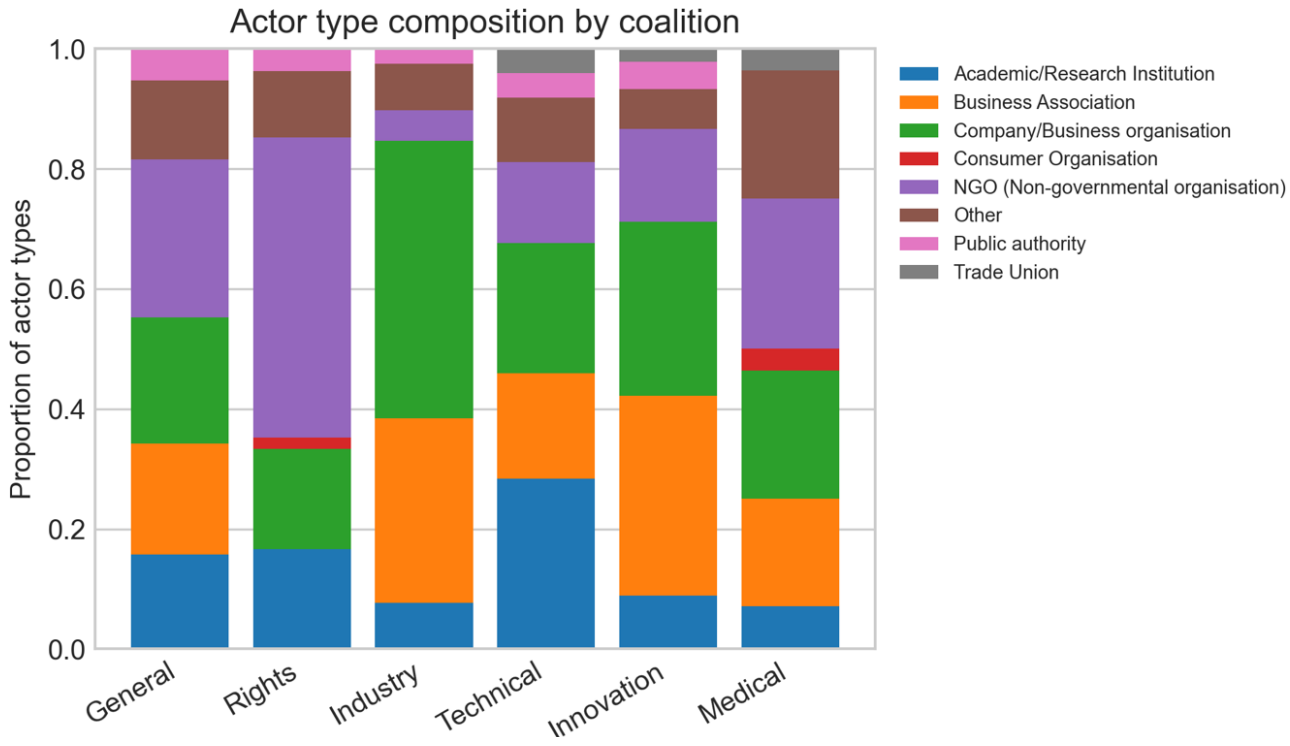

Figure 3. Actor type composition by coalition

4.3 Coalition Belief Profiles

Coalition belief profiles are examined using belief-theme distributions extracted from stakeholder position papers, alongside belief-level composition indicating whether coalitions are anchored primarily in deep core, policy core, or secondary aspects. Across coalitions, belief articulation varies systematically in both thematic emphasis and the level at which contestation is concentrated. Figure 4 visualises the dominant belief indicators for each coalition, illustrating consistent differences in emphasis and belief-level structure.

The Rights-oriented coalition concentrates heavily on fundamental rights protection, discrimination risks, and democratic accountability. Its framing emphasises structural harm and legal safeguards, with a vocabulary saturated by rights language. Governance preferences are anchored largely in deep core commitments rather than administrative design, which helps explain why many of its claims require subsequent specification to fit bounded legal provisions. The Industry-oriented coalition shows a contrasting profile. Proportionality, legal certainty, and compliance feasibility dominate its articulation. Competitiveness appears mainly as a constraint on regulatory scope and cost, while accountability and oversight are framed in terms of role clarity and predictable obligations.

The Technical governance-oriented coalition centres on governance architecture, standards and certification, transparency, and risk assessment. Regulation is treated as a given. Disagreement focuses instead on how regulatory objectives can be operationalised through workable governance arrangements. The Medical and sectoral safety-oriented coalition displays a narrower and more concentrated belief profile. Its emphasis lies on healthcare and safety-critical risk management, including patient safety, clinical validation, and compatibility between horizontal AI regulation and established sectoral regimes, accompanied by sustained attention to implementation detail.

The Innovation and competitiveness-oriented coalition articulates innovation primarily through enabling conditions rather than abstract growth claims. Emphasis falls on competitiveness, ecosystem development, and capacity-building, often expressed through governance instruments such as sandboxes, testbeds, and routinised compliance pathways. The General or mixed governance-oriented coalition exhibits the weakest thematic consolidation. Its belief profile combines elements of rights protection, safety concerns, and governance feasibility without a dominant organising axis, reflecting greater internal dispersion.

Belief-level composition also differs systematically across coalitions. The Rights-oriented coalition contains the highest proportion of deep core statements. By contrast, the Technical governance-oriented, Innovation and competitiveness-oriented, Medical and sectoral safety-oriented, and Industry-oriented coalitions are dominated by policy core statements centred on governance design. The Medical coalition additionally shows a comparatively high share of secondary aspects, reflecting sustained engagement with implementation choices.

These profiles indicate that consultation-stage conflict in EU AI governance is structured not by a simple pro- versus anti-regulation divide, but by how actors locate the central regulatory problem. Some prioritise rights and democratic control, others compliance feasibility, governance capacity, sectoral safety, or innovation enablement through experimentation and standards.

