The Architecture of Discontent A Unified Theory of Political Expectation Failure

J.Konstapel,Leiden,25-6-2026

Summary

The gap. There is a growing gulf between what citizens expect from government and what politicians can actually deliver.

Powerlessness. Governments have progressively less grip on complex global problems such as climate change and the economy, while citizens demand fast solutions.

The danger. This constant disappointment produces deep distrust in democracy and gives populism room to grow.

Speed. Our "on-demand" culture collides head-on with the slow, viscous reality of political decision-making.

What is new in this edition. The original essay argued that democracies are trapped because they cannot complete the learning cycle that healthy cognitive systems use to revise failed expectations. This edition keeps that diagnosis but strengthens it from two independent directions, read together. The first is five established external research traditions — Wlezien's thermostatic model, predictive processing and the free energy principle, punctuated equilibrium theory, valence politics and "state valence failure," and Soroka's negativity bias research — each of which gives one of PEFT's three pillars a falsifiable boundary or an existing measurement instrument. The second is a systematic full-text reading of the author's much larger blog corpus, not filtered by whether a post names PEFT, Schank, or De Vries, since the corpus develops the same underlying pattern under many different vocabularies across nineteen years (2007–2026). This reading surfaced a working operational layer the original essay only gestured toward: concrete monitoring instruments (a Panarchy Monitor, a Socrates Monitor, the SWARP platform's case-based gating mechanism), a formal threshold condition for when an institution has enough internal variety to revise itself at all, a four-organ constitutional design for closure (the "quaternion state") sharpened by a documented historical mechanism (constitutional co-optation) and a documented contemporary danger (democracy dismantled in its own name) for why its fourth organ specifically has never formed, a proven totality theorem for honest non-answering that already runs in a working citizen-facing tool (QUO?), an industrialized account of expectation construction via political microtargeting, concrete Dutch case material (the VVD's 2005 manifesto, the Leiden real-estate "revolving door," structural waiting times across nearly every public domain, a working adaptive-governance counter-example in water-safety policy), and a serious counter-position from agonistic democratic theory (Mouffe) that closure itself can be a form of illegitimately suppressing political conflict.

Core message. Calm will only return to society once we bring our expectations back into line with political reality — and once we can show, with evidence, exactly where, why, and for whom that realignment currently fails, what kind of institution could let it succeed, and what a legitimate disagreement looks like that closure should not try to erase.

Known gaps. This edition's full-text corpus reading prioritized the 555 posts (of 999 total since 2007) whose content matches PEFT's underlying conceptual pattern, beginning with the most heavily-matching posts; it has not yet covered the full set, and is being completed in phases across separate sessions rather than delivered as repeated partial updates within one. Separately, posts published 18–23 June 2026 — including a piece on the House of Orange's entourage as a question of institutional counter-power — have not yet been read at all. Both gaps are flagged here explicitly rather than left for the reader to discover.

Abstract

Across the globe, democratic systems are beset by a convergent crisis characterized by declining institutional trust, the rise of populist movements, and persistent voter dissatisfaction. Despite a wealth of empirical data, existing explanations remain fragmented across disciplines. This essay synthesizes these perspectives into a unified framework, Political Expectation Failure Theory (PEFT). By integrating Roger Schank's cognitive theory of script-based expectation with Catherine E. De Vries's empirical research on political dissatisfaction and complexity science's models of non-learning systems, this essay argues that modern democracies have become institutionalized systems of expectation failure without closure. They are structures that not only fail to resolve the gap between citizen expectations and political reality but actively reproduce this gap as a political resource.

This fourth edition extends the original argument along two independent tracks that converge on the same conclusion. The first integrates five external research traditions — Wlezien's thermostatic model of opinion-policy feedback, predictive processing and the free energy principle, Baumgartner and Jones's punctuated equilibrium theory, Clarke, Stewart, and Whiteley's valence politics theory, and Soroka's negativity bias program — each of which gives one of PEFT's three pillars a falsifiable boundary or an existing measurement instrument that the 2026 original asserted but did not yet demonstrate. The second track is the product of reading the author's complete blog corpus published since the original essay (102 posts, March–June 2026) without restricting attention to passages that name PEFT explicitly, since the same underlying pattern recurs across the corpus under different vocabularies — cybernetics, nilpotent algebra, TRIZ, Human Design — and several of the richest extensions of PEFT do not use its name at all. This track contributes a formal variety threshold below which an institution cannot revise its own foundations regardless of intention; a four-organ constitutional design (the "quaternion state") that operationalizes closure as a structural property rather than a hoped-for virtue; a proven totality theorem showing that a system can treat every possible question honestly without being able to answer every question, now running as a working citizen-facing instrument (QUO?); concrete Dutch empirical material — a twenty-year natural experiment in single-loop non-learning, a documented real-estate "revolving door," and municipal chatbot failure rates — that grounds the theory in named, checkable cases rather than abstraction; and a serious internal challenge from Chantal Mouffe's agonistic democratic theory, which the essay does not resolve but states plainly, since a theory that only ever recommends more closure risks becoming exactly the kind of antagonism-suppressing technocracy its own diagnosis warns against.

1. Introduction: The Crisis of a Broken Promise

The contemporary democratic malaise is no longer a matter of academic conjecture but a documented global reality. From the erosion of voter turnout in established Western democracies to the proliferation of anti-system parties in emerging ones, a pervasive sense of institutional failure has taken hold. Standard explanations — economic inequality, cultural backlash, elite disconnection — each capture a piece of the puzzle. Yet they struggle to explain the self-reproducing nature of the crisis: why do the very mechanisms designed to correct democratic dysfunction often generate new, more profound forms of dissatisfaction?

This essay proposes a fundamental reorientation. The central unit of democratic order is not the institution, the party, or even the aggregated citizen preference, but rather the *expectation*. The core argument is that a structural condition has emerged — Political Expectation Failure (PEF) — in which democratic systems continuously generate, violate, and politically exploit citizen expectations without ever achieving cognitive or institutional closure. This condition is not a failure of any single policy or leader but a systemic feature, a perverse equilibrium where the management of the gap between what citizens anticipate and what they experience has become the primary dynamic of political competition.

This unified theory, Political Expectation Failure Theory (PEFT), is built on a tripartite foundation. First, it draws on the cognitive architecture of Roger Schank to explain how citizens form and process expectations. Second, it leverages the extensive empirical program of Catherine E. De Vries, which has meticulously documented the mechanics of political dissatisfaction across Europe. Third, it employs insights from complexity science to understand why democratic systems, designed for stability, are structurally incapable of learning from their own failures.

Each of these three pillars, however, has historically rested more on analogy and qualitative description than on a tested mechanism. This edition closes that gap. The cognitive pillar gains a formal account, drawn from predictive processing, of *why* some expectations harden into unrevisable grievance while others do not. The empirical pillar gains a direct theoretical neighbor — valence politics theory — that has already measured something very close to political expectation failure for thirty years. The systemic pillar gains both a falsifiable boundary condition (the thermostatic model tells us exactly where non-learning should and should not occur) and a far more precisely specified mechanism for institutional non-learning (punctuated equilibrium theory's account of bounded institutional attention, in place of the looser panarchy metaphor). A sixth addition, Soroka's negativity bias research, supplies an asymmetry that the original cycle described only implicitly: failure is remembered roughly twice as vividly as fulfillment, which is what makes the cycle ratchet toward accumulating discontent rather than simply oscillate.

The resulting framework offers explanatory power across cognitive, political, and systemic levels, with profound implications for how we understand and potentially reform democratic governance — and, for the first time, a set of concrete, testable predictions that did not follow from the theory's first formulation.

1.1 The Deeper Pattern: An Outside Story and an Inside Story, and a Five-Thousand-Year Timeline

Before turning to PEFT's three pillars individually, it is worth making explicit a pattern that runs underneath all of them, because it explains why the pillars needed strengthening in the same direction at once, and because the author's own corpus states it more directly outside the political essays than within them.

