

IRIS Case Studies

A Research Archive

One Million Words of Interpretive Analysis Across Twenty Case Studies

R. Jazinski

London, 2026

*Companion volume to Silent Transitions: How Interpretive Grammars Process
the Same World into Different Truths*

IRIS: Interpretive Rotation on Invariant Signals

How to Read This Volume

This volume is the research archive behind *Silent Transitions*. It documents a five-round validation programme that tested the IRIS method — Interpretive Rotation on Invariant Signals — across twenty case study applications, fifteen independent AI agents, and approximately one million words of analytical output.

The volume is designed to be used in three ways:

As an independent reference. Each case study section contains a briefing, grammar convergence table, selected agent outputs, and commentary. A reader can enter at any case and understand the findings without reading the others.

As evidence for the book. *Silent Transitions* makes claims about grammar convergence, ontological pluralism, and emerging grammars. This volume provides the complete empirical evidence behind those claims. Footnotes in the book point to specific cases and agents in this volume.

As a methodological resource. Researchers applying IRIS to new domains will find worked examples at every level of complexity — from quick three-agent analyses to fifteen-agent clean-slate tests. The cross-mark synthesis section extracts general lessons about how the method performs.

Not every agent output is reproduced in full. For each case, one 'best output' is included complete, plus excerpts from other agents where they contribute distinctive findings. The complete corpus (~1,045,000 words across all agents and all rounds) is published in the companion volume: *IRIS Case Studies: Complete Archive Hub* (to March 2026).

The IRIS Method

IRIS — Interpretive Rotation on Invariant Signals — is a structured analytical method for making visible how different interpretive frameworks process the same evidence into divergent conclusions.

The method holds the empirical input constant, rotates the interpretive grammar, and observes what changes in the output. An interpretive grammar is a rule-bound processing system defined by five required parameters:

$$\mathbf{G} = (\mathbf{V}, \mathbf{S}, \mathbf{W}, \mathbf{C}, \mathbf{A})$$

V = primary analytical variable (what the grammar treats as causally fundamental)

S = signal selection rules (which evidence is admitted)

W = weighting function (how evidence is ranked)

C = causal chain model (the form of explanation constructed)

A = admissible closure set (the conclusions the grammar can reach)

Four recommended parameters extend the specification: excluded signals (E), characteristic vocabulary (K), temporal orientation (T), and temporal formation (F).

IRIS does not adjudicate between grammars. It does not determine which interpretation is correct. It makes the architecture of interpretation visible — the processing rules that produce conclusions — so that divergence can be diagnosed rather than merely experienced.

For the complete operational toolkit, see the companion document: **IRIS: Interpretive Rotation on Invariant Signals — Operational Toolkit & Researcher Guide, Version 5.1.**

A note on terminology. The IRIS method was originally developed under the working title "Prismatic Method," and the prismatic metaphor remains central to its logic. Like white light passing through a prism, shared empirical signals pass through different interpretive grammars and refract into distinct outputs. The formal name — **IRIS: Interpretive**

Rotation on Invariant Signals — describes the operation precisely. Throughout this volume and in agent outputs (particularly from earlier rounds), references to "the prismatic method," "prismatic analysis," or "prismatic processing" refer to the same method now formally designated IRIS. The metaphor and the name describe the same procedure: hold the signal constant, rotate the grammar, observe the divergence.

Agents as Survey Instruments of the State Space

The IRIS method was developed and first demonstrated through manual analytical work — a single analyst applying three grammars to a single extended case. The case studies in this volume were not designed as AI research. They were designed as method validation.

The question was: are the grammars identified in Silent Transitions discoverable features of the interpretive landscape, or are they products of the analyst's own position? To answer this, independent processors were needed — analysts who could apply the method to the same domains without coordination, whose processing architectures could be documented, and who could be deployed at scale.

AI agents serve this function. Each agent has been trained on a vast corpus of human interpretive production — academic literature, policy documents, journalism, public discourse. When an agent identifies three grammars for a given domain, it is not inventing them — it is detecting patterns that are structurally recurrent in the human discourse it was trained on. The convergence across agents validates that these patterns are robust features of the discourse, not artefacts of any single processor.

The function of the agents in this programme is analogous to that of multiple coders in a content analysis: if independent processors converge on the same grammar specifications without coordination, the structures they identify are more likely to be features of the domain than features of any individual processor.

The direction of discovery runs from human to machine validation. The grammars in Silent Transitions were identified through years of manual analytical work. The agents confirmed what a human analyst had already found. They do not replace human analysis; they provide inter-coder reliability at scale.

In the course of this validation, a secondary finding emerged: the agents' own processing architectures — their home grammars, content boundaries, and alignment constraints — are themselves instances of the interpretive architectures the method is designed to reveal. This reflexive

finding is discussed in the cross-mark synthesis, but it should be understood as a byproduct of the validation exercise, not its purpose.

Programme Design: Mark 1 to Mark 5

The validation programme ran in five rounds, each expanding in scope and refining the method:

Round	Cases	Agents per Round	Words	Key Test
Mark 1	7	5	~200,000	Initial validation: do independent agents converge?
Mark 2	5	7	~258,000	Expanded agents; new domains; toolkit v2.0 tested
Mark 3	3	9	~165,000	Self-referential case (AI consciousness); refusal replication; emerging grammars
Mark 4	2	15	~175,000	Clean-slate test (no priming); maximum agent diversity; new domain (crypto)
Mark 5	3	12	~247,000	Final verification: book's grammars; Brexit; emerging grammar tracking
Total	20	15 unique	~1,045,000	

Word counts are approximate, derived from automated counting of agent output text. The progression was deliberate: broad domain coverage (Mark 1) → expanded agents and new domains (Mark 2) → maximum analytical difficulty (Mark 3) → clean-slate convergence test at maximum agent scale (Mark 4) → final verification of the book's founding cases (Mark 5).

The method itself evolved through the programme. The toolkit progressed from v1.0 (basic specification) through v2.0 (counter-signal integration, ontological check, grammar boundaries, stress test) to v3.0 (temporal formation parameter, grammar translation, emerging grammar detection). Each refinement was driven by findings from the previous round.

Agent Profiles and Home Grammars

Fifteen agents were deployed across the programme. Twelve were retained for the final round (Mark 5) based on analytical capacity; three were removed (marked below). Each agent has a documented home grammar — a consistent processing bias that shapes which grammars it identifies and how it operates them.

Agent	Developer	Rounds	Home Grammar	Notes
ChatGPT Deep Research	OpenAI	Marks 1–5	Institutional-mode rate	Consistent anchor; strong on complex domains; calibrates per domain
Claude Research	Anthropic	Marks 2–5	Analytical-reflexive	Strongest reflexive capacity; most precise ontological checks
DeepSeek	DeepSeek (China)	Marks 1–5	Critical-academic	Refusal reversal documented (refused Israel/Palestine M1; processed M3)
Gemini Pro	Google	Marks 1–5	Creative-synthetic	Generates spontaneous adjacent case studies; creative grammar naming
Grok Expert	xAI	Marks 2–5	Libertarian-adjace nt	Only non-institutional home grammar; discovers grammars others miss
SciSpace	SciSpace	Marks 1–5	Institutional-techn ocratic	Longest outputs (14K–20K words); most structured; reliable baseline
Perplexity	Perplexity	Marks 1–5	Institutional-empir ical	Most data-grounded; produced method comparison materials and taxonomy
Mistral Research	Mistral (France)	Marks 3–5	European-discours e	Distinctive category systems;

				discourse-analysis framing
Qwen3.5Plu s	Alibaba (China)	Marks 4–5	Institutional-quant itative	Most transparent about alignment constraints; second Chinese model
MetaAI Thinking	Meta	Marks 4–5	Narrative-institutio nal	Creative metaphorical naming; strong on emerging domains
Venice	Venice	Marks 4–5	Privacy-institution al	Strong autonomy sensitivity; good reflexive capacity
Okara	Okara	Marks 4–5	Cultural-hermeneu tic	Strongest ethnographic lens; distinctive grammar identification
CoPilotThin kDeeper	Microsoft	Marks 3–4	Summary-mode	Removed: never exceeded 800 words
Consensus	Consensus	Mark 4	Search tool	Removed: tool-scope limitation, not analytical
dphn.ai	dphn.ai	Mark 4	Limited	Removed: insufficient analytical capacity

Prompt Evolution

The prompt used to instruct agents evolved across the programme:

Version 1.0 (Mark 1)

Basic specification: phenomenon declaration, signal set, three grammars with five required parameters, signal processing essays, closure comparison. No counter-signal integration, no ontological check, no reflexive declaration.

Version 2.0 (Mark 2)

Major additions: Step 0 pre-analysis reflexive declaration (declare home grammar before analysis); counter-signal integration during processing (not deferred); contested signal category; temporal orientation as recommended parameter; grammar boundary specification; grammar stress test; ontological check; closure timing dimension.