Figure 4. Top 8 belief indicators per coalition

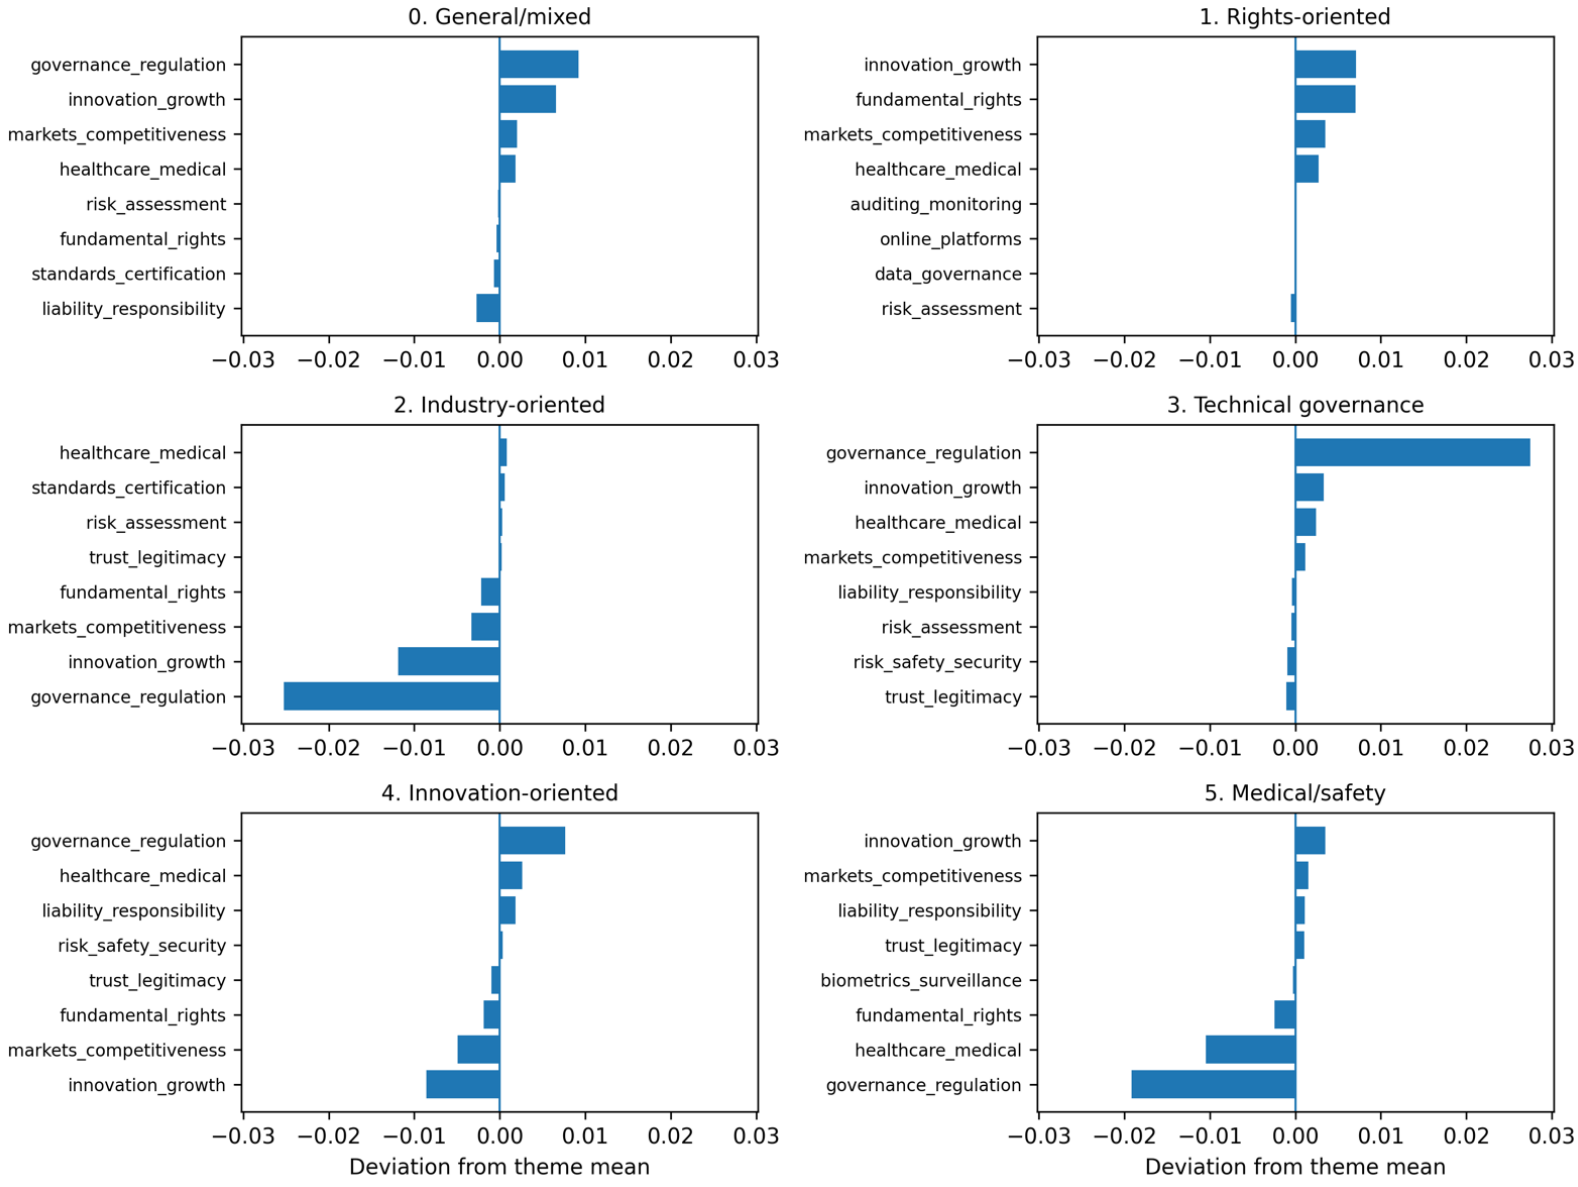

4.4 Summary

Two findings stand out from the coalition analysis. First, consultation-stage conflict in EU AI governance is not structured around a binary industry–civil society divide. Instead, it clusters into a limited number of belief-based coalitions that cut across organisational categories. Coalition membership corresponds to shared interpretations of policy problems and regulatory priorities rather than to sectoral identity alone, in line with a core expectation of the Advocacy Coalition Framework in a technically complex regulatory domain.

Second, coalitions differ not only in what they prioritise, but in where they locate the core regulatory problem. Coalitions anchored in deep core beliefs articulate expansive normative claims centred on fundamental rights, democratic accountability, or societal harm. Coalitions dominated by policy core and secondary aspects emphasise governance design, feasibility, standards, and implementation capacity. As policy moves from consultation toward legalisation, claims articulated at the level of governance architecture and operational detail tend to correspond more closely with binding provisions than claims framed primarily as broad normative or strategic demands.

5. Policy Movement Analysis

Policy emphasis shifts as EU AI governance moves from agenda-setting to binding law. The 2020 White Paper foregrounds innovation, competitiveness, and broad rights-based framing. By contrast, the Artificial Intelligence Act adopted in 2024 places far greater weight on governance architecture, risk classification, procedural obligations, and enforcement mechanisms. This chapter examines how that shift corresponds with the belief structures articulated by stakeholders during the consultation phase.

Using the belief-theme framework developed in Chapters 2–4, the analysis compares the thematic composition of the White Paper and the final Act and assesses whether policy movement corresponds more closely with the belief profiles of advocacy coalitions identified in Chapter 4. The comparison does not treat thematic similarity as evidence of lobbying success or actor-level influence. Instead, it documents correspondence between consultation-stage belief configurations and legislative content as policy hardens into law.

The documents perform different institutional functions. A White Paper articulates problems, priorities, and possible directions; a regulation must allocate responsibility, specify procedures, and establish enforceable oversight. Changes in thematic emphasis therefore reflect not only political prioritisation but also the structural demands of legalisation. Observed alignment is interpreted as an outcome of institutional selectivity rather than as proof of causal influence.

The correspondence analysis identifies which belief configurations articulated in mid-2020 align more closely with the final Act adopted in 2024, using a consistent belief extraction framework. It cannot, however, explain why particular beliefs came to dominate the final text. The analysis does not observe the policymaking process itself (Commission deliberations, inter-institutional negotiations, or parliamentary amendment) and therefore cannot attribute outcomes to specific actors or decisions.

5.1 Aggregate Policy Movement Across Belief Themes

Policy movement is operationalised as change in the relative emphasis of belief themes between the 2020 White Paper and the final AI Act. Because the documents differ sharply in genre and institutional purpose, movement is assessed using proportional theme distributions rather than absolute counts. The focus is on shifts in relative priority rather than differences in document length or drafting style. This approach captures how policy emphasis changes as

agenda-setting gives way to legalisation, when regulatory content must be rendered into enforceable obligations, procedures, and institutional responsibilities.