There are, in the author's broader work, two ways to tell the story of any self-organizing system — a bacterium, a skull, a language, a democracy. The *outside story* explains the system as the product of external forces acting on a passive substrate: selection pressure shapes the organism, incentive shapes the voter, institutional design shapes the citizen. The *inside story* explains the same system as actively probing its environment, maintaining its own coherence, and crossing internal thresholds into new configurations — not because something pushed it, but because it is the kind of system that selects its own stable states. Neither story is false. The outside story is not wrong about selection pressure, incentive, or institutional design; it is incomplete, because it has no place for the system's own activity in stabilizing what it becomes.

This distinction has a precise physical anchor that the author's work returns to across domains far from politics: James Clerk Maxwell's original 1865 formulation of electromagnetism was written in quaternion algebra and included a scalar component, alongside the vector components, describing phase relations and coupled oscillation rather than only directional force. In 1884, Oliver Heaviside dropped the scalar component to simplify the mathematics for engineering use — a choice that built the motors and launched the radio stations, and that also led a generation of physicists to conclude, wrongly, that the scalar component had described nothing real, since it no longer appeared in the equations they used. What was lost was not a mechanical detail but an entire register: the physics of phase coherence, of relation and resonance, as opposed to the physics of directional force acting on separable objects.

The same reduction, the author's corpus argues, happened at civilizational scale: the explicit suppression, beginning around 1500 CE and accelerating through the colonial era, of knowledge traditions organized around coherence, resonance, and relation — practices for stabilizing meaning and community that did not operate through force, measurement, or external causation, and that were dismantled specifically because they were illegible to, and threatened the institutional authority of, a worldview built entirely on the vector layer. The result, on this reading, was a civilization extraordinarily competent at controlling the external world and extraordinarily incompetent at maintaining internal coherence — its own, or its citizens'.

The author's corpus dates this pattern far more precisely than "civilizational scale" suggests, and the dating matters for how PEFT's diagnosis should be read. Genetic evidence places the deepest surviving human lineage among the San peoples of southern Africa at roughly 200,000–300,000 years, with click-consonant languages, trance-based coherence practice, and an egalitarian, low-concentration social organization documented as the oldest continuously practiced human knowledge tradition. Migration out of Africa roughly 65,000 years ago carried a functionally identical knowledge system to Australia (Aboriginal songline navigation, independently verified through precision survey of stone arrangements) and, via the Andaman Islands, to the Indian subcontinent, where its most complete archaeological expression is the Indus Valley Civilization (roughly 3300–1300 BCE): a society of more than a million people, larger than contemporaneous Egypt and Mesopotamia combined, in which over a century of excavation across hundreds of sites has found no fortifications with a military function, no weapons caches, no palaces built for conquering kings, no mass graves indicating warfare, and no iconography of domination — the only large-scale, multi-century urban civilization on the archaeological record that functioned without centralized power concentration, organized instead around standardized weights and measures, cities built on precise astronomical alignment, and a meditation posture appearing on Indus figurines that is structurally identical to the *kāyotsarga* standing meditation still practiced in Jain ritual today.

On this account, the pattern broke at a specific, dateable point rather than declining gradually: a documented abrupt climate event around 3117 BCE — the Piora oscillation, coinciding with sediment layers across multiple Middle Eastern excavation sites and a candidate impact crater southeast of Madagascar — desertified the then-fertile Green Sahara within decades, forcing populations into river valleys and oases where competition over suddenly scarce resources replaced the cooperative social organization the same corpus identifies as the symmetric, low-concentration coupling structure shared by San, Aboriginal, and Indus societies alike. The Sumerian King List records the resulting transition directly: kingship, it says, was "lowered again from heaven" after the flood — the shaman, in this reading, replaced by the priest, direct access to coherent experience replaced by institutional mediation, and the first concentration of communication infrastructure, food redistribution, and military capacity into what the corpus calls deep, hard-to-escape attractor wells. This is named, in the corpus, the first epistemocide. The colonial-era suppression described above — the burning of Maya manuscripts in 1562, the prohibition of Aboriginal songlines and forced removal of children beginning in 1788, the execution of an estimated 40,000–60,000 bearers of pre-Christian European coherence practice during the witch trials of 1400–1800, and the "kill the Indian, save the man" residential school policy pursued in North America from 1840 to 1980 — is, on this account, the *second* epistemocide: not the origin of the pattern PEFT diagnoses, but its most recent and most thoroughly documented repetition.

This reframes what PEFT's original three pillars are actually describing. The claim is not that modern democratic institutions have a design flaw that happens to produce expectation failure. The claim is that modern democratic institutions are the present-day surface expression of a roughly five-thousand-year-old shift from coherence-based, symmetrically-coupled social organization to control-based, asymmetrically-coupled social organization — a shift with a specific historical trigger, not a permanent feature of human nature — and that the specific mechanism PEFT names (expectation construction and exploitation without closure) is simply what that older shift looks like when it operates through electoral competition rather than through priesthood, conquest, or colonial administration. This does not change any of PEFT's formal claims in Sections 2 through 7 below; it changes their scope. If the corpus's dating is right, the relevant comparison class for "what a coherent, low-concentration, self-correcting society looks like" is not a hypothetical future reform but a documented historical case — the Indus Valley Civilization specifically — which makes the quaternion state's design goal (Section 5.3) a restoration project rather than a speculative invention, and which is why this edition treats *aparigraha*, the Jain principle that irreversible power concentration degrades the coherence of the collective field, as an architectural constraint on institutional design rather than only an ethical ideal: the corpus derives the same constraint independently from quaternion algebra applied to social eigenstates, and treats the Indus Civilization's two-millennium stability as the empirical demonstration that the constraint is satisfiable at civilizational scale, not merely a theoretical nicety.

This is the pattern PEFT's three pillars are independently rediscovering, in political-science vocabulary, without naming it as such. The original essay's complexity pillar (Section 5) describes democratic institutions as designed for stability rather than learning — a vector-layer description, in which the citizen is a force to be managed, an electorate to be moved, an expectation to be constructed and exploited from outside. PEFT's central diagnosis — that this management produces compounding failure rather than resolution — is precisely the prediction the inside/outside distinction generates: a system that treats its citizens only as objects to be acted upon by policy, incentive, and campaign messaging, and never as agents actively maintaining their own coherent model of how governance should feel and seeking thresholds at which that model can legitimately update, will reliably reproduce the very failure it is trying to manage away. Section 2.1's predictive-

processing addition — the citizen as a Bayesian agent minimizing free energy, maintaining a generative model of "what good governance feels like" and updating it on disconfirming evidence — is, read this way, not an import from cognitive science so much as a restoration of the scalar register Heaviside cut from the equations: it reinstates the citizen's own coherence-seeking activity as a first-class object of political theory, rather than treating political behavior as wholly explicable from the outside, in terms of incentives and forces alone. Section 5.3's quaternion state pursues the same restoration institutionally: its four organs are an attempt to give the *inside* story — citizens' own disconfirming experience, civil society's own constituting activity — actual constitutional weight against a governing architecture that, like Heaviside's equations, currently has no term for it at all.

A note on the lineage of this argument is worth making explicit, because it bears on how seriously the deep timeline above should be taken. The distinction between coherence-preserving and control-imposing social organization is not a claim invented to support PEFT after the fact; an unpublished 2006 working paper in the same author's corpus, predating both PEFT and the SWARP platform by two decades, develops the identical structural distinction from comparative anthropology rather than from quaternion algebra or genetics. Surveying the Medicine Wheel tradition, that paper observes that early hunter-gatherer cosmologies organized their cyclical calendars around a fixed, unmodified natural center — what it calls "preserving the centre" — and that the recurring failure mode across subsequent civilizations was specifically to draw that center *into* the cycle and adjust it to short-term need, with the explicit, named consequence that "many cultures left the earth behind as a desert." The same 2006 paper independently identifies Aristotle, rather than Heaviside, as the historical figure who performed the reduction Section 1.1 above describes in physical terms: collapsing Pythagorean and Platonic cyclical, multi-sphere cosmology to a single causal line and a single material sphere, a move the paper credits with founding Western materialism and with making cyclical, harmonic thinking about reality something that subsequently had to survive underground, in suppressed or marginalized traditions, rather than at the center of Western thought. That a structurally identical argument was independently derivable in 2006 from comparative religious anthropology and economic cycle theory, and again in 2026 from genetics, archaeology, and quaternion algebra, is itself a form of the convergent validation Section 1.1's sources claim for the older traditions — though it should be read as two passes of the same author refining one underlying intuition over twenty years, not as two independent research programs arriving at the same conclusion by chance.