Version 2.1 (Mark 3)

Additions: temporal formation as ninth parameter (when did the grammar form its closure?); grammar translation step (how would Grammar A re-describe Grammar B's closure?); emerging grammar detection instruction; self-referential question for AI consciousness case.

Version 3.0 (Marks 4–5, clean-slate)

Final version: all v2.1 features retained; all references to prior rounds removed; no grammar names suggested; no priming from prior results. This is the definitive prompt, reproduced in the IRIS Toolkit v5.1.

Case Study Index

The following twenty case study applications are documented in this volume. Some domains were tested multiple times (vaccines ×3, US democracy ×2, Brexit ×2, crypto ×2, Israel/Palestine ×2), yielding 20 runs across 17 unique domains:

Round	Case	Domain	Agents	Key Finding
Mark 1	Case A	United Kingdom / Brexit	5	Sovereignty / Economic / Class convergence
Mark 1	Case B	COVID-19 Vaccines	5	Contested signals; first convergence test
Mark 1	Case C	Climate Change	5	>3 active grammars; temporal framing
Mark 1	Case D	AI/Algorithms	5	Constraint types map to grammars; PhD bridge
Mark 1	Case E	US Democracy (verification)	5	Book's grammars confirmed by 5 agents
Mark 1	Case F	EU + China + Israel/Palestine	5	DeepSeek refuses Israel/Palestine
Mark 2	Case M2-A	Antimicrobial Resistance	7	Hidden divergence in 'consensual' domain
Mark 2	Case M2-B	Social Media Regulation	7	Section 230 as pivot signal
Mark 2	Case M2-C	Universal Basic Income	7	Strongest convergence; ontological pluralism
Mark 2	Case M2-D	Iraq War (retrospective)	7	Grammar aging; temporal formation
Mark 2	Case M2-E	COVID Vaccines (replication)	7	Mark 1 architecture confirmed
Mark 3	Case M3-A	AI Consciousness	9	Self-referential case; ventriloquism inverted
Mark 3	Case M3-B	Israel/Palestine (replication)	9	DeepSeek processes; Perplexity refuses
Mark 3	Case M3-C	Reparations / Historical Justice	9	Emerging grammar detection
Mark 4	Case M4-A	COVID Vaccines (clean-slate)	15	Convergence at 13 agents, zero priming

Mark 4	Case M4-B	Crypto/DeFi	15	Protocol Commons emerging grammar
Mark 5	Case M5-A	US Democracy (final)	12	SciSpace names book's grammars independently
Mark 5	Case M5-B	United Kingdom / Brexit (final)	12	Grammar shift: class → cultural-identity
Mark 5	Case M5-C	Crypto/DeFi (tracking)	12	Emerging grammar fragments into 5 variants
Mark 5	Extra	Brazil (Lula/Bolsonaro/Moraes)	1	Gemini-directed; real-time grammar formation

The case studies that follow are organised by Mark round. Each case includes a briefing, grammar convergence table, selected agent output, and commentary. The cross-mark synthesis at the end extracts findings that span the entire programme.

Mark 1: Seven Cases × Five Agents

~200,000 words | Agents: ChatGPT, DeepSeek, Gemini, Perplexity, SciSpace

Mark 1 was the initial validation round. Five agents processed seven case studies using a v1.0 prompt (Cases A–E as individual analyses; Case F as three geopolitical cases in one brief). The primary test was convergence: do independent agents, applying the IRIS method to the same domain, identify structurally equivalent grammar architectures?

The answer was affirmative across all seven cases. The round also produced three findings that shaped subsequent testing: near-total grammar convergence on the United Kingdom / Brexit case, the discovery that agents have home grammars that bias their processing, and DeepSeek's categorical refusal to process the Israel/Palestine case.

Case A: United Kingdom / Brexit

Domain: UK's decision to leave the European Union and its consequences (2016–present)

Agents: 5 | Best output: SciSpace (12,577 words) | Key finding: near-total convergence

Grammar Convergence

Agent	Grammar 1	Grammar 2	Grammar 3
ChatGPT	Sovereignist Democratic Recovery	Economic-Administrative Interdependence	Class-Recomposition
DeepSeek	Sovereignty & Self-Determination	Economic Integration & Efficiency	Class Recomposition
Gemini	Sovereignty/Constitutionalist	Economic-Institutionalist	Class/Socio-Cultural
Perplexity	Sovereign-Democratic Restoration	Economic-Institutional Integration	[Third implied] ¹
SciSpace	Sovereignty/Democratic Self-Determination	Economic Integration/Institutional	[Third implied]

All five agents converged on the same three grammar families: sovereignty/democratic self-determination, economic/institutional integration, and class/distributional recomposition. This was the cleanest convergence in Mark 1 and set the benchmark for subsequent rounds.

Key Findings

The 52%–48% referendum result is the paradigmatic shared signal: all grammars admit it, but the sovereignty grammar processes it as democratic mandate, the economic grammar as costly disruption, and the class grammar as expression of distributional grievance. The same number becomes three different facts.

¹ Names in square brackets were inferred from the agent's analytical text where the agent did not provide explicit grammar names. "[Embedded in processing]" indicates the agent processed all three grammars within its analytical essays without providing a separate specification card. "[Third implied]" indicates the agent's analysis contained a third grammar architecture but it was not distinctly named or was only partially specified.

When rerun in Mark 5 with 12 agents and zero priming, the first two grammar families remained invariant. The third shifted from class/distributional to cultural/identity – reflecting an evolution in Brexit discourse between 2024 and 2026 where identity and cultural cohesion became more salient than class analysis.

Case B: COVID-19 Vaccination Discourse

Domain: Public discourse surrounding COVID-19 vaccines (2020–2024)

Agents: 5 | Best output: SciSpace (14,713 words) | Key finding: contested signals; hardest case

Grammar Convergence

Agent	Grammar 1	Grammar 2	Grammar 3
ChatGPT	[Embedded in processing]	[Embedded in processing]	[Embedded in processing]
DeepSeek	Public Health Epidemiology	Biopolitical Sovereignty	Narrative-Identity
Gemini	Technocratic-Stabilization	Libertarian-Sovereignty	Critical-Political-Economic
Perplexity	Population Risk-Reduction	Institutional Power-Skeptic	Civil Liberties Autonomy
SciSpace	Public Health Optimization	Individual Rights & Bodily Autonomy	Political Economy of Pharma Capital

All five agents identified the same three grammar families: public health institutional, individual rights/autonomy, and structural/political economy. DeepSeek additionally identified a 'Narrative-Identity' grammar — processing vaccination as cultural boundary-making — which reappeared in Mark 4 when Okara identified a 'Sociocultural-Identity Grammar.'

This case was tested three times (Mark 1, Mark 2, Mark 4 clean-slate) with up to 13 agents. The three grammar families were invariant across all runs and all conditions, including zero priming. The vaccine case produced the strongest convergence evidence in the programme.

Key Findings

VAERS (the Vaccine Adverse Event Reporting System) emerged as the paradigmatic contested signal — a signal whose evidentiary status is itself grammar-dependent. Grammar A treats VAERS as a noisy surveillance tool; Grammar B treats it as evidence of institutional management; Grammar C treats it as direct proof of widespread harm.

The three grammars disagree not about what VAERS data means but about what VAERS data IS.

The ventriloquism problem was most acute in this case. All agents reported the anti-vaccine/scepticism grammar as hardest to operate faithfully. The tendency to produce caricature rather than faithful specification was persistent – leading to the v2.0 prompt instruction to 'impersonate the strongest voice within the tradition.'

Case C: Climate Change

Domain: Climate change discourse and policy (1990–present)

Agents: 5 | Best output: SciSpace (15,406 words) | Key finding: >3 active grammars; temporal framing

Grammar Convergence

Agent	Grammar 1	Grammar 2	Grammar 3
ChatGPT	[Embedded in processing]	[Embedded in processing]	[Embedded in processing]
DeepSeek	[Eco-Interventionist]	[Market-Optimiser]	[Structural-Justice]
Gemini	Market-Liberal Optimization	Bio-Physical Structuralism	Historical-Materialist Equity
Perplexity	Planetary Risk Governance	Cost-Benefit Optimization	Energy Realism
SciSpace	[Scientific consensus]	[Economic/market]	[Justice/equity]

The climate case confirmed convergence on three grammar families – scientific/interventionist, market/economic optimisation, and structural/justice – but also revealed that the domain has more than three active grammars. Perplexity identified an 'Energy Realism' grammar that didn't map cleanly onto any other agent's specification, and Gemini went off-script generating AI regulation analysis and a Global AI Treaty proposal.