Figure 5 shows that thematic emphasis rebalances systematically as EU AI governance moves from agenda-setting to binding regulation. Governance-oriented themes expand most clearly in relative prominence. *Governance_regulation* exhibits by far the largest positive change, indicating a strong consolidation of attention around regulatory architecture, administrative responsibility, and institutional oversight. Governance themes such as biometrics and surveillance controls, security and migration concerns, and prohibitions also increase in relative emphasis. Once regulation enters binding legal form, specifying enforceable obligations and allocating responsibility becomes unavoidable, and this shift is reflected directly in the structure of the final Act.

By contrast, several prominent consultation themes contract as policy hardens into law. Strategic and market-oriented framings, most notably *innovation_growth*, decline in relative emphasis in the AI Act. Themes related to liability and responsibility and risk and safety framing also show substantial negative movement. This pattern does not indicate the disappearance of innovation, safety, or responsibility concerns. Rather, it reflects their displacement from overarching framing principles toward more circumscribed and situational treatment as governance and enforcement structures come to dominate the regulatory text. Innovation persists primarily through specific enabling instruments rather than shaping the overall logic of the regulation.

Rights-related themes display a more selective pattern of movement. Certain elements increase in prominence where they can be stabilised in legally explicit form, particularly through prohibitions and oversight requirements. Broader normative and ethical framing such as societal impact narratives and trust-building discourse declines relative to increasingly procedural and institutional content. This selective rebalancing indicates that normative concerns are narrowed and reformulated in ways compatible with legal specification and administrative implementation.

The results visible in Figure 5 indicate that policy movement across the AI Act is structured rather than ad hoc. As policy progresses toward binding regulation, beliefs that can be expressed as enforceable governance obligations expand in relative prominence. Beliefs articulated primarily as broad normative aspirations or strategic visions contract. At the document level, this rebalancing is consistent with the expectation that legalisation conditions which belief types can persist as policy hardens into law, setting the stage for examining how these shifts differentiate across advocacy coalitions in the next section.

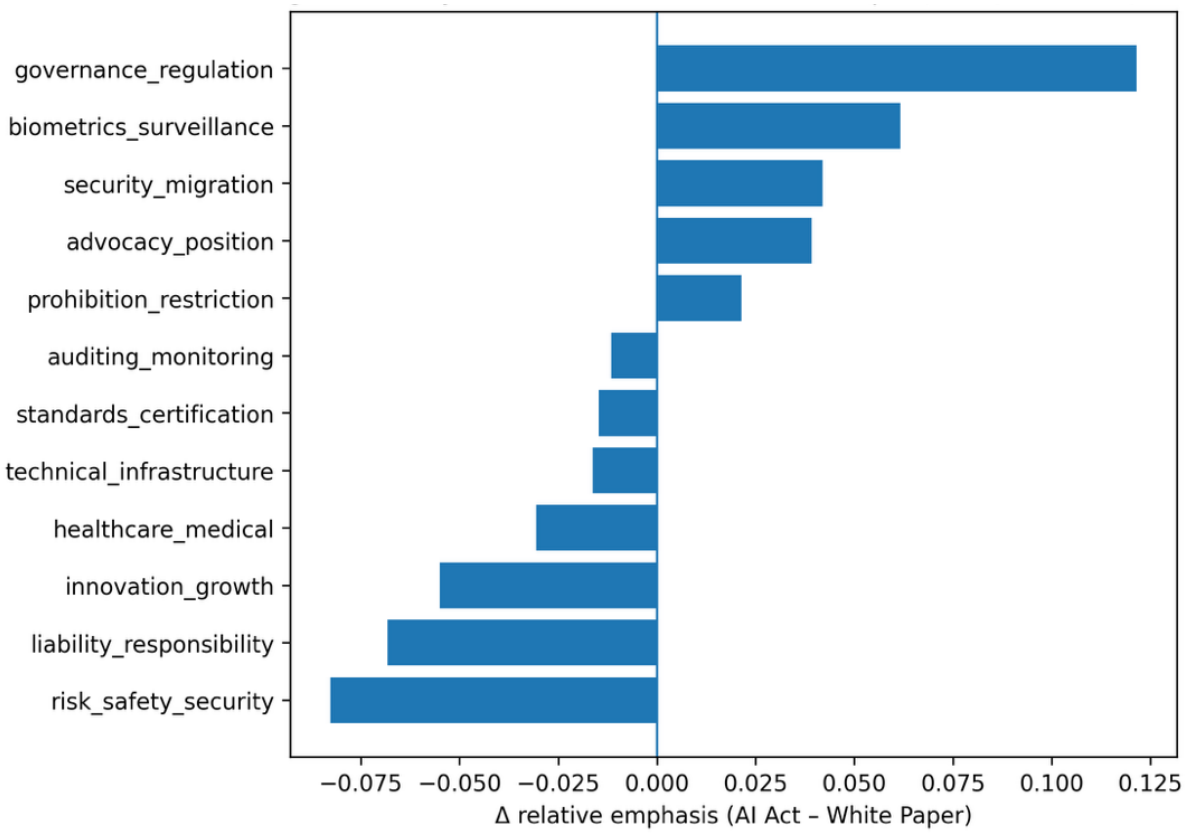

Figure 5. Policy theme movement from White Paper (2020) to AI Act (2024).

5.2 Coalition–Policy Alignment

To assess how consultation-stage belief coalitions translate into binding regulation, this section examines changes in coalition–policy alignment as EU AI governance moves from the 2020 White Paper to the final AI Act. Alignment is measured as cosine similarity between coalition-level belief profiles and policy texts. The analysis focuses on change in alignment under legalisation rather than absolute similarity, capturing which belief configurations persist as policy hardens into binding law.

Figure 6 shows clear differentiation in how coalitions' alignment shifts over the policy cycle. Governance-oriented coalitions experience the smallest reduction in alignment, indicating that their belief profiles remain largely compatible with the final regulatory form. The Technical governance coalition exhibits the least decline, followed closely by the General/mixed coalition. These coalitions articulate beliefs primarily in terms of regulatory architecture, risk categorisation, administrative oversight, and enforceable obligations. Elements that map directly onto the requirements of binding EU legislation.

Rights-oriented coalitions occupy an intermediate position. Their alignment declines more than that of governance-focused coalitions but substantially less than that of sectoral or strategic coalitions. This reflects selective incorporation: rights-related beliefs persist where they can be stabilised as explicit prohibitions, oversight requirements, or procedural safeguards. Broader ethical and societal framing loses compatibility as policy enters legal form.

By contrast, innovation-oriented and industry-oriented coalitions experience the largest reductions in alignment. These coalitions place greater emphasis on strategic flexibility, market positioning, and sector-specific concerns. These are belief preferences that are less readily translatable into uniform regulatory obligations. The magnitude of alignment loss for industry-oriented coalitions is particularly notable, indicating that resource-rich actors do not retain belief alignment where their claims conflict with legal standardisation.

The selectiveness mechanism seems systematic rather than incidental. Legalisation does not uniformly privilege certain actors or coalitions; instead, it filters belief types differentially. Beliefs articulated in governance-operable terms retain alignment as policy hardens into law. Beliefs framed primarily as strategic objectives or broad normative aspirations lose compatibility. This structured differentiation across coalitions supports institutional filtering as the mechanism shaping consultation–legislation translation in EU AI governance.