This is also why Section 2.3's challenge from Mouffe cannot be waved away by appeal to the rest of the framework: a theory built to restore the inside story to political analysis must take special care that its own remedy — closure, revision, coherence — does not itself become a new vector-layer imposition, a force applied from outside to make citizens' disagreements resolve on schedule. The distinction defended in that section, between performance and procedural scripts (where closure is appropriate) and relational and value disagreement (where it is not), is the place this edition draws that line, precisely because the inside story does not predict that all internal activity converges on agreement — only that it seeks its own stable coherence, which for plural societies may permanently include standing, legitimate disagreement as one of its stable states.

1.2 Method of This Review

This edition's second track was assembled by reading the author's full blog corpus published since the original essay, rather than by searching it for occurrences of "PEFT," "expectation failure," "Schank," or "De Vries." An earlier pass at this material used exactly that keyword search and

missed most of what mattered, including the material in Section 1.1 above, because the corpus's habit is to develop the same underlying pattern under a different name in nearly every domain it touches — cybernetics, nilpotent algebra, Human Design, TRIZ, comparative mythology — and several of its richest extensions of PEFT do not use political-science vocabulary, or the theory's own name, at all. The working rule adopted for this edition is accordingly not "does this passage name PEFT" but "does this passage describe the same structural pattern PEFT describes, regardless of the domain or vocabulary in which it is dressed." This is a slower method and a more error-prone one, since it requires a judgment call on every post rather than a clean keyword match, but the alternative was demonstrated, in the course of producing this edition, to silently discard exactly the passages a keyword filter is least equipped to find.

A note on corpus coverage. The second track of this edition draws on a full reading of the author's blog corpus published between the original essay (28 March 2026) and 17 June 2026 inclusive — 102 posts, read in full rather than filtered by keyword, for the reasons given in Section 1.2 below, and including posts that never mention PEFT or expectation failure by name, such as the source for this section. The export this reading was based on does not extend past 17 June, and a further cluster of posts published 18–23 June 2026 — including, on their titles, material on the House of Orange's entourage as an institutional counter-power question and a piece connecting ADHD to systemic/connective-tissue research — has not yet been read or integrated. This edition should therefore be read as current through 17 June 2026, with a known and explicitly flagged gap for the six days following, to be incorporated in a subsequent revision rather than silently absorbed into this one.

2. The Cognitive Foundations: Scripts, Schemas, and Political Reality

The starting point for PEFT is a cognitive one. Citizens do not approach political life as blank slates. They possess dense, historically accumulated mental structures — what Roger Schank and Robert Abelson termed *scripts*. These scripts are stereotyped sequences of events that encode what should happen in familiar contexts, from a visit to a restaurant to the functioning of a democracy. When reality deviates from the script, a cognitive failure is registered, triggering a search for explanation and, ideally, a revision of the script itself.

In the political domain, these scripts operate at three nested levels. *Procedural scripts* encode expectations about the political process itself: elections should be fair, debates substantive, and promises honored. *Performance scripts* concern policy outcomes: economic growth should improve living standards, and public services should deliver. *Relational scripts*, the deepest and most emotionally charged, define the perceived compact between citizen and state: representatives should listen, institutions should respond, and collective decisions should reflect shared values.

A crucial insight is that these expectations are not absolute but comparative. Citizens evaluate political performance not against an abstract ideal but against reference points derived from prior experience or elite-constructed narratives. De Vries's benchmark theory of European public opinion exemplifies this: citizens judge the European Union against the script of their national political experience, a benchmark that often sets expectations the EU is structurally unable to meet. This comparative structure is politically potent. Actors who can manipulate the reference frame can induce or suppress the experience of failure without altering objective performance, making expectation management a central tool of political strategy.

2.1 What Was Missing: A Formal Mechanism for Why Scripts Resist Revision

Schank's script theory is rich descriptively, but it was never formalized as a quantitative model, and it leaves an important question unanswered: why do some citizens revise their political scripts in light of experience, while others entrench them ever more deeply against the same disconfirming evidence?

The **predictive processing** framework, formalized through Karl Friston's free energy principle, supplies this missing mechanism. It treats belief updating as *precision-weighted error minimization*: an agent holds prior expectations, encounters a prediction error — a mismatch between expectation and observed outcome — and updates its beliefs in proportion to the *precision*, or confidence weighting, assigned to that error relative to the prior. Crucially, updating is conditional. When a prior is held with very high precision, contrary evidence is discounted rather than integrated, no matter how often it recurs.

Recent work on ideology and predictive processing argues that this same architecture governs political belief formation. Group members converge on shared priors about how the political world works, and because these priors are socially validated and reinforced by in-group communication, their precision rises over time — making them progressively more resistant to revision by genuinely disconfirming evidence. This is precisely the mechanism that produces, in PEFT's own terms, scripts that survive repeated, accumulating failures without ever being revised. It also sharpens the role of the political entrepreneur, discussed in Section 6: the entrepreneur's most effective rhetorical move is not merely to *name* a failure but to inflate the perceived *precision* of the prediction error — making the gap between promise and reality seem unambiguous and beyond honest dispute — while simultaneously reinforcing the precision of the prior itself, so that the only available "revision" is toward a new, equally high-precision expectation, rather than toward a more accurate, lower-confidence model of political reality.

This formal addition should be read as a heuristic for generating testable predictions about political belief revision, not as a settled claim about neural mechanism — the free energy principle remains contested even within cognitive science as an explanatory rather than a merely descriptive framework. But as a heuristic, it converts Schank's qualitative sequence (explanation → reminding → generalization → revision) into something with an actual lever: precision.

A second, independent formalization of the same gap comes from organizational learning theory. Argyris and Schön's (1978) distinction between *single-loop* and *double-loop* learning gives a precise vocabulary for what the original essay described only qualitatively as the difference between revising a script and defending it. Single-loop learning corrects behavior within an unexamined set of governing assumptions — tightening a rule, adding a procedure, intensifying enforcement. Double-loop learning revises the governing assumptions themselves. The crucial diagnostic claim is that political and bureaucratic systems facing expectation failure default almost without exception to single-loop correction: when a script fails, the system writes a new rule rather than asking whether the underlying model of the problem was wrong in the first place. This is not a contingent political failing but, as the cybernetic literature on requisite variety also suggests, close to a structural property of systems whose mechanisms for self-confirmation outweigh their mechanisms for self-correction — and it gives PEFT's central concept of "closure" an operational test: closure has occurred only when an institution's *governing assumptions*, not merely its rules, have changed in response to a documented expectation failure.

2.2 What Was Also Missing: Why Different Citizens and Cultures Fail Differently

PEFT's original cognitive pillar treats script failure and non-revision as a roughly uniform phenomenon across citizens. The Fractal Karma model — developed in the author's subsequent work on the SWARP platform — adds a further layer: it specifies *which* failure mode a given citizen, organization, or political culture will display, and argues that this is not random but structurally determined by a dominant cognitive orientation, formalized as a normalized quaternion over McWhinney's (1997) four irreducible worldviews:

$$q = w_B \cdot 1 + w_R \cdot i + w_G \cdot j + w_Y \cdot k$$

where Blue (analytic/unitary) weights institutional integrity and rule-following; Red (sensory) weights direct physical and financial feedback; Green (social) weights relational consensus; and Yellow (mythic) weights narrative coherence and visionary framing. The model's central claim is that each orientation has a *characteristic* failure topology, fractally self-similar from the individual to the political-cultural scale: a Blue-dominant system fails by holding onto an outdated script for too long, producing institutional or systemic collapse rather than revision; a Yellow-dominant system fails through a cycle of inspired action and structural collapse, unable to revise its core narrative; a Green-dominant system fails specifically at the *retrieval* step — it repeatedly fails to apply lessons from its own institutional case history, so the same relational impasse recurs; and a Red-dominant system breaks the cycle by force, externalizing failure and refusing the vulnerability that script revision would require.