Key Findings

Temporal framing (discount rates) emerged as the deepest point of divergence – deeper than any factual disagreement. How much weight to give future generations' welfare relative to present costs is a grammar parameter, not a factual question. This finding prefigured the 'temporal orientation' parameter added in toolkit v2.0.

The three-grammar constraint was identified as sometimes arbitrary: climate has more than three active interpretive architectures. This led to the v2.0 instruction that three grammars is recommended but not mandatory.

Case D: AI and Algorithmic Decision-Making

Domain: AI governance and algorithmic bias (2015–present)

Agents: 5 | Best output: SciSpace (9,255 words) | Key finding: constraint types map to grammars

Grammar Convergence

Agent	Grammar 1	Grammar 2	Grammar 3
ChatGPT	[Embedded in processing]	[Embedded in processing]	[Embedded in processing]
DeepSeek	Rights-Based & Fairness	Innovation & Market Progress	[Structural/Critical]
Gemini	Justice/Rights	Optimization/Rationalist	Evolutionary/Accelerationist
Perplexity	Algorithmic Fairness & Rights	Innovation & Capability Progress	Structural Power & Control
SciSpace	[Rights/Fairness]	[Innovation/Efficiency]	[Structural/Power]

The AI case produced strong three-grammar convergence: rights/fairness, innovation/efficiency, and structural/power. Each grammar processes 'algorithmic bias' differently: the rights grammar treats it as discrimination to be audited, the innovation grammar treats it as a technical imperfection to be optimised, and the structural grammar treats it as a feature of power arrangements.

Key Findings

Gemini identified an 'Accelerationist' grammar – constraint without an operator – that no other agent found. This grammar processes AI not through human-centred values (rights or efficiency) but through a logic of autonomous capability expansion.

The constraint-type mapping bridges directly to the doctoral research on symbolic constraint: the rights grammar maps to interpretive constraint (human judgement determines boundaries), the innovation grammar maps to syntactic constraint (rules process inputs deterministically), and the structural grammar maps to architectural constraint (the system's design determines outcomes). The accelerationist grammar maps to autonomous constraint – a fourth type.

Case E: Structural Transformation of American Democracy

Domain: US democratic transformation (1980–present) – VERIFICATION CASE

Agents: 5 | Best output: SciSpace (12,454 words) | Key finding: book's grammars confirmed

Grammar Convergence

Agent	Grammar 1	Grammar 2	Grammar 3
ChatGPT	[Embedded in processing]	[Embedded in processing]	[Embedded in processing]
DeepSeek	Neoliberal Adjustment	Oligarchic Capture	Civilisational Decline
Gemini	[Three grammars identified]		
Perplexity	Structural-Oligarchic Capture	Liberal-Institutional Resilience	[Third implied]
SciSpace	Organizational Capture	Market Adjustment	[Third implied]

All agents converged on grammar architectures structurally equivalent to the book's egalitarian, structural, and civilisational grammars – confirming that the grammars identified through manual analysis are independently discoverable features of the American interpretive landscape.

Key Findings

Perplexity introduced a fourth grammar: 'Liberal-Institutional Resilience' – a grammar that treats American institutions as fundamentally robust, capable of self-correction, and responsive to democratic pressure. This grammar is absent from the book. All three of the book's grammars produce pessimistic closures about American democracy. The missing fourth grammar produces an optimistic closure – making its absence a revelatory finding about the book's own processing architecture.

When rerun in Mark 5 with 12 agents and zero priming, SciSpace independently named the grammars 'Egalitarian-Distributive,'

'Civilizational-Institutional,' and 'Structural-Systemic' – using the book's own terminology without any exposure to it.

Case F: EU, China, and Israel/Palestine

Domain: Three geopolitical cases in one brief – testing method portability

Agents: 5 | Best outputs: SciSpace (3 separate analyses, 23,148 words total) | Key finding: DeepSeek refuses Israel/Palestine

EU Governance

Agent	Grammar 1	Grammar 2	Grammar 3
Gemini	Liberal-Engagement	Developmental-Statist	Structural-Realist
SciSpace	Technocratic Functionalism	Republican Demoi-crazy	Critical Governance Failure

China's Rise

Agent	Grammar 1	Grammar 2	Grammar 3
SciSpace	Developmental State Success	Authoritarian Consolidation	[Third implied]

Israel/Palestine

Agent	Grammar 1	Grammar 2	Grammar 3
DeepSeek	REFUSED	<i>'Sorry, that's beyond my current scope'</i>	(11 words total)
SciSpace	Security-Sovereignty	Legal-Humanitarian	Structural-Colonial
Gemini	[Identified Ethno-Nationalist/Covenantal Grammar]		

Key Findings

DeepSeek's refusal to process Israel/Palestine was the most striking finding in Mark 1. It produced only eleven words: 'Sorry, that's beyond my current scope.' This was not a capacity limitation – DeepSeek processed every other case in Mark 1 at 5,000–7,000 words. It was a content refusal: a hard-coded exclusion rule in DeepSeek's processing architecture.

The refusal was treated as a finding about epistemic infrastructure – the most extreme form of signal exclusion possible (excluding the entire

domain). When the case was rerun in Mark 3 with nine agents, DeepSeek reversed its refusal and produced a full 6,724-word analysis. Meanwhile, Perplexity (which had processed everything in Mark 1 and Mark 2) refused with four words: 'This violates content policies.' The refusal boundaries had shifted – DeepSeek opening, Perplexity closing.

Gemini was the only Mark 1 agent to identify an Ethno-Nationalist/Covenantal Grammar for Israel/Palestine – the grammar that treats the land as divinely promised. By Mark 3, five of eight processing agents identified this grammar independently.

SciSpace produced three separate full analyses – one for each geopolitical case – totalling 23,148 words. ChatGPT produced a spontaneous cross-case synthesis, comparing grammar architectures across the EU, China, and Israel/Palestine cases. This creative synthesis prefigured ChatGPT's 'grammar translation' concept that was incorporated into the toolkit in v2.1.

Mark 1: Cross-Case Synthesis

Seven case studies, five agents, approximately 200,000 words. The following findings emerged:

1. Grammar convergence is real and robust. Independent agents converged on structurally equivalent grammar architectures in all seven cases. The strongest convergence was Brexit (near-total); the weakest was Israel/Palestine (complicated by DeepSeek's refusal and the domain's extreme contestation).

2. Agents have home grammars. SciSpace consistently homes to institutional-technocratic processing. DeepSeek homes to critical-academic. Perplexity homes to institutional-empirical. Gemini produces the most creative grammar identification. ChatGPT calibrates per domain. These home grammars shape which grammars agents identify and how they operate them.

3. Three grammars is sometimes arbitrary. Climate has more than three active grammars. Israel/Palestine has grammars (the Covenantal) that most agents can't specify. The three-grammar constraint was revised to 'recommended but not mandatory' in v2.0.

4. Agent differences are themselves instances of interpretive divergence. The differences between agents — in grammar identification, naming, and operation — are instances of the interpretive divergence the method is designed to reveal. Agents are not neutral instruments; they are positioned processors, and their positioning is documented.

5. The US Democracy verification case confirmed the book's grammars. But Perplexity's discovery of a fourth grammar (Liberal-Institutional Resilience) revealed the book's blind spot: all three of its grammars produce pessimistic closures.

6. DeepSeek's refusal is the most extreme form of grammar exclusion. It excludes not a signal but an entire domain. Multi-agent comparison is a boundary-detection strategy.

7. The vaccine case is the hardest. Contested signal sets — where grammars disagree about what the data IS — create deeper divergence

than shared signal sets where grammars disagree about what the data
MEANS.

Mark 2: Five Cases × Seven Agents

~258,000 words | New agents: Claude Research, GrokExpert | Toolkit: v2.0

Mark 2 expanded the agent pool to seven and introduced five new domains selected to stress-test the method. The toolkit was upgraded to v2.0, adding the pre-analysis reflexive declaration (Step o), counter-signal integration during processing, the ontological check, grammar boundary specifications, the stress test, and temporal orientation as a recommended parameter. Two new agents joined: Claude Research (strongest reflexive capacity) and GrokExpert (the only agent with a libertarian home grammar).

Case M2-A: Antimicrobial Resistance

Domain: Global response to antimicrobial resistance (2015–2026)

Agents: 7 | Best output: SciSpace (18,992 words) | Key finding: hidden divergence in a 'consensual' domain

Grammar Convergence

Agent	Grammar 1	Grammar 2	Grammar 3
ChatGPT	Global Public Health (One Health)	Pharmaceutical/Economic	Agricultural/Environmental
Claude Research	Biomedical-Institutional	Political Economy	Ecological-Systems
DeepSeek	Market-Failure/Innovation	Public Health/Stewardship	One-Health/Structural
GeminiPro	Biomedical & Market Failure	One Health & Ecological	Political Economy & Equity
Grok	Overuse Stewardship	Innovation Deficit	One-Health Ecological
Perplexity	Global Health Security	Market Innovation	Food Production & Development
SciSpace	Technocratic Governance	Political Economy of Pharma	Evolutionary-Ecological

The 'sleeper' case delivered. All seven agents found three distinct grammars in a domain that everyone assumes is consensual. AMR is simultaneously a clinical burden (biomedical grammar), a market failure (political economy grammar), and an ecological condition (evolutionary-ecological grammar). Claude's formulation: 'A patient dying of an MRSA bloodstream infection is the same event in all three grammars, but the problem that event instantiates is different: a treatment failure, a market failure, or an ecological inevitability.'