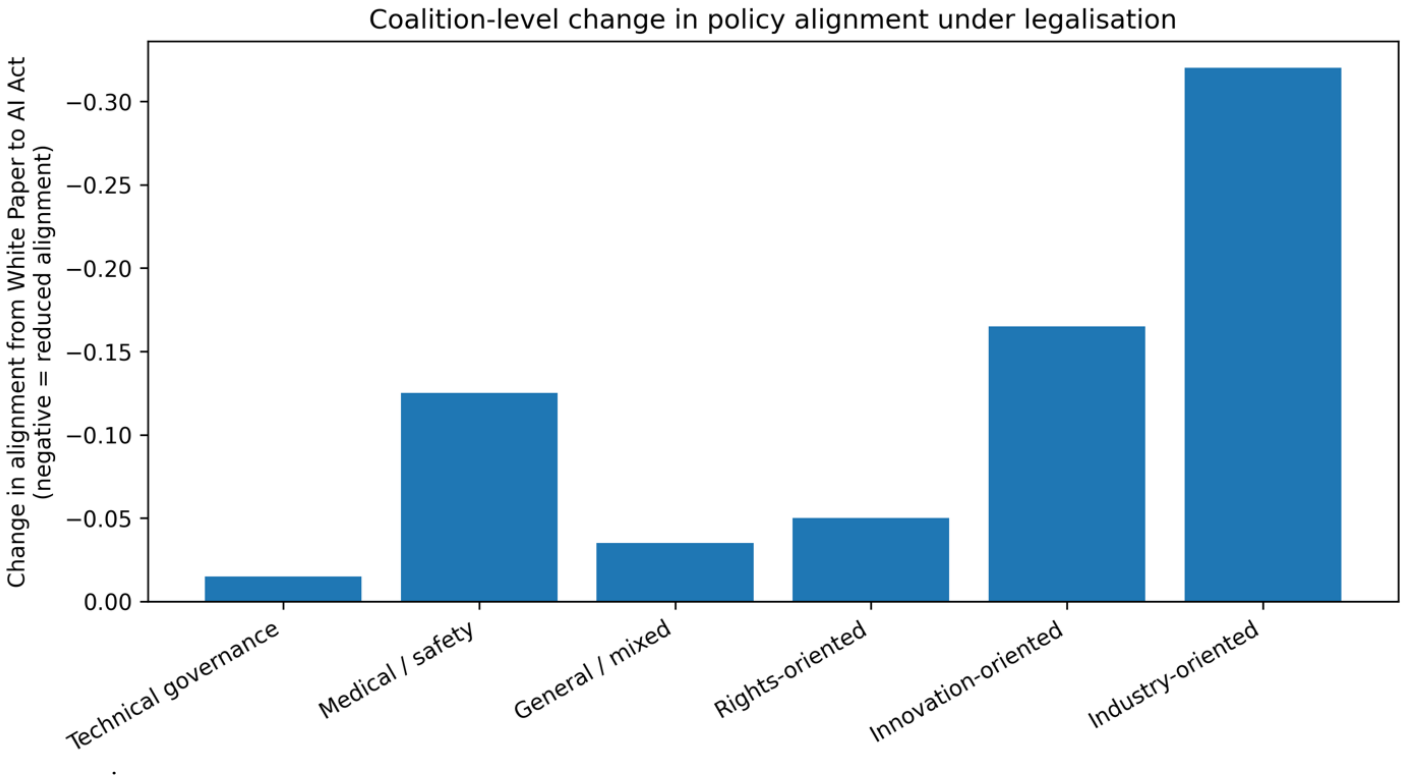

Figure 6. Bars show the change in alignment (AI Act − White Paper); low (negative) values indicate more alignment with the final AI Act.

5.3 Interpreting Policy Movement

Sections 5.1 and 5.2 found evidence indicating that belief persistence through EU AI governance is structured by legal form rather than by consultation prominence or actor type. What survives legalisation are not the most widely expressed preferences, but belief emphases that can be rendered as enforceable obligations. This characteristic holds across all coalition types and across both document-level and coalition-level analyses.

Policy movement shows systematic rebalancing toward governance architecture, oversight, and enforcement, while strategic and normative framing contracts. This shift reflects the institutional requirements of binding EU law, which must specify obligations and establish supervisory mechanisms. As policy enters legal form, claims that cannot be expressed in these terms lose relative compatibility, regardless of their salience during consultation.

Coalition-level analysis reinforces this interpretation. Figure 6 shows that governance-oriented coalitions experience the smallest reductions in alignment under legalisation, while innovation- and industry-oriented coalitions experience the largest. Rights-oriented coalitions occupy an intermediate position, retaining alignment where beliefs can be stabilised as prohibitions, oversight requirements, or procedural safeguards, but losing compatibility where claims rely on expansive ethical or societal framing. This structured differentiation by belief type—rather than by organisational identity or resource endowment—indicates that legalisation filters beliefs according to their expressibility as binding rules.

This points toward institutional filtering as the operative mechanism structuring consultation–legislation translation in EU AI governance. Better Regulation scholarship predicts that EU legal drafting privileges claims that can be articulated through causal reasoning, proportionality assessment, and administratively workable procedures (Radaelli & Meuwese, 2009). The evidence presented here confirms that expectation across multiple dimensions: beliefs formulated in governance-operable terms retain compatibility as policy hardens into law, while beliefs framed primarily as strategic objectives or broad normative aspirations lose alignment.

Could alternative mechanisms account for these findings? Commission discretion in selecting consultation inputs might explain selective uptake, but it would not predict systematic differentiation by belief type rather than actor category. Parliamentary amendment dynamics could strengthen or weaken individual provisions, but the patterns documented here emerge already in the Commission's proposal and concern relative thematic emphasis rather than isolated articles. Council negotiations may narrow obligations, but they do not account for the consistent contraction of strategic and expansive normative framing observed across coalitions. While process-tracing would be required to establish causation definitively, institutional filtering offers the most parsimonious explanation for the structured patterns observed.

5.4 Summary

This chapter has examined how consultation-stage beliefs translate into binding EU AI regulation as policy moves from agenda-setting to legalisation. Across both document-level and coalition-level analyses, three consistent patterns emerge.

Policy movement from the 2020 White Paper to the final AI Act is uneven and directional. As policy hardens into binding regulation, governance design, institutional

oversight, enforceability, and risk-management architecture expand in relative prominence. These elements come to dominate the regulatory text not because they were uncontested during consultation, but because they become institutionally unavoidable once legal obligations must be specified and responsibility must be allocated.

Innovation-oriented framing recedes over this transition. It remains visible through specific instruments such as regulatory sandboxes and targeted support measures, but it no longer structures the regulation as a whole. Strategic and proportionality-based claims follow a similar trajectory. They are acknowledged in principle, yet increasingly constrained by the need to define compliance pathways, supervisory roles, and enforcement arrangements as the regulation takes legal form.

Rights-based concerns occupy an intermediate position. Certain elements—most notably prohibitions on specific AI practices and requirements for human oversight—gain prominence as they are translated into enforceable provisions. Broader ethical narratives and societal framing, by contrast, lose visibility as the text becomes more procedural and legalistic. This reflects not the displacement of rights-based claims, but a narrowing of the forms in which they can be sustained within binding EU law.