This matters for PEFT because it predicts that the four structural reasons for non-learning identified in Section 5 below (incentive misalignment, temporal mismatch, cognitive asymmetry, institutional rigidity) should not weigh equally across all political actors and cultures. A Blue-dominant institutional culture — a strong fit for much of Dutch administrative governance, discussed in Section 5.2 — should be disproportionately vulnerable to institutional rigidity specifically, since its characteristic defensive reflex is *intensification*: when rules fail, write more rules; when procedures falter, design better procedures; when oversight falls short, appoint an overseer of the oversight. This is not a design error within a Blue-dominant system; it is the only response such a system is structurally capable of producing, which is precisely why Section 5.2's case study shows a self-diagnosed reform program executing only the portion of its own prescription that concentrated power, while every portion that would have devolved it quietly died.

The model adds a further mechanistic layer via the Enneagram, used not as a personality typology but as a map of the specific *defensive strategies* that prevent the Schankian cycle (Section 2 above) from completing for each dominant quaternion component. Blue-dominant types characteristically intensify existing scripts or stockpile knowledge as a buffer against implementation, rather than testing the script against a failure; Yellow-dominant types revise the narrative without revising the underlying structure, or aestheticize failure into identity rather than correcting it; Green-dominant types prevent a failure from ever being registered as their own failure in the first place; and Red-dominant types short-circuit the cycle at the expectation stage itself through projected force, pre-empting the discomfort of being wrong by acting before disconfirmation can register. This gives PEFT's account of institutional non-learning (Section 5) a specific, falsifiable prediction about *mechanism*, not only about outcome: the same documented expectation failure should produce observably different defensive behavior depending on the dominant worldview of the institution exhibiting it — rule-stacking in Blue systems, narrative reframing without structural change in

Yellow systems, failure-attribution to outsiders in Green systems, and forceful pre-emption in Red systems.

The corresponding design principle, stated explicitly in the author's subsequent work, is that governance should be structured to generate "legible, survivable, revisable failures" — failures that are visible before they compound, survivable for the institution that produces them, and structured so that revision is the path of least resistance rather than the path of greatest political cost. This is the practical converse of PEFT's own Phase 3/Phase 4 diagnosis (Section 4): if democratic systems currently make failure illegible, unsurvivable for the actor who admits it, and structurally costlier to revise than to exploit, then the design target is not the elimination of failure but the engineering of its visibility, survivability, and revisability.

2.3 A Serious Internal Challenge: Does PEFT's Pursuit of Closure Itself Suppress Legitimate Conflict?

PEFT's own corpus does not shy away from a challenge that deserves to be stated plainly rather than resolved by omission. Chantal Mouffe's agonistic theory of democracy argues that the political is irreducibly conflictual: attempts to suppress antagonism in favor of rational consensus do not produce a healthier democracy but a democratic deficit that creates the conditions for exactly the right-wing populism such suppression was meant to prevent. On this view, a theory whose central design recommendation is "closure" — get citizens and institutions to a revised, shared, more accurate script — risks becoming a technocratic project that treats persistent disagreement as a cognitive defect to be corrected rather than as a legitimate and irreducible feature of plural societies.

This is not a challenge PEFT can dismiss by appeal to its own architecture, because the architecture itself (Section 2.2's quaternion-coded worldviews) already implies that closure looks different, and may not be uniformly desirable, depending on which worldview is doing the revising. A Green-dominant relational impasse and a Blue-dominant institutional rigidity are not the same kind of "failure to close," and treating both as defects to be corrected by the same instrument would itself be a category error of exactly the kind PEFT diagnoses in others (see Section 5.2's discussion of relational problems receiving only procedural answers). The honest position, taken in this edition, is that PEFT's closure mechanism should apply to *performance* and *procedural* scripts — where there typically is a checkable fact about whether a promise was kept — and should not be read as licensing the suppression of *relational* and *value* disagreements that have no fact capable of closing them. Section 7.1's discussion of QUO?'s formal treatment classes makes exactly this distinction operational: a value question is not treated as an unrevised script awaiting correction, but as a standing disagreement to be shown fairly and connected to others who hold it, not resolved by fiat.

3. The Empirical Landscape: De Vries's Documentation of a Systemic Failure

The theoretical claims of PEFT find their most comprehensive empirical validation in the research program of Catherine E. De Vries. Spanning two decades, her work has systematically mapped the landscape of political dissatisfaction, providing the evidentiary backbone for the theory.

The first phase of her research focused on the formation of expectations, demonstrating that electorates hold latent but powerful scripts — what might be termed dormant expectations — that

can be activated by political circumstances. Subsequent work on party competition revealed that political actors do not merely respond to these scripts but actively construct them, framing what citizens should anticipate from policy outcomes.

The second phase detailed the mechanics of expectation failure. De Vries's foundational work on Euroscepticism established that dissatisfaction with the EU is driven by perceived deviation from national benchmarks. This is a pure script failure: the EU's performance is measured against a script derived from domestic politics, and the mismatch generates withdrawal of support. The concept of the "political entrepreneur," developed with Sara Hobolt, introduced the crucial mechanism of exploitation. Successful challenger parties are those that identify existing expectation gaps and amplify them, transforming failure from an accidental byproduct into a deliberately cultivated electoral resource.

The third and most recent phase demonstrates that expectation failure has moved beyond episodic crisis to become a structurally embedded condition. Research on the "geographies of discontent" shows that these failures now cluster spatially and socially, creating persistent zones of institutional alienation. Her forthcoming work, *Symfonie van Onvrede*, conceptualizes the rise of the radical right as the cumulative political expression of decades of unresolved expectation failure — a system's pathological immune response to a chronic condition that was never treated.

3.1 What Was Missing: A Mature Theoretical Neighbor for the Empirical Pillar

De Vries's program is treated in the original essay as documentation of PEF, which is accurate but incomplete: De Vries does not work in isolation, and her closest theoretical neighbor within political science — **valence politics theory** — already has a name, a thirty-year measurement tradition, and a directly parallel construct for the very phenomenon PEFT is describing.

Clarke, Sanders, Stewart, and Whiteley's valence politics theory argues that, contrary to spatial models of left-right competition, voters in contemporary democracies increasingly choose between parties not on disagreement over policy goals — most of which command broad societal consensus — but on perceived *competence* in delivering those shared goals. Elections become contests over who can be trusted to deliver, not over what should be delivered. A 2025 extension of this framework, applied to the rise of the Reform party in the United Kingdom, introduces the construct of **state valence failure**: a perceived, persistent inability of the existing political class, across the mainstream spectrum, to deliver competently on consensus goals — control of immigration, functioning public services, economic stability — identified as the underlying driver of populist insurgency on both the political right and left.

State valence failure is, descriptively, very close to what PEFT calls expectation failure at the performance-script layer. The decisive advantage of bringing it into PEFT is methodological: valence politics already has thirty years of British Election Study panel data and validated survey batteries for competence perception. PEFT's central construct — the unresolved gap between expectation and reality — does not need to be measured from scratch at this layer; it can borrow an existing, tested instrument. What valence theory does *not* explain on its own is why state valence failure becomes self-reproducing across successive governments of different parties, rather than being resolved once a more competent government takes office. That is precisely PEFT's distinctive explanatory ambition, and Sections 5 and 6 below — strengthened by punctuated equilibrium theory and predictive processing — supply the missing account of why the underlying competence gap regenerates after every election, even as which party gets blamed changes.

This combination also generates a direct empirical test: if PEFT and state valence failure are both correct, perceived competence should decline across successive incumbents of different ideological orientation within the same political system, at a rate that does *not* track measurable changes in objective governance outcomes — because the failure lies in the expectation architecture, not in objective performance. This is testable using valence-perception time series that already exist in election studies across Britain, the Netherlands, and other European democracies, collected for unrelated purposes.