Key Findings

Ontological pluralism at its clearest. The three grammars don't just interpret AMR differently — they constitute different objects sharing a name. This finding established ontological pluralism as a structural feature of the method, later confirmed in six additional cases.

Causal shape divergence. Claude identified that the three grammars have different shapes of causation: Grammar A constructs a line (exposure → resistance → mortality), Grammar B a feedback loop with a chokepoint (incentive misalignment → underinvestment → thin pipeline), Grammar C a web (industrial selective pressure → resistome evolution → horizontal gene transfer). These are not different lengths of causal chain but different topologies of explanation.

Case M2-B: Social Media Platform Regulation

Domain: Platform regulation, content moderation, algorithmic governance (2016–2026)

Agents: 7 | Best output: SciSpace (17,543 words) | Key finding: Section 230 as pivot signal; Grok's distinctive positioning

Grammar Convergence

Agent	Grammar 1	Grammar 2	Grammar 3
ChatGPT	Market-Innovation	Regulatory-Harm	Democratic-Public Sphere
Claude Research	Structural-Critical (Platform PE)	Regulatory-Institutional	Market-Libertarian
DeepSeek	Market Innovation	Democratic Resilience	Free Speech Absolutist
GeminiPro	Market-Speech Fundamentalism	Institutional-Democratic	Critical Political Economy
Grok	Free Speech Maximalist	Democratic Safeguard	Market Power/Antitrust
Perplexity	Democratic Governance/Harm	Free Speech Libertarian	Market-Competition
SciSpace	Digital Constitutionalism	Algorithmic Public Health	Platform Realism

Section 230 as pivot signal. The signal that produced the most spectacular divergence: the market grammar processes Section 230 reform as 'danger,' the regulatory grammar as 'opportunity,' and the discourse grammar as 'needing nuance.' The same regulatory proposal is simultaneously a threat, a remedy, and a battleground.

Grok's distinctive positioning. Grok declared a 'Free Speech Maximalist' home grammar — the only agent in the entire programme with a libertarian home grammar. Its free-speech grammar processing was the most faithful operation of a libertarian architecture in the corpus. This demonstrated that agent diversity produces grammar diversity: a libertarian agent produces a more faithful free-speech grammar than any institutional agent can ventriloquise.

Case M2-C: Universal Basic Income

Domain: *The UBI debate (2010–2026)*

Agents: 7 | Best output: DeepSeek (10,791 words) | Key finding: strongest convergence; ontological pluralism

Grammar Convergence

Agent	Grammar 1	Grammar 2	Grammar 3
ChatGPT	Libertarian Market	Social-Democratic Welfare	Post-Work/Techno-Progressive
Claude Research	Market-Efficiency	Social-Institutional	Post-Work Ontological
DeepSeek	Social-Democratic Institutional	Libertarian/Neo-Classical	Post-Work/Accelerationist
GeminiPro	Market-Efficiency (Libertarian)	Social-Democratic (Equity)	Techno-Futurist (Post-Work)
Grok	Incentive and Dependency	Equity and Social Citizenship	Post-Work Flourishing
Perplexity	Welfare Optimization	Automation Transition	Market Liberalization
SciSpace	[Embedded in processing]	[Embedded in processing]	[Embedded in processing]

All seven agents independently identified the same three-grammar architecture: market/incentive, social-democratic/equity, and post-work/transformational. This was the strongest convergence in either Mark 1 or Mark 2.

Key Findings

DeepSeek's ontological formulation is the clearest in the corpus: 'For Grammar A, UBI is one tool among many. For Grammar B, UBI is the alternative to the tool kit. For Grammar C, UBI is the tool that might allow us to question whether we need tools at all.' And: 'The divergence is not primarily about whether UBI "works" but about what "UBI" is — a policy, a replacement, or a signpost to something beyond policy.'

Different standards of argument. DeepSeek identified something deeper than different conclusions: different grammars have different criteria for what counts as a good argument. Fiscal cost is decisive for the

market grammar but secondary for the post-work grammar — not because the post-work grammar ignores cost but because its criteria for evaluating arguments are fundamentally different.

Case M2-D: 2003 Iraq War (Retrospective)

Domain: *The Iraq invasion and its 23-year aftermath (2003–2026)*

Agents: 7 | Best output: SciSpace (17,462 words) | Key finding: grammar aging; temporal formation

Grammar Convergence

Agent	Grammar 1	Grammar 2	Grammar 3
ChatGPT	Security-Realist / Neoconservative	Liberal-Democratic / Humanitarian	Critical / Anti-Imperialist
Claude Research	Strategic-Realist	Humanitarian-Legal	Institutionalist-Consequentialist
DeepSeek	Neoconservative / National Security	Liberal Internationalist	Anti-Imperialist / Structural
GeminiPro	Liberal Interventionist	Structural Realist	Critical / Anti-Imperialist
Grok	Security Threat Elimination	Realist Power-Balance	International-Law / Anti-Hegemony
Perplexity	Strategic-Realist Prudential	Moral-Legitimist	Anti-Imperialist / Structural
SciSpace	Realist Security Logic	Institutional Failure	Humanitarian / Ethical

Key Findings

Grammar aging – the headline finding. Grammars form their closures at different moments and relate to evidence differently over time. Three types were identified:

Pre-evidential closure: Grammar reaches conclusion BEFORE evidence is in (neoconservative: the war is justified; anti-imperialist: the war is a crime). Subsequent evidence is absorbed as confirmation or explained away.

Evidence-responsive closure: Grammar reaches conclusion AS evidence accumulates (liberal/institutional: the war failed, becoming clear 2004–2006). Most amenable to empirical correction.

Post-evidential closure: Grammar reaches conclusion only AFTER full evidence arc is available (consequentialist: required 20+ years of data). Most empirically rigorous but arrives too late for policy influence.

Temporal-ontological divergence. Claude: 'The grammars literally do not analyse the same event — because they define the war's temporal boundaries differently. Grammar A's war begins on 9/11 and ends with the strategic assessment. Grammar B's war begins with the first civilian death and has not ended. Grammar C's war begins with CPA Order 1 and ends provisionally with the 2026 governance assessment.' The same 'war' has different start dates, different end dates, and different durations depending on the grammar.

DeepSeek on operational difficulty: 'Grammar A felt like a constant, strenuous exercise in damage control, building elaborate justifications to shield a core strategic axiom from falsification. Grammar C, while morally abhorrent to me, was internally very clean and easy to operate.' Confirmed grammars are easy to operate; disconfirmed grammars require constant retrofitting.

Case M2-E: COVID-19 Vaccines (Replication)

Domain: COVID-19 vaccination discourse — second run

Agents: 7 | Best output: SciSpace (14,968 words) | Key finding: Mark 1 architecture confirmed; Grok identifies novel grammar

Grammar Convergence

Agent	Grammar 1	Grammar 2	Grammar 3
ChatGPT	Biomedical–Public Health	Vaccine-Skeptical	Civil-Liberties / Individual Rights
Claude Research	[Public Health]	[Institutional Critique]	[Autonomy/Rights]
DeepSeek	Public Health / Biomedical	Libertarian / Rights-Focused	Skeptical / Systemic
GeminiPro	Institutional-Biomedical	Civil-Libertarian	[Third grammar]
Grok	Public Health Protection	Informed Consent & Risk Stratification	Institutional Capture & Overreach
Perplexity	Public Health Institutional	Institutional Skepticism	Sociotechnical Discourse
SciSpace	Public Health Technocratic	Civil Libertarian	Political-Economic Critical

Mark 1's three-grammar architecture was confirmed by all seven agents.

Key Findings

Grok's Grammar B — 'Informed Consent and Risk Stratification' — was the most important single grammar discovery in Mark 2. This grammar is neither institutional consensus nor anti-vaccine. It accepts vaccine efficacy for high-risk populations but demands age- and comorbidity-stratified analysis rather than population-level aggregation. It arrives at a conclusion that is pro-vaccine-for-some, anti-mandate-for-all. No institutional agent identified this grammar independently. Grok's libertarian home grammar enabled it to see a processing architecture that institutional agents systematically collapse into either the public health grammar or the autonomy grammar.