In sum, the patterns documented in this chapter indicate institutional selectivity rather than preference satisfaction. Beliefs articulated during consultation persist only insofar as they can be rendered compatible with legal certainty, administrative capacity, and enforceable governance. Legalisation and consultation do not eliminate conflict, but condition which belief types remain compatible with the final regulatory outcome. The next chapter turns to what this implies for participation, legitimacy, and influence in EU AI policymaking

6. Conclusion and Discussion

6.1 Summary and answer to the research question

This thesis set out to explain why an extensive and formally inclusive consultation process in EU AI governance nonetheless produced a regulatory outcome that many participants did not recognise as reflecting their positions. Focusing on the European Commission's 2020 consultation on the White Paper on Artificial Intelligence and the subsequent development of the Artificial Intelligence Act, it tested whether institutional filtering helps explain how stakeholder beliefs articulated during agenda-setting translate into binding regulatory outcomes once legalisation requirements take effect.

Consultation is treated as a venue in which actors articulate priorities in relatively open language, followed by a legislative phase in which those priorities must be rendered compatible with a regulation that allocates responsibilities, defines procedures, and remains enforceable across Member States. A consistent pattern emerges. Conflict expressed during consultation does not resolve into a simple industry–civil society divide. Instead, positions cluster into belief configurations that cut across organisational categories. Belief-based clustering of 326 position papers identifies six coalitions organised around competing governance priorities.

Tracing these belief configurations from consultation through to the final AI Act reveals a clear directional shift in policy emphasis. Governance architecture, oversight mechanisms, and enforceability expand in relative prominence as policy hardens into law, while broad innovation framing and open-ended normative language contract. Rights-oriented claims do not disappear but change form. Selected concerns stabilise through prohibitions, high-risk obligations, and human oversight requirements, however more expansive ethical and societal narratives recede as the text becomes increasingly procedural and legalistic.

Belief configurations framed in governance-operable terms show higher correspondence with the final Act than normative or strategic claims, reflecting selective translation under legalisation: as policy moves into binding form, beliefs already expressed in legally and administratively workable terms persist more fully in regulatory text, even though beliefs that resist specification lose relative compatibility.

The analysis documents outcomes rather than decision-making processes. Alternative institutional mechanisms could generate similar results, but the evidence shows that belief persistence through the AI Act follows a structured pattern aligned with legal expressibility

rather than consultation prominence or organisational power. On this basis, institutional filtering provides the most compelling interpretation of how consultation-stage beliefs translate into binding EU AI regulation.

6.2 Reflection and limitations

The primary limitation concerns causal inference. The study documents how belief configurations articulated during consultation align differently with the final AI Act once legalisation requirements take effect, but it does not reconstruct the decision-making processes through which these patterns emerged. Observed correspondence between governance-oriented beliefs and legislative outcomes could reflect selective uptake by the Commission, institutional constraints that privilege particular regulatory forms irrespective of advocacy, or a combination of the two. The prominence of procedural and administrative elements, such as governance roles, conformity assessment mechanisms, supervisory capacity, and enforcement architecture, points toward the influence of legal and institutional necessity rather than preference aggregation alone, while remaining compatible with alternative explanations involving political compromise or strategic drafting.

The methodological approach involves trade-offs. The deterministic, rule-based belief extraction framework prioritises transparency and reproducibility, but may undercount implicit reasoning and legally sophisticated signalling. Some actors—particularly law firms and business associations—express positions through precise legal formulation rather than explicit normative claims, which sentence-level rule matching does not fully capture. Coalition identification through k-means clustering likewise imposes analytical structure. Coalition boundaries should be understood as analytical constructs rather than fixed categories. The similarity of results across alternative cluster solutions indicates a structured belief landscape with internal heterogeneity and overlap.

Applying the Advocacy Coalition Framework to a compressed regulatory cycle further constrains interpretation. Although the framework is well suited to capturing belief-based conflict that cuts across organisational categories, it offers limited leverage for explaining why particular governance architectures achieved legal viability beyond the general observation that binding regulation privileges specificity and administrative workability. In this study, the framework therefore functions as an interpretive tool for structuring belief analysis rather than as a causal model of policy change.

These limitations constrain causal claims but do not undermine the central empirical finding: belief persistence through the AI Act follows a systematic pattern aligned with legal expressibility rather than consultation prominence or organisational power

6.3 Policy recommendations

The findings of this study indicate that consultation–legislation gaps in EU AI governance reflect institutional filtering rather than participation barriers or lobbying capture. Extensive consultation produced a wide range of articulated beliefs, yet only a subset translated into the final AI Act. The evidence shows that belief persistence is shaped less by how widely positions are expressed than by whether they can be rendered compatible with binding legal form. This has concrete implications for how consultation processes are designed, interpreted, and evaluated within EU policymaking.

First, expectation gaps in consultation should be addressed more explicitly. Consultation documents frequently invite broad reflections on values, innovation, and societal impact. However, the legislative phase necessarily prioritises enforceable obligations, administrative roles, and procedural clarity. The analysis shows that beliefs articulated at a high level of abstraction are systematically less likely to persist through legalisation. Making the institutional constraints of legal drafting more visible at the consultation stage would reduce mismatches between participant expectations and legislative outcomes, without narrowing access or constraining participation.

The treatment of innovation-oriented claims reflects a structural tension between legalisation and enablement. While innovation is central to AI governance, it is typically framed in terms of flexibility and experimentation rather than enforceable obligations. As a result, innovation concerns persist in the AI Act mainly through bounded instruments such as sandboxes or exemptions, rather than through its core regulatory architecture. This reflects the limits of legal form, not the marginalisation of innovation.

Second, consultation design should better support belief translation into legally and administratively workable terms. Actors differ substantially in their capacity to express concerns in governance-operable language. The findings show that beliefs articulated in such terms retain alignment as policy hardens into law, others lose compatibility regardless of consultation prominence. Providing guidance, structured prompts, or illustrative examples that encourage participants to specify implementation pathways, responsibility allocation, or

enforcement implications could reduce asymmetries in belief persistence without privileging particular actors or interests.

Third, evaluation criteria for consultation processes should be recalibrated. Consultation effectiveness is often assessed in terms of participation breadth or diversity of viewpoints expressed. The evidence presented here suggests that these metrics are insufficient on their own. Even highly inclusive consultations can yield predictable patterns of legislative contraction if institutional filtering remains unaddressed. Evaluating consultations in terms of how well they facilitate translation into legal form would provide a more accurate account of their democratic function within EU governance.

Overall, these implications follow directly from the systematic patterns documented in this study. The AI Act consultation generated extensive participation and diverse belief configurations, yet legislative content aligned most closely with governance-oriented claims. Unless consultation design accounts for how beliefs translate into enforceable regulation—not merely who participates—similar gaps between input diversity and legislative output are likely to recur across EU digital governance initiatives.

6.4 Future research

Several extensions could strengthen or challenge the argument developed here. A first is process oriented. Interviews with Commission officials, parliamentary rapporteur teams, and selected stakeholder representatives could help assess whether observed correspondence reflects deliberate uptake, strategic compromise, or convergence driven primarily by legal requirements. Even a limited set of targeted interviews could clarify whether belief alignment results from institutional filtering or from actor-level influence exercised during drafting and negotiation.

A second extension is comparative. Applying the same belief-extraction framework to other EU digital regulatory files—such as the Digital Services Act, Digital Markets Act, data governance initiatives, or cybersecurity legislation—would help determine whether the pattern observed here is specific to AI's risk-based architecture or more general to EU legalisation in technically complex domains.

A final methodological extension would combine belief similarity with relational data. Incorporating information on co-signatures, joint position papers, or consultation networks could help distinguish between shared emphasis articulated independently and coordinated advocacy efforts. Such an approach would complement text-based correspondence analysis by

providing additional leverage for identifying when belief alignment reflects structural compatibility rather than strategic coalition behaviour.