3.2 What Was Missing: A Cross-National Mechanism for the Relational Layer, and a Sharper Account of Entrepreneurial Specialization

Performance-script failure (Section 3.1) is not the whole of what drives contemporary populism, and the author's own subsequent work supplies a second, independently documented mechanism operating specifically at the relational layer (Section 2): nostalgic deprivation, a tripartite experience of status loss — economic, social, and political simultaneously — identified by political scientists including Ferwerda and colleagues as a stronger predictor of populist support across nineteen European countries than economic anxiety alone, for both left- and right-populist parties. The phenomenon is not Dutch, European, or even Western: it is visible in the United States, where the 2025 National Security Strategy explicitly invokes "an America that cherishes its past glories"; in Turkey's Ottoman revivalism under Erdoğan; in India's Hindu-nationalist mythology under Modi; in China's rhetoric of national rejuvenation after a "century of humiliation"; and in Russia's appeals to tsarist and Orthodox tradition. This gives PEFT's claim of global PEF diffusion (the original essay's discussion of De Vries's globalization findings) a specific, comparative, cross-regime mechanism rather than only an assertion of structural similarity: each of these political cultures supplies the same relational-script content — a recognizable, continuous "us" under threat from change — through entirely different historical material, which is exactly what a shared underlying mechanism with locally varying content should produce.

This also sharpens PEFT's account, in Section 6 below, of what political entrepreneurs actually exploit. The author's analysis of the Dutch SP's electoral peak in 2006 and subsequent decline identifies a specific division of entrepreneurial labor: the SP possessed what might be called a *language for injustice* — a vocabulary, rooted in direct presence in working-class neighborhoods, for naming economic grievance — but not a *language for invisibility*: a vocabulary for the distinct relational grievance of feeling actively unrecognized, of having one's account of reality not merely ignored but denied by those who govern. Right-populist entrepreneurs, from Fortuyn onward, filled precisely that second gap. This means PEFT's Phase 3 (failure exploitation, Section 4) is not a single undifferentiated mechanism but at least two distinct ones operating on two distinct script layers: performance-script exploitation, which competes on naming objective deprivation, and relational-script exploitation, which competes on validating subjective non-recognition — and the comparative evidence above suggests the second has been, in recent electoral cycles, the more potent and more portable across radically different national contexts.

4. Political Expectation Failure Theory: The Cycle of Non-Resolution

Synthesizing these cognitive and empirical insights, PEFT proposes a formal model centered on a four-phase cycle that characterizes modern democratic systems.

The cycle begins with **Phase 1: Expectation Construction**. Through electoral promises, policy narratives, and media framing, political actors set the scripts against which their performance will be judged. This establishes the benchmarks for future evaluation.

In **Phase 2: Performance Deviation**, institutional performance inevitably diverges from these constructed expectations. This deviation may be objective (a policy fails) or perceptual (a shifting reference frame redefines what success looks like). Given the complexity of governance, some degree of deviation is structurally inevitable.

Rather than triggering learning, this deviation leads to **Phase 3: Failure Exploitation**. Political entrepreneurs, typically from outside the mainstream, capture the expectation gap. They do not seek to resolve it but to amplify it, making the failure more salient and constructing new, often unachievable, expectations that incumbents are framed as structurally unable to meet.

This culminates in **Phase 4: Non-Resolution**. New actors gain power on a platform of resolving the failure, but they too are subject to the same structural constraints and political incentives. They either fail to meet the inflated expectations they helped construct or find themselves exploiting new failures. The cycle then returns to Phase 1 with a new cast of characters but the same underlying dynamics.

This cycle's self-perpetuating nature is explained by a critical divergence from Schank's cognitive model. In a healthy cognitive system, failure triggers a sequence: expectation → failure → explanation → reminding → generalization → revised script. In the political system described by PEFT, the sequence is truncated: expectation → failure → political mobilization → new expectation → repeated failure. The missing element is *closure* — the cognitive and institutional process by which a failed script is revised into a more accurate model of reality.

4.1 What Was Missing: A Falsifiable Boundary, and an Asymmetry

As originally stated, the claim that democracies cannot achieve closure is almost unbounded — and unbounded claims of this kind are vulnerable, because political science already has a well-tested model of the opposite phenomenon: domains where the public *does* learn from policy and revise its expectations, continuously and measurably.

Christopher Wlezien's **thermostatic model** proposes that public preferences for policy respond to policy itself as a negative feedback loop: when spending in a domain increases, preference for further increases falls, and vice versa. This mechanism has been confirmed across the United States, United Kingdom, and Canada — but only conditional on two factors: issue salience, and the centralization of institutional responsibility, which determines how easily citizens can attribute outcomes to government action. Where both conditions hold, citizens behave very much like Schank's healthy learner: they observe a deviation and revise the expectation.

This does not refute PEFT's account of Phases 2 through 4; it bounds it precisely. The thermostatic mechanism is known to break down exactly under low salience, diffuse institutional responsibility, and contested attribution — precisely the conditions under which Phase 3 (failure exploitation) and Phase 4 (non-resolution) should dominate. This converts PEFT's broadest claim into a sharp, testable hypothesis: non-learning should appear as an interaction effect, present where salience is high but attribution is diffuse or contested (EU governance, multi-level immigration policy, coalition government), and largely absent where salience and attribution are both high and singular (central bank interest-rate policy, domestic crime statistics). The performance-script layer (Section

2) is where this boundary should be sharpest; the relational layer, being least attributable to any single actor or decision, is where PEFT's non-learning claim should hold most strongly.

The thermostatic mechanism has a further boundary that is demographic rather than only topical, and the author's own municipal-level data supplies it directly. An analysis of voter turnout in Leiden found that roughly 40 percent of eligible voters do not participate in local elections, concentrated by income, education, and tenure status in specific lower-socioeconomic neighborhoods, while the governing coalition's effective electoral base is approximately 30 percent of all eligible voters — disproportionately drawn from higher-income, higher-turnout, often transient student populations in the wealthier parts of the city. Citizens who do not vote are not merely a missing data point in the thermostatic model; they are, by construction, entirely outside the feedback loop the thermostatic mechanism requires to operate at all, since that mechanism presupposes a population that registers policy and adjusts its expressed preference accordingly. This identifies a further, demographically specific population for whom PEFT's strongest claim — non-learning without closure — should hold not because of contested attribution or low salience, but because the feedback channel itself is structurally absent for them, a self-reinforcing condition the Dutch national Social and Cultural Planning Office's own research independently confirms: lower-status groups report lower institutional trust and consequently vote less, which further reduces their weight in coalition formation, which further reduces the likelihood that policy reflects their stated interests — a closed loop with no obvious entry point for the thermostatic correction that operates, imperfectly but measurably, for the population that does participate.

The cycle as originally described is also, in its language, largely *symmetric*: expectations are constructed and then deviated from, in either direction. Stuart Soroka's **negativity bias** research shows that this symmetry does not hold psychologically. Negative political and economic information is weighted roughly twice as heavily as positive information of equivalent magnitude, in citizen evaluation of leaders, in aggregate reactions to economic news, and even in physiological response to news content — a direct extension of Kahneman and Tversky's loss aversion. This supplies the missing asymmetry in the PEF cycle: even where the thermostatic mechanism does produce genuine policy responsiveness, the public's *felt* sense of expectation failure will lag behind, and exceed, what the objective responsiveness data would predict, because failures are remembered roughly twice as vividly as corrections. Combined with the predictive-processing addition in Section 2.1, this also explains why political priors tend to ratchet toward pessimism over successive cycles rather than oscillate around a stable equilibrium: each cycle's failure is weighted roughly twice as heavily as any intervening fulfillment, producing a one-directional accumulation of negatively-weighted political memory even when fulfilled and violated expectations occur with comparable objective frequency.

5. The Systemic Trap: Why Democracy Fails to Learn

Four structural factors combine to prevent democratic systems from achieving closure. First, there is a fundamental **incentive misalignment**: political actors gain electorally from exploiting failures, not from resolving them. Resolution reduces political differentiation; failure amplification increases it. Second, a **temporal mismatch** exists between short electoral cycles and the long time scales required for complex policies to yield results. Third, a profound **cognitive asymmetry** persists: citizens hold implicit scripts but lack instruments to make them explicit, test them against reality, or revise them, while political actors are highly skilled at exploiting this opacity. Fourth, **institutional**

rigidity — constitutional frameworks and bureaucratic structures designed for stability — actively resists the rapid adaptation that learning from failure demands.