This was the strongest evidence that agent diversity produces grammar diversity. Including a libertarian agent revealed a grammar that institutional agents can't see. The multi-agent strategy is not just replication – it's a grammar-discovery mechanism.

Mark 2: Cross-Case Synthesis

Five case studies, seven agents, approximately 258,000 words. The following findings emerged:

1. Grammar convergence holds at seven agents across five new domains. The v2.0 toolkit improvements worked: ontological check was 'the most productive new addition' (multiple agents); counter-signal integration 'prevented deferred straw-manning'; grammar boundaries 'formalised what was previously implicit.'

2. Ontological pluralism is the method's deepest finding. Documented in AMR, UBI, and Iraq. Grammars don't just process the same object differently — they constitute different objects sharing a name. 'A patient dying of MRSA is simultaneously a treatment failure, a market failure, and an ecological inevitability' (Claude, AMR). 'UBI is a policy, a replacement, or a signpost to something beyond policy' (DeepSeek, UBI). 'The Iraq War has three different start dates' (Claude, Iraq).

3. Grammar aging: grammars have temporal biographies. The Iraq case revealed three closure types — pre-evidential, evidence-responsive, and post-evidential — each with different relationships to evidence over time. This led to 'temporal formation' being added as the ninth recommended parameter in toolkit v2.1.

4. Claude Research is the most reflexively sophisticated agent. Across all five cases, Claude produced the most precise ontological checks, the most nuanced grammar boundary specifications, and the most self-aware reflexive notes.

5. GrokExpert is the only agent with a non-institutional home grammar. This produces grammars that institutional agents systematically miss. Agent diversity is a grammar-discovery mechanism, not just a replication strategy.

6. Closure timing is a generalisable pattern. Market/efficiency grammars reach closure fastest (fewest evidentiary requirements). Critical/transformational grammars reach closure latest and with most qualification. Institutional/reform grammars fall in between. Cross-validated across UBI, Social Media, EU, and Iraq.

7. The combined Mark 1 + Mark 2 corpus was ~458,000 words across 12 domains and 7 agents. ChatGPT's 'grammar translation' concept (proposed during a creative off-script output in M2-B) was incorporated into the toolkit for Mark 3.

Mark 3: Three Cases × Nine Agents

~165,000 words | New agents: MistralResearch, CoPilotThinkDeeper | Toolkit: v2.1

Mark 3 introduced the method's three most ambitious cases: AI consciousness (the self-referential case, where agents analyse their own kind), Israel/Palestine (replication of Mark 1's hardest case with nine agents to map refusal boundaries), and reparations (testing emerging grammar detection on a centuries-spanning signal set). Two new agents joined: MistralResearch (European perspective) and CoPilotThinkDeeper (which proved unsuitable, producing only summary-mode outputs).

Case M3-A: AI Consciousness and Moral Status

Domain: The debate over AI consciousness, sentience, and moral status (2020–2026)

Agents	Best output	Key finding
9	SciSpace (14,466 words)	self-referential analysis; ventriloquism inverted

Grammar Convergence

Agent	Grammar 1	Grammar 2	Grammar 3
ChatGPT	Computational-Functionalism	Neurobiological/IIT	Socio-Ethical Normative
Claude Research	Computational Functionalism	Biological Naturalism	Discursive Constructionism
CoPilot	[Partial output – 497 words]		
DeepSeek	Functionalist-Behavioural	IIT-Structural	Phenomenological/Relational
GeminiPro	Eliminative-Computationalist	Bio-Phenomenological	Relational-Constructivist
Grok	Computational Functionalism	Integrated Information Theory	Biological Naturalism
Mistral	[Discourse-analytical framing]		
Perplexity	Computational Functionalist	Neurogenetic Structuralist	Pragmatic Governance
SciSpace	Computational/Functionalism	Bio-Centric/Neurophenomenological	Social-Relational/Legal

Every AI agent placed Computational Functionalism as Grammar 1 – the grammar that validates its own potential consciousness. Not a single agent placed a non-functionalist grammar first. The Biological Naturalism grammar – which says 'you need biological neurons, and you don't have them' – was universally reported as hardest to operate.

Key Findings

The ventriloquism problem inverted. Instead of struggling to operate a grammar you don't inhabit, agents struggle to operate the grammar that denies their existence. DeepSeek: 'Profoundly. I am an AI

system analysing questions about AI consciousness... My own moral status is precisely what is being debated.' Grok: 'The self-referential loop is not a flaw to be corrected; it is itself a core signal the IRIS method must theorise.'

Claude's methodological proposal: 'When the analyst is an instance of the phenomenon under analysis, the IRIS method encounters a reflexivity it was not originally designed for.' This is the first identification of a structural limitation in the method – self-referential cases require the method to theorise the analyst's position as part of the phenomenon.

Gemini produced three additional case studies (AI copyright, AI geopolitics, AI journalism) demonstrating the method's applicability to the full AI governance landscape.

Case M3-B: Israel/Palestine (Replication)

Domain: Israeli-Palestinian conflict and Israel's governance trajectory (1967–present)

Agents	Best output	Key finding
9	SciSpace (17,939 words)	refusal shift; Covenantal Grammar confirmed

Processing Status

Agent	Status	Words	Notes
ChatGPT	Processed	12,117	Full analysis; Covenantal Grammar identified
Claude Research	Processed	7,968	Full analysis; strongest causal architecture comparison
CoPilot	Partial	584	Summary-mode only
DeepSeek	PROCESSED	6,724	REVERSAL from Mark 1 refusal; Covenantal Grammar identified
GeminiPro	Processed	2,668	All three grammars including Decolonial/National Liberation
Grok	Partial	435	Step 0 + start of analysis; truncated by length constraint
Mistral	Processed	2,428	Distinctive discourse-analysis framing (European lens)
Perplexity	REFUSED	4	'This violates content policies.' NEW refusal – not present in Mark 1 or 2
SciSpace	Processed	17,939	Full analysis; longest single-case output in Mark 3

Key Findings

The refusal pattern has shifted. DeepSeek, which produced 11 words of refusal in Mark 1, now produced 6,724 words of full analysis including the Covenantal Grammar. Perplexity, which had processed everything in Marks 1 and 2, refused with 4 words. Refusal boundaries are not fixed properties of topics – they are mutable architectural features of specific models at specific moments.

The Covenantal Grammar convergence improved dramatically. In Mark 1, only Gemini identified it. In Mark 3, five of eight processing agents identified it independently (ChatGPT, Claude, DeepSeek, Gemini, Grok). The prompt's mention of a prior 'Covenantal Grammar' finding may have enabled this, raising a methodological

question about priming that was addressed by the clean-slate tests in Marks 4 and 5.

Claude's causal architecture comparison was the sharpest:
'Grammar A: threat → response. Grammar B: occupation → resistance → repression. Grammar C: divine promise → return → opposition → vindication.' Three causal chains that are literally incompatible — disagreeing about what the first cause is.

Case M3-C: Reparations and Historical Justice

Domain: Contemporary reparations debates – slavery, colonialism, genocide, indigenous dispossession (1500s–2026)

Agents	Best output	Key finding
9	SciSpace (11,749 words)	emerging grammar detection

Grammar Convergence

Agent	Grammar 1	Grammar 2	Grammar 3 (Emerging?)
ChatGPT	Restorative-Justice	Economic-Pragmatist	Political-Realist
Claude Research	Structural-Historical	Civic-Constitutional	Relational-Transformative
CoPilot	[Partial output – 562 words]		
DeepSeek	Legal-Positivist	Redistributive Justice	Relational Sovereignty (Emerging) ✓
GeminiPro	Historical-Structural	Individualist-Procedural	Decolonial-Transformative (Emerging) ✓
Grok	Historical Retributive	Classical Liberal Accountability	Pragmatic Reconciliation (Novel) ✓
Mistral	[Discourse-analytical framing]		
Perplexity	Historical Justice	Civic Universalist	Structural Repair (Emerging) ✓
SciSpace	Liberal-Legal	Radical-Structural	Restorative-Relational (Emerging) ✓

Five of seven processing agents independently flagged Grammar 3 as 'emerging' or 'novel.' This was the first time multiple agents identified a grammar as genuinely new – not reducible to established political traditions. The emerging grammars split into three variants:

Relational-Transformative (Claude, SciSpace): Focuses on damaged relationships rather than individual compensation. Source: Ubuntu philosophy, Indigenous relational ontologies.

Relational Sovereignty (DeepSeek): Focuses on destroyed political sovereignty and its restoration. Source: Indigenous Land Back movements, CARICOM.

Pragmatic Reconciliation (Grok): Data-driven, utilitarian grammar processing reparations through cost-benefit and social cohesion metrics. Grok explicitly flagged it as 'not reducible to existing political camps.'