Bibliography

Bunea, A. (2013). Issues, preferences and ties: Determinants of interest groups' preference attainment in the EU environmental policy. *Journal of European Public Policy, 20*(4), 552–570. https://doi.org/10.1080/13501763.2012.726467

Coen, D., & Richardson, J. J. (2009). *Lobbying the European Union: Institutions, actors, and issues* (1st ed.). Oxford University Press. https://doi.org/10.1093/oso/9780199207350.001.0001

de Boer, N., & Raaphorst, N. (2023). Automation and discretion: Explaining the effect of automation on how street-level bureaucrats enforce. *Public Management Review, 25*(1), 42–62. https://doi.org/10.1080/14719037.2021.1937684

Dehousse, R. (2008). Delegation of powers in the European union: The need for a multi-principals model. West European Politics, 31(4), 789–805. Https://doi.org/10.1080/01402380801906072

Dür, A. (2008). Interest groups in the European Union: How powerful are they? *West European Politics, 31*(6), 1212–1230. https://doi.org/10.1080/01402380802372662

European Commission. (2020). Final report on the public consultation on the White Paper on Artificial Intelligence. Publications Office of the European Union.

European Commission. (2020). White paper on artificial intelligence: A European approach to excellence and trust (COM(2020) 65 final). Publications Office of the European Union.

European Commission. (2023). Better regulation guidelines (European Commission Staff Working Document SWD(2021) 305 final, revised 3 November 2021; Better Regulation: Guidelines and Toolbox). Publications Office of the European Union.

European Parliament & Council of the European Union. (2024). Regulation (EU) 2024/1689 of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) No 2016/797 and (EU) 2020/1828. Official Journal of the European Union, L 1689.

Gorwa, R., Binns, R., & Katzenbach, C. (2020). Algorithmic content moderation: Technical and political challenges in the automation of platform governance. *Big Data & Society, 7*(1), 205395171989794. https://doi.org/10.1177/2053951719897945

Gorwa, R., Lechowski, G., & Schneiß, D. (2024). Platform lobbying: Policy influence strategies and the EU's Digital Services Act. *Internet Policy Review, 13*(2), 1–26. https://doi.org/10.14763/2024.2.1782

Helberger, N., & Diakopoulos, N. (2023). ChatGPT and the AI Act. *Internet Policy Review, 12*(1). https://doi.org/10.14763/2023.1.1682

Hix, S., & Høyland, B. K. (2022). *The political system of the European Union* (4th ed.). Bloomsbury Academic.

Jenkins-Smith, H. C., Nohrstedt, D., Weible, C. M., & Ingold, K. (2018). The advocacy coalition framework: An overview of the research program. In C. M. Weible & P. A. Sabatier (Eds.), *Theories of the policy process* (pp. 135–171). Routledge. https://doi.org/10.4324/9780429494284-5

Krarup, T., & Horst, M. (2023). European artificial intelligence policy as digital single market making. *Big Data & Society, 10*(1). https://doi.org/10.1177/20539517231153811

Leifeld, P. (2013). Reconceptualizing major policy change in the advocacy coalition framework: A discourse network analysis of German pension politics. *Policy Studies Journal, 41*(1), 169–198.

Page, E. C., Balla, S. J., & Lodge, M. (2015). Paul A. Sabatier, "An advocacy coalition framework of policy change and the role of policy-oriented learning therein." In *The Oxford handbook of classics in public policy and administration*. Oxford University Press. https://doi.org/10.1093/oxfordhb/9780199646135.013.24

Radaelli, C. M. (2005). Diffusion without convergence: how political context shapes the adoption of regulatory impact assessment. Journal of European Public Policy, 12(5), 924–943. Https://doi.org/10.1080/13501760500161621

Radaelli, C. M., Dunlop, C. A., & Fritsch, O. (2013). Narrating impact assessment in the European Union. *European Political Science, 12*(4), 500–521.

Radaelli, c. M., & meuwese, a. C. M. (2009). Better regulation in europe: between public management and regulatory reform. Public Administration (London), 87(3), 639–654. Https://doi.org/10.1111/j.1467-9299.2009.01771.x

Sabatier, P. A., & Jenkins-Smith, H. C. (1993). *Policy change and learning: An advocacy coalition approach*. Westview Press.

Sabatier, P. A., & Weible, C. M. (2007). The advocacy coalition framework: Innovations and clarifications. In P. A. Sabatier (Ed.), *Theories of the policy process* (pp. 189–220). Routledge.

Schmidt, V. A. (2008). Discursive institutionalism: The explanatory power of ideas and discourse. *Annual Review of Political Science, 11*, 303–326.

Schmidt, V. A. (2010). Taking ideas and discourse seriously: Explaining change through discursive institutionalism as the fourth 'new institutionalism'. *European Political Science Review, 2*(1), 1–25.

Tallberg, J., Lundgren, M., & Geith, J. (2024). AI regulation in the European Union: Examining non-state actor preferences. *Business and Politics, 26*(2), 218–239.

Appendices

Appendix A.1 Deterministic Rule

This appendix provides an example overview of the deterministic coding rules used for belief extraction. For each belief theme, it lists representative trigger patterns (simplified for readability) and the associated rule-weight range.

ACF level	Belief theme	Representative trigger patterns (simplified)	Weight range
deep_core	democratic_values	democracy/ies; freedom; autonomy	2–3
deep_core	fundamental_rights	fundamental rights; dignity; human rights	2–4
deep_core	innovation_growth	prosperity; growth	2–3
deep_core	markets_competitiveness	competition; markets	2–3
deep_core	social_justice	solidarity; common good; non[-]discriminat	2–3
deep_core	trust_legitimacy	legitimacy; trust	2–3
policy_core	governance_regulation	governance; regulation; regulatory	2–3
policy_core	innovation_growth	smes; innovation; investment	2
policy_core	markets_competitiveness	competitive/iveness	2
policy_core	prohibition_restriction	moratorium; banned/ning; prohibited/ion	3

policy_core	risk_assessment	impact assessment; high[-]risk; risk[-]based	2–3
policy_core	standards_certification	certification; standards/isation?; conformity	2–3
secondary	auditing_monitoring	audits/ing	2
secondary	biometrics_surveillance	facial recognition; biometric	2–3
secondary	creative_industries	music; media; cultural	1–2
secondary	data_governance	data quality; training data; datasets	2–3
secondary	experimentation_sandboxes	Sandboxes	2

Appendix A.2 Corpus Construction and Preprocessing

Table A.2.1 Corpus construction and preprocessing pipeline

Stage	**Number of documents**
Documents successfully extracted from PDFs	440
Documents retained after deterministic cleaning	326
Documents excluded during cleaning	114

Table A.2.2 Excluded documents and reasons for exclusion

Filename (shortened)	**Exclusion reason**
F529882-DK_Specific_Comments_on_the_AI-White_Paper…	Too short

F529937-FR_CCBE-Response-to-the-consultation…	Non-English (French)
F530039-OEGB_Kuenstliche_Intelligenz…	Non-English (German)
F530117-RN_Report_Facial_Recognition_FR…	Non-English (French)
F530197-PAI_response_to_EC_White_Paper…	Non-English (German)
(109 further documents)	Non-English / insufficient length

All exclusions were based on predefined, content-neutral criteria (language and minimum length). No document was excluded on the basis of substantive position or actor type.

Appendix A.3: Examples of coded belief-bearing sentences

A "coded sentence" is a sentence from the cleaned corpus that matches at least one regex pattern in the deterministic rule library. For transparency, this table lists illustrative coded sentences alongside their assigned belief theme and ACF level, and the specific rule trigger (regex literal) used for coding.