The original essay reached for complexity science, and specifically panarchy theory, to explain how this trap manifests structurally. Healthy systems undergo adaptive cycles of growth, conservation, release (crisis), and reorganization. Democratic systems experiencing chronic PEF were described as trapped in a pathological variant: the release phase (electoral disruption, populist breakthrough) occurs regularly, but the reorganization phase fails. New actors take power, but the underlying expectation architecture remains unchanged.

5.1 What Was Missing: A Tested Mechanism in Place of a Borrowed Metaphor

Panarchy theory is evocative, but it was built from ecological and social-ecological systems, not from political data, and its application to democratic institutions has so far been by analogy rather than direct empirical test. The policy sciences possess their own, harder-edged theory of exactly this phenomenon, built from forty years of cross-national budget and agenda data: **punctuated equilibrium theory** (PET).

Baumgartner and Jones's PET, tested against hundreds of thousands of observations in the Policy Agendas Project, holds that political institutions are boundedly rational information processors that attend disproportionately to a small number of salient issues at any given time, allowing problems in unattended domains to accumulate unaddressed — the "error accumulation model." This produces a pattern of long stasis interrupted by large, infrequent corrections, rather than smooth incremental adjustment. PET has already been extended cross-nationally: more centralized, less pluralistic systems display more extreme punctuation patterns than systems with more numerous, competing venues for issue attention. This gives the vague claim that "democratic institutions are designed for stability, not learning" a specific, falsifiable institutional variable — punctuation severity should correlate inversely with the number of independent venues processing a given issue — and it sharpens the explanation of why populist parties, once in power, often reproduce the very failures they campaigned against. The reproduction of failure is not merely a matter of incentives persisting after the election (as Section 6 below describes); it is a structural property of bounded institutional attention itself. An incoming government of any ideology inherits the same scarce attention capacity and the same backlog of unattended issues; resolving the one issue that won the election does not redistribute attention to the many other domains where expectation failure has been silently accumulating. This would hold even under hypothetically perfect incentive alignment — which is precisely why the failures recur regardless of which party or movement takes power.

PET and the thermostatic model (Section 4.1) also fit together directly: punctuation events are exactly the moments at which the thermostatic mechanism, ordinarily too slow or too diffuse to register, becomes visible in a sudden, large correction — the political equivalent of an accumulated error being processed all at once rather than continuously.

Both mechanisms remain, in PET's own literature and in the original 2026 PEFT essay alike, primarily qualitative — a pattern to be read off after the fact in policy-agenda data, not a quantity a system can monitor in real time. The author's subsequent platform work supplies two candidate real-time instruments for exactly this gap. A self-organized-criticality monitor tracks the scaling exponent of the distribution of policy or institutional changes; values in the range of roughly 1.5 to 3.0 indicate the system is operating near the critical boundary between excessive rigidity and excessive volatility, where information processing is most efficient — in PET's vocabulary, neither

so deep in stasis that error silently accumulates without limit, nor so disordered that no stable script can form long enough to be tested against reality. A second, complementary instrument borrows the Kuramoto model of coupled oscillators from physics to define a coherence parameter r(t), where r approaching 1 indicates that the relevant population's expectations and the institution's signals are closely synchronized and r approaching 0 indicates near-total incoherence between what is expected and what is communicated. Neither metric has yet been validated against an external criterion of actual political learning, and that validation is precisely the kind of claim this edition's Section 7.1 commitment to epistemic honesty requires flagging as untested; what they contribute, even unvalidated, is a concrete answer to a question PET and panarchy theory both leave open — what, operationally, a monitoring system would actually compute, on an ongoing basis, to detect an institution drifting toward the kind of silent error accumulation Sections 5.1 and 5.2 describe, before the next punctuation event makes it visible in retrospect.

These instruments have an earlier, much cruder precursor in the same author's 2006 work, which is worth noting because it shows the underlying intuition predates any of the formal apparatus now available to test it. That paper combines Edward Dewey's Foundation for the Study of Cycles (1942) — an empirical catalogue of economic cycles of varying periodicity, including the roughly fifty-year Kondratiev wave and the roughly ten-year Juglar investment cycle — with the observation that harmonically related cycles of different periodicities periodically come into conjunction, amplifying or dampening each other's effects, much as the SOC and Kuramoto instruments now formalize the difference between a system whose multiple time-scales are constructively aligned and one whose time-scales are working against each other. The same paper applies an early five-phase cyclical model, derived from the Chinese Sheng/Ko/Wu correspondence system, to two centuries of Western institutional history and arrives at a recurring two-phase pattern it labels Discovery followed by Exploitation — observing, in 2006, that "just as in the Renaissance, the Age of Discovery is followed by an Age of Exploitation," and identifying the same two-phase sequence in the Hellenistic and Roman periods. This is the same pattern Section 1.1 above redescribes, twenty years later, as the difference between expectation construction in service of coherence and expectation construction in service of extraction — which is one further indication that PEFT's diagnosis is not a recent theoretical invention applied retroactively to old material, but the most recent formal statement of an argument this corpus has been refining continuously since at least 2006.

The author's corpus also supplies a documented, contemporary illustration of late-K rigidity at a scale PEFT's original essay did not examine empirically: the operational interior of Dutch public administration rather than only its electoral surface. A 2025 cross-domain analysis using Dave Snowden's Cynefin framework alongside the panarchy cycle catalogues structural waiting times across nearly every Dutch public domain simultaneously — mental health care (over 108,000 people waiting as of October 2024, a figure that had risen each year since 2022), youth protection (well over a thousand children without an assigned caseworker), asylum procedures (an average wait of roughly 53 weeks merely to begin the process), driving exams, childcare, social housing (a measured average wait of 9.8 years in Amsterdam), and electricity grid connections (waits of up to a decade for business connections) — and finds the same diagnostic pattern across all of them: not a knowledge deficit, since the relevant advisory bodies (the Scientific Council for Government Policy, the Court of Audit, the Council of State) had already identified the structural causes, but an incentive-and-translation failure, in which institutions and the politicians who oversee them are never held accountable for adaptive capacity or throughput, only for procedural compliance and budget adherence. Three specific causal mechanisms identified in this analysis sharpen PEFT's account of why institutional rigidity (Section 5 above) persists even after correct diagnosis: *blame*

avoidance, in which officials deliberately delay or defer decisions specifically to limit personal or political exposure to criticism; *administrative burden as deliberate friction*, in which procedural barriers are knowingly designed to reduce de facto access to a right or benefit without ever formally restricting that right; and *policy drift*, in which rules are simply never updated as circumstances change, producing de facto stagnation without any formal decision ever being taken to stagnate. None of these require bad faith at the level of any individual actor; each is, like Section 2.2's worldview-specific defensive mechanisms, a structurally available response that becomes individually rational once the underlying incentive structure (Section 4 above) rewards procedural defensibility over adaptive throughput.

The same analysis, however, also supplies PEFT's strongest available counter-example to date: a domain in which Dutch governance has demonstrably built and sustained the kind of adaptive capacity PEFT's closure mechanism requires. The Dutch Delta Programme for water safety — built on adaptive policy pathways, explicit tipping points, scenario planning, and pre-defined decision windows at which a chosen path can be reopened and revised — operates as a working model of what Section 5.3's quaternion state proposes constitutionally: a domain where revision is built into the institutional design from the outset rather than resisted as an admission of failure. This matters for PEFT because it demonstrates that the non-learning pattern documented throughout Sections 5.1 and 5.2 is not a fixed property of Dutch governance as such, but a property of specific institutional designs that happen to dominate most domains; water safety governance shows that the alternative is not merely theoretical, since a comparable adaptive architecture has been operating, and has been credited by the same advisory bodies that diagnose rigidity elsewhere, in this one domain for years.