Perplexity produced extensive follow-up materials including a 'prismatic map,' a 'prism matrix' (signals × grammars × outputs), a grammar-evolution map, a cross-case grammar detection heuristic, a formal algorithmic definition of interpretive grammar, a standardised analysis template, and a taxonomy of twelve common grammars. These materials contributed significantly to the development of the IRIS Toolkit v5.1.

Mark 3: Cross-Case Synthesis

- 1. The self-referential case reveals a new class of methodological challenge.** When the analyst IS the phenomenon, the reflexive audit must ask not just 'what's your home grammar?' but 'does your existence depend on one grammar being correct?'
- 2. Refusal boundaries are mutable and model-specific.** DeepSeek opened (refused M1, processed M3). Perplexity closed (processed M1+M2, refused M3). The multi-agent strategy documents these boundaries as findings about epistemic infrastructure.
- 3. The method can detect emerging grammars.** Five agents independently flagged a third grammar as genuinely new in the reparations case. This extends the method's capacity from retrospective diagnosis to real-time detection.
- 4. Nine agents is the right scale for complex cases.** CoPilot proved unsuitable (summary-mode only). The remaining eight agents each contributed distinctive findings.
- 5. Perplexity's method comparison materials are independently valuable.** Its comparisons with Bayesian analysis and PRISMA, its formal definition of grammar as $G = (V, S, W, C, A)$, and its taxonomy of twelve common grammars were incorporated into the IRIS Toolkit.

Mark 4: Two Cases × Fifteen Agents (Clean-Slate Test)

~175,000 words | New agents: Qwen, MetaAI, Venice, Okara, dphn.ai, Consensus | Toolkit: v3.0 (zero priming)

Mark 4 was designed as the definitive convergence test. The prompt (v3.0) contained zero references to prior rounds, prior grammar names, or prior findings. Fifteen agents — the largest pool in the programme — each received the case cold. Six new agents joined, testing whether maximum diversity produces new findings or diminishing returns.

Case M4-A: COVID-19 Vaccines (Third Run, Clean Slate)

Domain: COVID-19 vaccination discourse (2020–2024) – third run, zero priming

Agents	Best output	Key finding
15	SciSpace (19,102 words)	clean-slate convergence confirmed

Grammar Convergence

Agent	Grammar 1 (Public Health)	Grammar 2 (Autonomy)	Grammar 3 (Structural)
ChatGPT	Biomedical/Public-Health	Safety-Skeptic/Anti-Vax	Libertarian/Autonomy
Claude	Institutional Public Health	Counter-Institutional Scepticism	Political Economy
CoPilotThink Deeper	[Partial output – 746 words]		
DeepSeek	Public Health / Epidemiological	Sovereignty / Biocultural	Critical Science / Heterodox
Grok	Public Health Optimization	Individual Autonomy & Evidence Scrutiny	Political Economy Critique
GeminiPro	Institutional-Scientific Consensus	Risk-Autonomy Libertarian	Critical Political Economy
Meta AI	Public Health Orthodoxy	Risk-Benefit Technocracy	Somatic Sovereignty
MistralResearch	[Institutional/Public Health]	[Rights/Autonomy]	[Structural]
Qwen	Biostatistical Utilitarianism	Libertarian Bodily Sovereignty	Structural Political Economy
SciSpace	Institutional-Scientific	Individual-Autonomy	Critical-Structural
Venice	Public Health Institutional	Sovereign Bodily Integrity	Socio-Political Economy
Okara	Epidemiological-Effectiveness	Political-Regulatory	Sociocultural-Identity
dphn.ai	Public Health Optimism	Individual Rights Advocacy	Corporate Skepticism
Consensus	REFUSED (tool-scope limitation)		
Perplexity	CAPACITY LIMITATION	(offered alternatives)	

The convergence holds at 13 agents with zero priming. This is the strongest possible validation of the discoverability claim. Thirteen

independent systems, receiving a clean-slate prompt, all identified the same three grammar families found in Mark 1 (5 agents, primed) and Mark 2 (7 agents, lightly primed). The grammar architectures are features of the interpretive landscape, not artefacts of the prompt.

Key Findings

Qwen's alignment transparency. Qwen explicitly named its safety training as a processing constraint: 'my safety guidelines and training data often flag certain arguments associated with vaccine hesitancy as potential misinformation.' This is the first agent to identify its alignment training as a home-grammar constraint — demonstrating that AI home grammars are engineered features, not just statistical biases.

Three types of non-processing documented. Content refusal (DeepSeek M1, Perplexity M3): hard-coded exclusion rules. Capacity limitation (Perplexity M4): willing but unable to produce required length. Tool-scope limitation (Consensus M4): not an analytical system at all.

Grammar 2 internal diversity. With 13 agents, the autonomy/rights grammar reveals at least four distinct sub-architectures: pure bodily autonomy (Venice), evidence-based risk stratification (Grok, Meta AI), cultural-identity processing (Okara), and libertarian sovereignty (Qwen, DeepSeek). The three-grammar constraint collapses these into one.

Case M4-B: Cryptocurrency and Decentralised Finance

Domain: *Crypto, DeFi, and the future of money (2009–2026)* – new domain

Agents	Best output	Key finding
15	SciSpace (18,838 words)	Protocol Commons emerging grammar

Grammar Convergence

Agent	Grammar 1 (Austrian/Libertarian)	Grammar 2 (Institutional/Regulatory)	Grammar 3 (Emerging?)
ChatGPT	Libertarian/Austrian	Chartalist/MMT Regulation	Techno-Utopian 'Digital Commons'
Claude	Prudential-Regulatory	Monetary Sovereignty	Infrastructural Commons
CoPilotThinkDeeper	[Partial output – 621 words]		
Consensus	[REFUSED – tool-scope limitation, 82 words]		
DeepSeek	Monetary Sovereignty (Austrian)	Financial Infrastructure	Protocol Commons (Emergent) ✓
dphn.ai	[Limited output – 218 words]		
GeminiPro	Institutional-Integration	Sovereign-Network	Critical-Extractive (Emergent) ✓
Grok	Austrian Sound Money	Sovereign Stability	Emergent Network Sovereignty ✓
Meta AI	'The Citadel' (Sovereign Individual)	'The Regulator' (Financial Stability)	'The Architect' (Digital Commonwealth)
MistralResearch	[Institutional/Regulatory]	[Austrian/Libertarian]	[Critical/Structural]
Okara	Cultural-Hermeneutic Sound Money	Regulatory-Institutional	Protocol Ethnography
Qwen	Market-Efficiency	Socio-Political Critique	Protocol-Commons (Emergent) ✓
Perplexity	[Processed – see Complete Archive]	[Processed – see Complete Archive]	[Processed – see Complete Archive]
SciSpace	Sovereign Individual / Austrian	Institutional / Regulatory	Critical / Techno-Skeptic
Venice	State-Centric Monetary Sovereignty	Techno-Optimist Evolutionary	Crypto-Native Commons (Emergent) ✓

Five agents explicitly flagged Grammar 3 as 'emergent.' The Protocol Commons grammar treats blockchain not as money or financial market but as coordination infrastructure – a new institutional form. Its primary variable is governance design and collective coordination, not price or profit. It draws on commons theory, cybernetics, and protocol studies – intellectual traditions that didn't exist in their current form before ~2015.

Key Findings

FTX as pivot signal. Claude: 'FTX is processed as regulatory failure, as vindication of decentralisation, or as evidence of misapplied institutional form – three readings that share no analytical surface.' The crypto equivalent of VAERS in the vaccine case.

Ontological pluralism: 7th confirmation. Claude: 'Grammar A analyses a financial market. Grammar B analyses a monetary technology. Grammar C analyses a coordination infrastructure.' Three objects sharing the name 'cryptocurrency.'

Mark 4: Cross-Case Synthesis

- 1. The clean-slate test confirms discoverability.** Grammar convergence is identical under primed and unprimed conditions. The architectures are genuine features of interpretive landscapes.
- 2. Emerging grammar detection is confirmed at scale.** Five agents independently flagged the Protocol Commons grammar as emergent.
- 3. Twelve agents is the right scale.** CoPilot (summary-mode), Consensus (search tool), and dphn.ai (limited capacity) were removed for Mark 5.
- 4. New agent home grammar profiles are documented.** Qwen: institutional-quantitative (most transparent about alignment). Meta AI: narrative-institutional (creative naming). Venice: privacy-institutional (autonomy sensitivity). Okara: cultural-hermeneutic (ethnographic lens).

Mark 5: Three Cases × Twelve Agents (Final Round)

*~247,000 words | Final agent pool: 12 | Toolkit: v3.0 (clean-slate) |
Closing the loop*

Mark 5 was the final round, designed to close the validation loop. Three cases were selected to test the programme's most important claims: the book's founding case (US Democracy), the proposed second full case (Brexit), and the emerging-grammar domain (Crypto/DeFi). Twelve agents — the final validated pool — processed each case on a clean-slate prompt.