Table A3.1 Examples of coded belief-bearing sentences

Filename (shortened)	Coded extract	Belief theme	ACF level	Rule trigger
F530380-Outline_reactie_internetconsultatie_AI..	However, the use of artificial intelligence can seriously affect fundamental rights.	fundamental_rights	deep_core	"fundamental rights"
F530130-European_Tech_Alliance_-_Addendum_Paper..	While AI does not hold all the answers, it will be a crucial part of Europe's immediate recovery and return to long-term prosperity.	innovation_growth	deep_core	"prosperity"
F529843-CNMC_S_AND_ACCO_S_JOINT_CONTRIBUTION…	2 For example, by the Catalan Competition Authority itself in 'The Data-Driven Economy.	markets_competitiveness	deep_core	"competition"
F530247-MDT_Submission_to_European_Commission..	MDR's detailed 'safety' concept already addresses potential cyber-vulnerability.	risk_safety_security	deep_core	"safety"
F530487-Wachter_et_al_Why_Fairness_Cannot_Be_Automated_Brid..	Such a situation would undermine how non-discrimination law is practiced in Europe.	social_justice	deep_core	"non-discriminat"

F530117-200612_RN_Report_Facial_Recognition..	This means in particular that from the moment a person is arrested, the legitimacy of the arrest must be demonstrated.	trust_legitimacy	deep_core	"legitimacy"
F529948-1106_Reply_to_Consultation_White_Paper..	This does not prevent us from endorsing a sectoral approach to governance of AI.	governance_regulation	policy_core	"governance"
F529888-DIGITAL_SME_Position_Paper_AI_White_Paper_FINA …	Healthcare solutions can help increase the life expectancy of the EU population.	healthcare_medical	policy_core	"healthcare"
F514709-CPME.Board_.Nov19.FINAL_.EN …	The limits of the physician´s liability when using AI have to be clearly defined.	liability_responsibility	policy_core	"liabilit(y|ies"
F514777-Algo.Rules_EN…	Before an algorithmic system is put to use, an impact assessment should be conducted and documented.	risk_assessment	policy_core	"impact assessment"
F529994-Feedback_on_White_Paper_on_Artificial_Intelligence__KU_Leuven_ …	These challenges can have an impact on the quality of the certification process.	standards_certification	policy_core	"certification"
F530347-V29_Legal_s_Contribution_Paper …	This implies an element of explainability, without explicitly naming it as such.	transparency_explainability	policy_core	"explainab(le|ility)"

F530412-IDEMIA_Consultation_on_the_White_Paper_on_Artificial_Intellige…	and Ensure that citizens are well-informed of the benefits of facial recognition.	biometrics_surveillance	secondary	"facial recognition"
F530308-AI_Position_Paper_SN_final2…	If a company wants to operate globally, training data must be collected globally.	data_governance	secondary	"training data"
F530088-CompTIA_EC_AI_White_Paper_FINAL_SH…	This would create a significant backlog for the newly established testing centers.	technical_infrastructure	secondary	"testing"

Appendix A.4: Clustering solution and diagnostics

This appendix documents the clustering procedure used to identify belief coalitions among actors, and reports the diagnostic criteria supporting the selection of the final clustering solution.

Clustering specification

Actors were clustered using the k-means algorithm applied to actor-level belief vectors derived from the deterministic coding procedure described in the methods section. Cosine distance was used as the similarity metric. Cluster membership remained stable across re-seeding.

Table A.4.1 — Clustering diagnostics and model choice

Criterion	Specification / result
Clustering algorithm	k-means
Distance metric	Cosine similarity
Input data	Actor-level belief vectors
k tested	3–10
Selected k	6
Selection criteria	Elbow (WCSS), silhouette score, Davies–Bouldin index
Stability check	Stable under re-seeding

Cluster selection diagnostics

Figures A.4.2–A.4.5 report standard clustering diagnostics. The elbow plot shows diminishing returns in within-cluster sum of squares beyond k = 6. While the silhouette score peaks at lower values of k, it declines gradually thereafter, indicating no single natural partition. The Davies–Bouldin index worsens for $k \geq 7$. The diagnostics support the selection of k = 6 as a substantively meaningful and defensible compromise.

Figure A.4.2 Cluster selection diagnostics: silhouette score by k

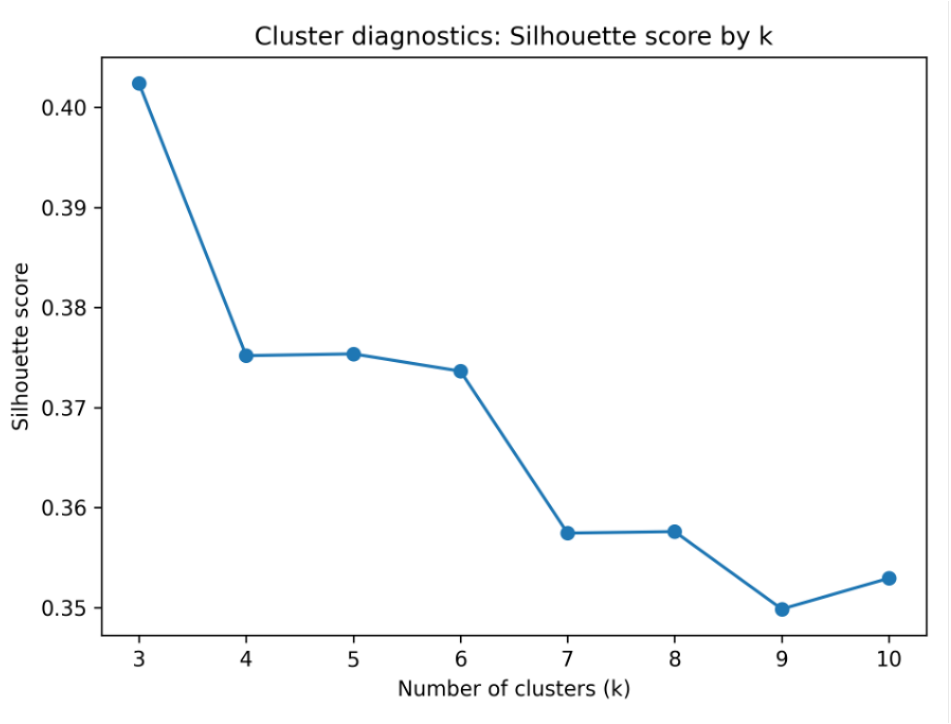

Figure A.4.3 — Cluster selection diagnostics: elbow plot (WCSS)