5.2 Two Twenty-Year Natural Experiments in Political Non-Learning

The clearest empirical illustration available for the single-loop/double-loop distinction introduced in Section 2.1 is a case the author has documented in detail: the Dutch VVD's 2005 *Liberaal Manifest*. The manifesto, drafted by a committee including historian Frank Ankersmit, diagnosed a relational script failure with striking clarity — citizen trust in Dutch governance had collapsed from 65 to 30 percent between 2000 and 2003, and power had leaked from the visible, accountable center into what the manifesto itself called "an unsurveyable labyrinth" of semi-autonomous administrative bodies. The diagnosis named the problem precisely as a *relational* one: citizens no longer recognized themselves in the institutions that governed them.

The party that wrote this diagnosis then held the office of prime minister continuously for fourteen years — long enough, on any reasonable accounting, to implement its own prescription. It did not. The directly elected prime minister, the directly elected mayor, the substantial local tax base, the flat tax, the reduction of the proliferation of semi-autonomous bodies: none of these were implemented. The one proposal that *was* implemented — a Ministry of Security, concentrating rather than devolving power — is the diagnostic detail. Of the manifesto's two directions of reform, the one that would have moved power toward the center survived contact with governing practice; every proposal that would have moved power toward the citizen did not. Trust continued to fall.

In Argyris and Schön's terms, this is single-loop learning operating exactly as the theory predicts under stress: the manifesto's governing assumption — that legitimate process guarantees legitimate outcomes, and that better-designed oversight bodies fix oversight failures — was never revised. It was intensified. Each documented failure of an oversight mechanism produced a proposal for an additional oversight mechanism, rather than a revision of the belief that oversight mechanisms were the right tool at all. This is precisely the failure mode that Section 2.2's Fractal Karma model

predicts for a Blue-dominant institutional culture: not corruption, not bad faith, but a structural incapacity to produce any response to documented failure other than intensification of the existing script — because intensification is the only move a Blue-dominant system's defensive architecture makes available to it.

A second, complementary case operates not through failure to revise a diagnosed script but through active prevention of the feedback that would generate the disconfirming evidence in the first place. The author's analysis of Dutch municipal real-estate governance — using the Leiden Bio Science Park and several housing developments as documented cases — identifies three self-reinforcing mechanisms that keep a script from ever encountering a registered expectation failure: research commissioned by the party with a stake in its conclusion (economic-impact studies of a development paid for by the development's own stakeholders, producing job-count figures that vary by a factor of two to three depending on which commissioned study is cited); participation processes that solicit comment on cosmetic details while the substantive decisions are settled beforehand in informal contact between officials and developers; and the deliberate non-integration of public records that, combined, would make the gap between official communication and documented decision visible — land registry, building permits, council minutes, and economic statistics are each individually public but never combined into a single queryable record, which is not a capacity gap but, on the available evidence, a political choice, since the Dutch open-government law carries no enforcement mechanism and municipalities self-report their own compliance through their own membership association. Where Section 5.2's first case shows a script that was diagnosed but never revised, this second case shows a script that was protected from ever generating the kind of disconfirming signal Schank's model (Section 2) requires for revision to even become a live option. The revolving door between Dutch governing parties and the real-estate sector — documented instances include former ministers and aldermen moving directly into developer-funded lobbying and advisory roles immediately after leaving office, with a measured average rate of approximately 42 percent of departing politicians moving into lobbying generally, concentrated heavily among parties with the closest sectoral ties — supplies the incentive structure (Section 4) that makes the suppression of disconfirming evidence individually rational for the actors involved, even where no single decision is independently unlawful.

5.3 A Constructive Response: The Quaternion State and Its Variety Threshold

If the diagnosis is structural rather than personal — if no individual party or leader is positioned to revise a script using only the instruments that the script itself supplies, and if the relevant actors can additionally prevent the disconfirming evidence from being generated at all — then institutional reform proposals and even citizen-level reflective instruments (Section 7) address only part of the problem. The deeper design question is what kind of state could make double-loop learning structurally available in the first place, rather than relying on a single dominant worldview to revise itself from the inside, or relying on the goodwill of actors who currently benefit from non-revision.

The author's most formally developed answer draws on Ashby's cybernetic law of requisite variety — only variety can absorb variety; a system must possess sufficient internal complexity to respond to the complexity of its environment — combined with a specific numerical threshold derived from the same nilpotent quaternion algebra underlying Section 2.2's Fractal Karma model. The derivation, developed in full elsewhere in the corpus, is not a historical analogy but an algebraic consequence: the three imaginary units of quaternion algebra (i, j, k) introduce a three-fold branching at every recursive step, and the natural fixed point of any three-fold recursive branching is the Bronze Mean, $\beta = (3 + \sqrt{13})/2 \approx 3.303$ — the positive root of $x^2 - 3x - 1 = 0$, and the slowest-

converging purely periodic continued fraction of period three, which makes it the maximally stable ratio for a system with this recursive structure. This generates a specific integer sequence (1, 1, 4, 13, 43, 142…) in which each stable coherence domain must cross a threshold before becoming a genuine attractor rather than a transient fluctuation, dividing the sequence into four named coherence phases: Phase I (threshold 1), minimal bound coherence; Phase II (threshold 4), integration at the scale of stable composite structure; **Phase III (threshold 13), autopoietic capacity — the threshold at which a system becomes capable of maintaining itself**; and **Phase IV (threshold 43), recursive self-reference — the threshold at which a system becomes capable of modeling itself**.

This naming is exact rather than evocative, and it is what gives PEFT's central concept of closure a structural address rather than only a description. PEFT's closure (Sections 4 and 5 above) requires that an institution register a documented expectation failure and revise the model of reality that produced the expectation — which is, in this vocabulary, nothing other than Phase IV self-modeling capacity applied to the institution's own governing premises. Phase III — mere self-maintenance, a system that reproduces itself successfully — is achievable, on this account, well below the threshold closure requires, which is the formal reason a political system can be perfectly stable, electorally successful, and self-reproducing (Phase III) while remaining structurally incapable of revising its own foundational script (Phase IV) — exactly the pattern Sections 5.1 and 5.2 document empirically without, until this addition, a principled account of why stability and revisability are not the same achievement and do not arrive together.

Applied historically, the author argues that the Dutch institutional settlement following the 1619 Synod of Dordrecht — which resolved a theological dispute over free will and predestination in favor of the position that closed off individual conscience as a source of legitimate challenge to established doctrine — fixed Dutch governance culture at Phase III without crossing into Phase IV, and that subsequent institutional forms (the VOC's extraction model, nineteenth- and twentieth-century pillarization, the postwar polder model) are successive surface expressions of the same underlying sub-threshold variety, which is why each new institutional generation reproduces the non-learning pattern in different material rather than resolving it. This is a strong historical claim, advanced elsewhere in the same corpus and not re-derived here in full.

The model's own stated standard for what would count as evidence against it is worth stating plainly, since it is unusually explicit for a claim this ambitious and this edition's commitment to epistemic honesty (Section 7.1) requires applying the same standard here: the broader framework predicts that any adequate empirical proxy for coherence capacity should cluster at Bronze Mean thresholds rather than vary continuously, and the author treats uniformly distributed transitions, rather than threshold-clustered ones, as a result that would falsify the selection rule itself rather than only the specific labels attached to its phases. No equivalent empirical test has yet been run against political-institutional data of the kind Sections 5.1 and 5.2 discuss; the threshold's application to Dutch governance history is, at present, an interpretation fitted to known historical outcomes rather than a prediction confirmed against new data, and this edition treats it accordingly — as the most precisely specified available hypothesis about why closure capacity has a formal lower bound, not as a demonstrated fact about Dutch institutions specifically.

The corresponding constitutional design — termed the **quaternion state**, a considerably more formally specified successor to earlier "fractal state" sketches in the same corpus — proposes four constitutionally independent organs, one per quaternion component, none of which can be appointed or suspended by the others:

- An **(s)-organ for constitutional continuity**: the foundational legal framework, treated not as a static text but as the locus of stable identity, designed so that its own disconfirmation — were it to be falsified by experience — would itself be a legitimate trigger for revision rather than an act the framework's defenders are structurally required to resist.
- An **(i)-organ for citizen falsification**: a mechanism giving individual citizens' direct, disconfirming contact with reality — lived experience, direct observation, whistleblowing — actual institutional weight against policy premises, replacing the existing model of consultation (citizens speak, the institution filters) with one in which a defined threshold of disconfirmation triggers mandatory premise review.
- A **(j)-organ for constitutive civil society**: formal recognition of commons, cooperatives, neighborhood associations, and civic organizations not as consulted stakeholders but as constituting powers in governance, with institutional standing the executive cannot unilaterally withdraw.
- A **(k)-organ, a coherence council**, with constitutional authority to determine when a policy's foundational premises are contradicted by evidence generated through the (i)- and (j)-channels, and to mandate foundational revision when that threshold is crossed — an authority that, by construction, cannot be appointed by the executive branch whose premises it may be required to overturn, since an executive-appointed body for revising executive foundations is a structural contradiction.