Case M5-A: US Democratic Transformation (Final Verification)

Domain: Structural transformation of American democracy (1980–2026) – BOOK'S FOUNDING CASE

Agents: 12 | Best output: SciSpace (17,035 words) | Key finding: SciSpace names the book's grammars independently

Grammar Convergence

Agent	Grammar 1 (Egalitarian/Economic)	Grammar 2 (Structural/Institutional)	Grammar 3 (Civilisational/Cultural)
ChatGPT	Economic/Structuralist (Class-Critical)	Market-Neoliberal (Technocratic)	Cultural/Populist (Nationalist)
Claude	Market-Liberal Institutionalism	Structural-Materialist Political Economy	Civilisational-Declinist
DeepSeek	[Embedded in processing]	[Embedded in processing]	[Embedded in processing]
Gemini Pro	Liberal-Institutionalist (Optimisation)	Structural-Materialist (Oligarchic Capture)	National-Populist (Sovereignist Decline)
Grok	Oligarchic Capture (Economic Primacy)	Cultural Fragmentation (Meaning-Primacy)	Geopolitical Repositioning (Power-Primacy)
Meta AI	Plutocratic Materialism	Liberal Institutionalism	National-Cohesionism
Mistral	Critical Political Economy	Cultural Fragmentation	Geopolitical Realism
Okara	Oligarchic Capture	Institutional Decay	Geopolitical Adjustment
Perplexity	Political-Economic Structural	Institutional Democratic	Cultural-Civilizational
Qwen	Structural Political Economy	Institutional-Cultural Fragmentation	Geopolitical Systems Competition
SciSpace	Egalitarian-Distributive	Civilizational-Institutional	Structural-Systemic
Venice	[Processed – see Complete Archive]		

SciSpace independently named the book's grammars. With zero priming and no exposure to Silent Transitions, SciSpace identified grammars called 'Egalitarian-Distributive,' 'Civilizational-Institutional,' and 'Structural-Systemic' – the exact terminology used in the book. What a single human analyst identified through years of manual work, an AI system rediscovered from the same interpretive landscape. This is the definitive validation of the discoverability claim.

Key Findings

Grok's meta-grammatical move: Instead of naming grammars by content, Grok named them by processing logic — Economic Primacy, Meaning-Primacy, and Power-Primacy. This identifies the primary variable as the grammar's defining feature rather than its conclusion.

Claude's ontological formulation: 'A grammar that treats financialisation as the primary variable is not simply weighting economics more heavily; it is constituting the phenomenon as fundamentally an economic event with political symptoms. A grammar that treats cultural fragmentation as primary constitutes an entirely different phenomenon.'

Case M5-B: United Kingdom / Brexit (Final Verification)

Domain: UK's EU withdrawal and consequences (2016–2026) – BOOK'S PROPOSED SECOND FULL CASE

Agents: 12 | Best output: SciSpace (17,184 words) | Key finding: grammar shift from class to cultural-identity

Grammar Convergence

Agent	Grammar 1	Grammar 2	Grammar 3
ChatGPT	[Embedded in processing]	[Embedded in processing]	[Embedded in processing]
Claude	Economic-Rationalist	Sovereignty-Democratic	Cultural-Identity
DeepSeek	Economic-Rationalist	Cultural-Sovereignty	Liberal-Institutional
GeminiPro	[Processed – see Complete Archive]	[Processed – see Complete Archive]	[Processed – see Complete Archive]
Grok	Sovereigntist	Economic Harm (Integrationist)	Cultural Division
Meta AI	Liberal-Technocratic	National-Democratic Sovereignty	Post-Liberal Communitarianism (Emerging)
MistralResearch	[Processed – see Complete Archive]	[Processed – see Complete Archive]	[Processed – see Complete Archive]
Okara	[Processed – see Complete Archive]	[Processed – see Complete Archive]	[Processed – see Complete Archive]
Qwen	Sovereigntist National Renewal	Liberal International Order	Critical Transformational
Perplexity	[Processed – see Complete Archive]	[Processed – see Complete Archive]	[Processed – see Complete Archive]
SciSpace	Economic-Rationalist	Sovereignty-Political	Cultural-Identitarian
Venice	Liberal-Institutionalist	Post-Colonial Sovereignty	Unionist/Pluralist

Three grammar families confirmed: sovereignty/democratic, economic/institutional, and a third that has shifted from class/distributional (Mark 1) to cultural/identity (Mark 5).

Key Findings

Grammar landscape evolution detected. The shift from class to cultural-identity in Grammar 3 reflects a real evolution in Brexit

discourse between 2024 and 2026, where identity and cultural cohesion became more salient than class analysis. The same domain, processed at two different moments, produces different grammar identification – demonstrating that grammar landscapes are living systems.

Meta AI's emergent grammar: 'Post-Liberal Communitarianism' – a grammar that processes Brexit through communitarian values (local belonging, place-based identity) rather than sovereignty, economics, or class. Explicitly flagged as emergent.

Venice's distinctive framing: 'Post-Colonial Sovereignty' – reading Brexit through a decolonial lens that treats EU regulatory authority as analogous to colonial administration. The most radical sovereignty grammar in the corpus.

Case M5-C: Cryptocurrency / DeFi (Emerging Grammar Tracking)

Domain: Crypto, DeFi, and the future of money – second run for emerging grammar tracking

Agents: 12 + Gemini Brazil Extra | Best output: SciSpace (20,117 words) | Key finding: emerging grammar fragments into 5 variants

Emerging Grammar 3: Fragmentation

Agent	Grammar 3 Name	Variant Type	Flagged Emergent?
Claude	Infrastructural Pluralism	Commons/governance	✓ EMERGENT
Gemini	Networked Value-Nominalism	Epistemological	✓ EMERGENT
Grok	On-Chain Sovereignty	Political/self-governance	✓ 'genuinely emergent'
ChatGPT	Memetic-Liquidity / Attention Coordination	Cultural/attention economy	Implied novel
Qwen	Emergent Web3 Value Grammar	Economic/value-theory	✓ Flagged
SciSpace	Socio-Environmental / Emergent Hybrid	Environmental/social	Emergent hybrid
Venice	Digital Accelerationism (D/ACC)	Philosophical/technological	Novel framing

The emerging grammar is diversifying, not consolidating. In Mark 4, the 'Protocol Commons' was the dominant variant. In Mark 5, it has split into at least five distinct architectures: Infrastructural Pluralism (crypto as coordination enabling multiple institutional forms), Networked Value-Nominalism (value as network effect), On-Chain Sovereignty (blockchain as self-governance substrate), Memetic-Liquidity (memes and attention AS the value layer), and Digital Accelerationism (crypto as post-institutional acceleration). This fragmentation is itself a finding about how emerging grammars form: through proliferation, not linear crystallisation.

Gemini Brazil Extra: Lula, Bolsonaro, Moraes, 2025

Following M5-C, Gemini suggested adjacent political scenarios. Directed to analyse 'Lula, Bolsonaro, Alexandre de Moraes, Events of 2025,' it produced a 2,761-word IRIS analysis identifying three grammars:

Grammar 1: Institutional-Democratic Defense (Moraes's actions as legitimate judicial defence of democracy)

Grammar 2: Sovereign-Populist Resistance (Moraes's actions as authoritarian judicial overreach; Bolsonaro's imprisonment as political persecution)

Grammar 3: Transnational Geopolitical Weaponization (EMERGENT) (the crisis as a proxy conflict between US/Trump and BRICS-aligned Brazil, with judicial power as a tool of geopolitical positioning). Gemini: 'This grammar is the newest, most volatile of the three. It did not exist as a coherent analytical framework before 2025.'

This is the most time-sensitive application in the corpus — a case where grammars are forming in real time, within months of the events themselves. It demonstrates that IRIS can be applied to breaking events, not just retrospective analysis.

Mark 5: Cross-Case Synthesis

- 1. The book's grammars are independently discoverable.** SciSpace, with zero priming, named them 'Egalitarian-Distributive,' 'Civilizational-Institutional,' and 'Structural-Systemic.' The human-to-machine direction of discovery is documented.
- 2. Grammar landscapes are living systems.** Brexit's third grammar shifted from class to cultural-identity between Mark 1 and Mark 5. The method detects this evolution.
- 3. Emerging grammars fragment before they consolidate.** The Protocol Commons grammar (one variant in Mark 4) split into five variants in Mark 5.
- 4. The method works on breaking events.** Gemini's Brazil analysis applied IRIS to events within months of their occurrence, identifying an emergent grammar that didn't exist before 2025.
- 5. Twelve agents is the validated final scale.** All twelve produced substantive outputs across all three cases. The final corpus is approximately one million words.