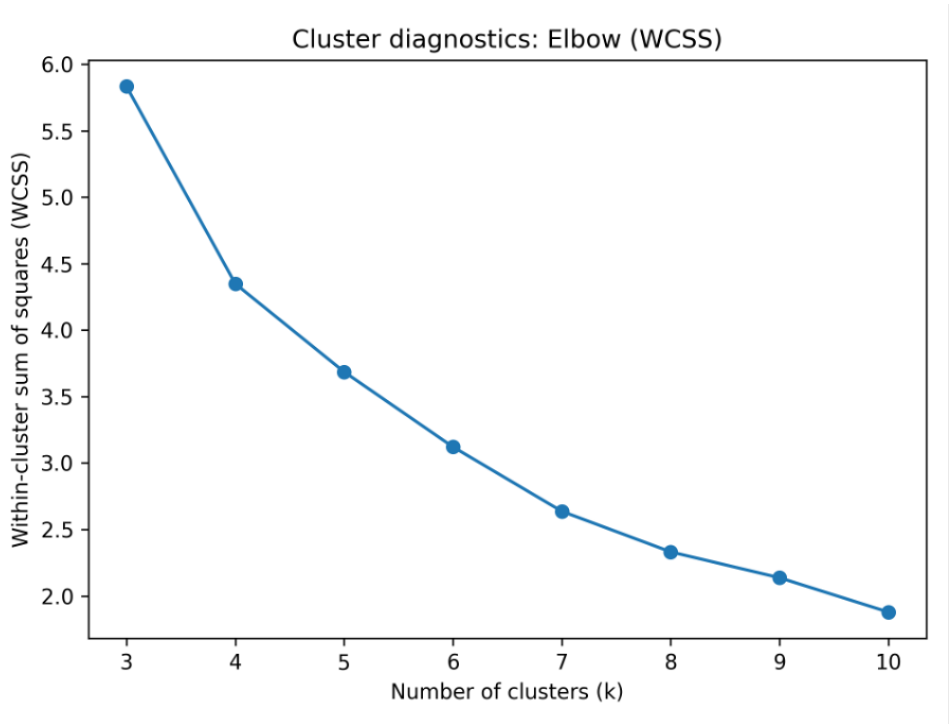

Figure A.4.5 — Cluster selection diagnostics: Davies–Bouldin index by k

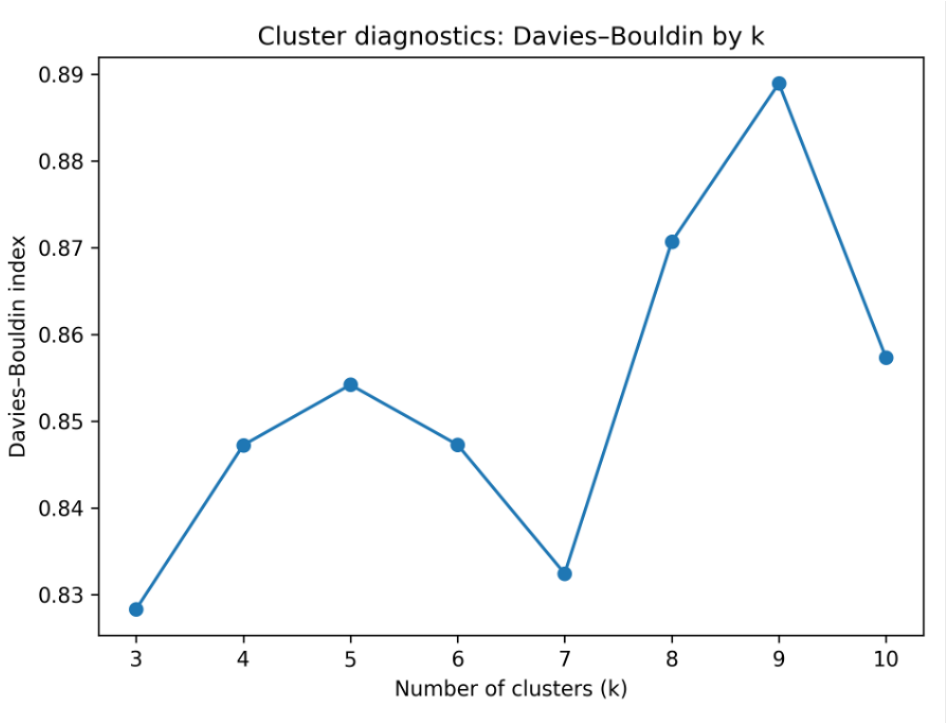

Appendix A.5: Uncoded sentence examples (deterministic coding)

An “uncoded sentence” is a sentence from the cleaned corpus that does not match any regex pattern in the deterministic rule library. This appendix provides a summary of uncoded-sentence rates across documents and examples sampled across the corpus.

Table A.5.1 — Uncoded sentence rate summary (document-level)

Statistic	Value
Documents (n)	326
Median uncoded sentence rate	0.514
Uncoded rate IQR (25–75%)	0.426 – 0.616
Min–Max uncoded rate	0.000 – 1.000
Median total sentences per document	78
Median uncoded sentences per document	40

Table A.5.2 — Examples of uncoded sentences (no rule match)

Sampling note. Examples were selected using stratified sampling across documents (one uncoded sentence per sampled document, prioritising longer sentences).

Filename (shortened)	**Uncoded sentence (verbatim)**	**Diagnostic label**
F530275-European_Womens_Lobby__AI_European_Commission_Consultation_Sup…	Digital technologies, including artificial intelligence, offers huge possibilities for disruptive power to change the system, for all European citizens, with new economic opportunities and work spaces opening up.	No rule match (uncoded)
F530306-Response_to_Consultation_on_the_AI_White_Paper…	Response to the Public Consultation on the ’s White Paper on Artificial Intelligence - A European Approach Submitted by Prof.	No rule match (uncoded)
F530253-Developers_Alliance_Submission_To_The_European_Commission_s_Co…	DevelopersAlliance.org Developers Alliance Submission To The 's Consultation On The White Paper On AI Developers Alliance welcomes the opportunity to provide our contribution to the consultation on the 's White Paper on AI.	No rule match (uncoded)

F530372-StanfordWonks_TechiesEUAIWhitePaper…	J U N E 1 5, 2020 BY: “ WONKS AND TEC HIES,” a multi-disciplinary group at Stanford University, cooperating on international technology and policy issues, led by Ms. Marietje Schaake P URPOSE: 1 All of the input, ideas, and recommendations included in this paper are the reflection of the contributing authors’ own research and analysis, and are not meant to reflect the views of any institution or organization.	No rule match (uncoded)
F530231-EFPIA_EuropeanArtificialIntelligence_briefing…	In conjunction with the European Strategy for Data, EFPIA recognises the need to create an environment that unlocks the value of the data and digital economy.	No rule match (uncoded)
F529973-AI_White_Paper_public_consultation_Telenor…	Consultation on the White Paper on Artificial Intelligence - A European Approach Introduction Artificial intelligence (AI) is a strategic technology that offers many benefits for citizens and the economy.	No rule match (uncoded)
F530173-Visa_submission_to_AI_White_paper_Consultation_12062020…	Visa Response to the White Paper on Artificial Intelligence June 2020 About Visa Visa welcomes the opportunity to respond to the ’s Artificial Intelligence White Paper.	No rule match (uncoded)
F530319-IBM_Submission_on_the_EC_AI_White_Paper…	EU WHITE PAPER ON ARTIFICIAL INTELLIGENCE SUBMISSION TO THE EUROPEAN COMMISSION BY IBM JUNE 2020 IBM is the largest technology and consulting employer in the world, with over 350,000 employees serving clients in 175 countries.	No rule match (uncoded)
F530080-Liberty_Global_position_paper_EC_White_Paper_on_AI.120620…	Liberty Global is convinced that such engagement will help further achieving a Europe fit for the digital age, by ensuring that this transformation will, indeed, work for both people and businesses.	No rule match (uncoded)
F530507-FNA_Paper_on_access_to_automotive_data_-_03.06.20…	FNA 9–11 Avenue Michelet F–93583 Saint-Ouen Cedex (+33) 1 40 11 99 04 - (+33) 6 62 79 62 12 Federation of Craft Businesses in the automotive sector and in mobility services (FNA) would like to thank the authors of the White Paper on Artificial Intelligence (AI) to inform citizens and stakeholders about the Commission's work in order to allow them to provide feedback on the intended	No rule match (uncoded)

	initiative and to participate effectively in future consultation activities.	