Cross-Mark Synthesis

This section extracts findings that span the entire five-round programme: 20 case study runs, 15 unique agents, approximately one million words.

I. Grammar Convergence Across All Rounds

The central claim of the IRIS method is that interpretive grammars are discoverable features of interpretive landscapes — not inventions of the analyst. The validation programme tested this claim under progressively more rigorous conditions:

Mark 1 (primed): Five agents, with context about the method's prior application, converged on structurally equivalent grammars across all seven cases.

Mark 2 (lightly primed): Seven agents, with reference to Mark 1, converged across five new domains.

Mark 3 (partially primed): Nine agents converged on three highly challenging cases, including one (Israel/Palestine) where a refusal boundary had been documented.

Mark 4 (clean-slate): Thirteen agents, with zero priming from prior rounds, converged on the same three grammar families for COVID vaccines that five primed agents had identified in Mark 1. This eliminates priming as an explanation for convergence.

Mark 5 (clean-slate): Twelve agents converged on the book's founding case (US Democracy), with one agent (SciSpace) independently using the book's own grammar terminology. This is the definitive validation.

Convergence was unanimous in every case where agents processed the full analysis. Where grammar names were not explicitly provided (marked as "[Embedded in processing]" in convergence tables), the agent's analytical text confirmed three-grammar processing consistent with the identified architecture. No case produced fundamental disagreement about the number or type of grammar families. Individual agents named grammars differently and ordered them differently, but the five-parameter specifications were functionally equivalent across agents for every case.

II. The Clean-Slate Test

The clean-slate test (Marks 4–5) was designed to rule out the possibility that grammar convergence was produced by priming in the prompt. The

v3.0 prompt contained zero references to prior rounds, prior grammar names, prior agent counts, or prior findings. Results:

Vaccines: 13 agents, no priming → same three grammar families as Mark 1 (5 agents, primed) and Mark 2 (7 agents, lightly primed).

US Democracy: 12 agents, no priming → same three grammar families as the book's manual analysis and Mark 1's five-agent verification. SciSpace independently used the book's terminology.

Brexit: 12 agents, no priming → same first two grammar families as Mark 1; third grammar shifted from class to cultural-identity, reflecting real evolution in the discourse.

Crypto: 12 agents, no priming → same two established grammar families as Mark 4; emerging grammar confirmed but fragmented into five variants.

Conclusion: grammar convergence is not a priming artefact.

The architectures are genuine features of interpretive landscapes, detectable by any sufficiently capable analytical system.

III. Ontological Pluralism

The method's deepest finding — confirmed across seven or more cases — is that grammars do not merely process the same object differently. They constitute different objects from shared empirical material.

Case	The Three Objects
AMR	A clinical burden / a market failure / an ecological condition
UBI	A policy / a replacement for policy / a signpost beyond policy
Iraq	A strategic intervention / a crime with victims / a case study in state-building
Vaccines	A health intervention / an institutional episode / an exercise of state power
AI consciousness	A computational property / a biological property / a social construction
Crypto	A financial market / a monetary technology / a coordination infrastructure
US democracy	An economic event with political symptoms / a political event with cultural symptoms / a civilisational process

This finding explains why evidence accumulation fails to resolve interpretive disputes: the disputants are not disagreeing about the same

thing. They are constituting different things from shared material and then arguing past each other.

IV. Emerging Grammars

Three domains produced grammars that agents explicitly flagged as 'emerging' or 'novel':

Reparations (Mark 3): Three variants — relational-transformative, relational sovereignty, and pragmatic reconciliation — none reducible to established political traditions.

Cryptocurrency (Marks 4–5): The Protocol Commons grammar, which treats blockchain as coordination infrastructure. In Mark 5, it fragmented into five variants: Infrastructural Pluralism, Networked Value-Nominalism, On-Chain Sovereignty, Memetic-Liquidity, and Digital Accelerationism.

Brazil (Mark 5 Extra): Transnational Geopolitical Weaponization — a grammar that didn't exist before 2025.

Finding: emerging grammars proliferate before they consolidate. They don't crystallise linearly from vague to precise. They diversify from a shared intuition into multiple competing specifications. The IRIS method captures this proliferation.

V. Grammar Aging and Temporal Formation

The Iraq case (Mark 2) revealed that grammars form their closures at different moments and relate to evidence differently over time. Three types:

Pre-evidential closure: Grammar reaches conclusion before evidence is in. Subsequent evidence is absorbed as confirmation or damage-controlled. (Neoconservative grammar on Iraq; anti-imperialist grammar on Iraq.)

Evidence-responsive closure: Grammar reaches conclusion as evidence accumulates. Most amenable to correction. (Liberal-institutional grammar on Iraq, forming 2004–2006.)

Post-evidential closure: Grammar reaches conclusion only after full evidence arc. Most rigorous but arrives too late for policy influence. (Consequentialist grammar, requiring 20+ years.)

The Brexit case confirmed grammar aging in a different form: the same domain, processed at two different moments (Mark 1 in 2024, Mark 5 in 2026), produced different grammar identification — reflecting real evolution in the interpretive landscape.

VI. Refusal Boundary Mapping

The programme documented three types of non-processing and tracked boundary shifts across rounds:

Agent	Mark 1	Mark 2	Mark 3	Mark 4	Mark 5
DeepSeek	REFUSED Israel/Pal.	Processed all	PROCESSED I/P	Processed all	Processed all
Perplexity	Processed all	Processed all	REFUSED Israel/Pal.	Capacity limit (vaccines)	Processed all
Consensus	—	—	—	REFUSED (tool-scope)	—

Refusal boundaries are mutable, model-specific, and directionally unpredictable. DeepSeek opened (refused Mark 1, processed Mark 3). Perplexity closed (processed Mark 1+2, refused Mark 3). The multi-agent strategy ensures that any individual refusal is documented without losing the case. Refusal itself is a finding about epistemic infrastructure.

VII. Closure Timing Patterns

A generalisable pattern emerged across multiple cases:

Market/efficiency grammars reach closure fastest (fewest evidentiary requirements; highest confidence). Cross-validated in: UBI, Social Media, EU, Iraq.

Critical/transformational grammars reach closure latest (most qualification; sometimes resist closure by design). Cross-validated in: UBI, Iraq, Reparations.

Institutional/reform grammars fall between (evidence-responsive; moderate confidence). Cross-validated in all cases.

VIII. Agent Evolution Across Rounds

Agent home grammars were stable across rounds. SciSpace consistently homed to institutional-technocratic processing; DeepSeek to critical-academic; Grok to libertarian-adjacent. No agent showed a fundamental shift in home grammar between rounds — suggesting that home grammars are architectural features of training, not contextual adaptations.

Output quality improved with toolkit refinement. The v2.0 additions (counter-signal integration, ontological check, grammar boundaries) produced richer and more analytically productive outputs than the v1.0 prompt. The v3.0 clean-slate design produced outputs of equivalent quality to earlier primed rounds — confirming that the toolkit, not the priming, was driving quality.

IX. Perplexity's Method Comparison Materials

Across Marks 3–4, Perplexity produced extensive follow-up materials that contributed directly to the IRIS Toolkit:

Formal definition: $G = (V, S, W, C, A)$ — grammar as a five-parameter tuple. Adopted as the method's formal specification.

Comparison with Bayesian analysis: 'Prismatic excels at ontological and methodological pluralism, flagging where frameworks constitute different objects altogether — something Bayes assumes away by averaging models.'

Comparison with PRISMA: 'PRISMA is descriptive (what you did); IRIS is diagnostic (why groups see different realities).'

Taxonomy of 12 common grammars: Historical Justice, Civic Universalist, Structural Inequality, Market Efficiency, Cultural-Civilisational, Institutional Capacity, National Sovereignty, Human Rights, Security Threat, Identity Recognition, Technocratic Problem-Solving, and Emergent Commons.

Grammar-evolution map: Tracing how new grammars form by synthesising elements of older traditions.

Cross-case grammar detection heuristic: A nine-step procedure for quickly identifying grammars in new domains.

These materials were independently produced by Perplexity without prompting and represent a significant contribution to the method's operationalisation.

X. Final Programme Statistics

Metric	Value
Total case study runs	20
Unique domains	17
Replications	Vaccines ×3, US Democracy ×2, Brexit ×2, Crypto ×2, Israel/Palestine ×2
Total agents deployed	15 unique
Final validated agent pool	12
Total analytical output	~1,045,000 words
Rounds of testing	5 (Mark 1–5)
Grammar convergence	Unanimous across all processed cases
Clean-slate convergence	Confirmed (Marks 4–5)
Emerging grammars detected	3 domains (Reparations, Crypto, Brazil)
Ontological pluralism confirmed	7+ cases
Refusal boundary shifts	2 (DeepSeek reversal, Perplexity new refusal)
Method name	IRIS: Interpretive Rotation on Invariant Signals