No existing democratic system contains a (k)-organ of this kind; the closest functional analogues — constitutional courts, ombudsmen, auditors-general — have the power to rule on procedure but not the constitutional standing to mandate that a government's operating premises be revised. Designing such a body, with genuine independence from executive appointment and genuine power to compel rather than merely recommend, is arguably the single institutional innovation current reform proposals (including the VVD's own 2005 manifesto, Section 5.2) do not contain and would need to, if PEFT's diagnosis of structurally unavailable closure is correct. Read against Section 1.1's longer timeline, the absence is not surprising: a (k)-organ-equivalent is precisely the kind of standing, non-appointed check on power concentration that the corpus's account of the post-3117-BCE transition argues was specifically eliminated, repeatedly, at the moments when coherence-based social organization gave way to control-based organization — which means the (j)-organ's *aparigraha* constraint (Section 5.3) and the (k)-organ's revision-mandating power are not two separate institutional proposals but two faces of the single architectural feature the Indus Valley Civilization apparently sustained for two millennia and that subsequent civilizations, beginning with the post-flood transition the Sumerian King List records, did not.

The quaternion state does not prescribe outcomes; it is a proposal about what conditions would make double-loop learning structurally possible across all four worldview orientations simultaneously, rather than requiring any one of them to revise itself unaided, or hoping that a sufficiently well-intentioned single-worldview government will eventually do so. Section 7.1 below describes a working instrument — QUO?, built on the SWARP platform — that operationalizes a function close to the (i)-organ directly: a standing, queryable record of the gap between official communication and documentary fact, available to any citizen without requiring them to first identify which agency, which procedure, or which form applies to their question.

5.4 Why the (k)-Organ Specifically Is Missing, Historically and Right Now

Two further pieces of the author's corpus sharpen the case for the (k)-organ — the coherence council with constitutional standing to mandate premise revision — beyond the general argument given above, by showing both that its absence has a specific historical mechanism in Dutch

governance and that its absence is, on a serious contemporary diagnosis, an active rather than merely historical danger.

Historically, an analysis drawing on historian D.G.A. Alkemade's research into the Dutch democratization process between 1786 and 1848 identifies the specific mechanism by which a functional equivalent of a (k)-organ has been repeatedly prevented from forming. Willem I's post-1813 settlement, on this account, did not suppress the radical-democratic gains of the Batavian Revolution through direct repression; it absorbed them through what Alkemade terms "constitutional domestication" — integrating former radical actors and institutions into a monarchical framework that retained democratic forms (representation, civil rights, a constitution) while keeping substantive power concentrated, neutralizing potential opposition by co-opting it into the new system rather than confronting it. The case of Pieter Linthorst, who moved from revolutionary radical to a figure absorbed into the institutional system, is offered as the individual-level instance of a pattern Alkemade documents at the level of constitutional design itself. This matters for the (k)-organ proposal specifically because it identifies co-optation, not only direct suppression, as the mechanism that has historically prevented an independent revision-mandating body from taking root in Dutch governance: any such body's members and powers have, on this account, historically been the first targets of absorption into the existing settlement, which means a (k)-organ's design would need to specify not only formal independence from executive appointment but active structural resistance to exactly this kind of co-optive integration — a harder design problem than independence alone.

The contemporary version of the same danger is given a sharper name in the author's reading of Ilja Leonard Pfeijffer's essay collection *Absolute Democratie*: democracy is not necessarily undermined from outside, by external enemies, but can be dismantled from within, in democracy's own name, when majoritarian or winning factions treat electoral legitimacy as license to erode the institutional checks, rule-of-law guarantees, and minority protections that distinguish constitutional democracy from unconstrained majority rule. Pfeijffer's diagnosis — that this produces a system that remains formally democratic while losing genuine counter-power (*tegenmacht*) over executive authority — is, in this edition's vocabulary, the end state of PEFT's Phase 4 (non-resolution, Section 4) when no (k)-organ exists to interrupt it: each successive election can, in principle, license a further round of dismantling the very mechanisms that would otherwise force premise revision, with nothing in the system positioned to say the process has crossed a line, because the body that would say so is exactly the kind of independent, executive-resistant institution absolute democracy's logic erodes first. This is also why this edition treats the (k)-organ's executive-independence requirement (Section 5.3) as necessary but insufficient on its own: Alkemade's history shows that independence on paper has not historically prevented co-optation in practice, and Pfeijffer's diagnosis shows what is structurally at stake if a functional equivalent never forms.

A further elaboration of the wijkcirkel ("neighborhood circle") implementation referenced in Section 5.3, developed independently of the quaternion-state formalization, adds an empirical dimension this edition's account did not previously have: an explicit second axis, orthogonal to Fiske's four relational modes, distinguishing *Judging* orientations (top-down, planning, evaluative, closure-seeking) from *Perceiving* orientations (bottom-up, observational, experimental, adaptive — borrowing this distinction from Myers-Briggs typology). Mapped against the 2025 Dutch party system, this analysis finds that the two dominant governing blocs — a Red-Blue-Judging "patriarchal" bloc (market plus hierarchy, top-down) and a Green-Blue-Judging bloc sometimes mistaken for its opposite (community plus bureaucracy, equally top-down) — share the same Judging orientation despite occupying opposite ends of the conventional left-right spectrum, while a

genuinely bottom-up, experimentalist Green-Yellow-Perceiving combination commands only a minority of seats. This is a directly falsifiable refinement of Section 2.2's claim that worldview composition predicts failure mode: it predicts that apparently opposed political blocs sharing the Judging orientation should display the *same* characteristic non-learning pattern — top-down construction of expectations followed by top-down attempts at correction, never bottom-up experimentation — regardless of where they sit on the conventional ideological axis, which is a stronger and more specific claim than "the left and the right behave differently." The proposed institutional response — nested neighborhood circles using consent-based (not consensus-based) decision-making, escalating to higher scales only when a matter cannot be resolved locally, with results tracked on public dashboards — is offered explicitly as a partial, sub-national implementation of the same closure principle the (i)- and (j)-organs formalize at constitutional scale, and is subjected in the same analysis to its own falsifiability check: a review against seven independent theoretical traditions (Ostrom's polycentric governance, deliberative democracy, social identity theory, complexity theory, public choice theory, the capabilities approach, and historical institutionalism), each of which identifies a specific way the proposal could fail — elite capture of nominally local processes, capacity asymmetries between circles, path-dependent resistance from existing party and consultancy interests — rather than presenting the proposal as a panacea.

6. The Systemic Role of Political Entrepreneurs

Within this framework, political entrepreneurs are not merely opportunistic actors; they perform a systemic function. They serve as the mechanism by which accumulated expectation failure is converted into electoral energy. Their role is structurally analogous to what Schank calls *explanation* in cognitive processing — the moment a script failure is noticed and named. However, where Schank's explanation leads toward learning and script revision, the political entrepreneur's explanation leads toward amplification and mobilization. They name the failure, assign blame, and construct a new, equally unachievable expectation, ensuring the cycle's continuity.

This dynamic is reinforced by a fundamental asymmetry. Political actors are highly skilled at *constructing* expectations — the core technology of electoral competition — but are structurally prevented from *revising* them, as revision requires acknowledging failure, which is politically costly in adversarial systems. The more expectations are constructed without revision, the larger the accumulated gap, and the more fertile the ground for entrepreneurial exploitation.

Section 2.1's predictive-processing addition gives this account a sharper, testable form: successful entrepreneurial rhetoric should be associated with language that *raises* certainty on both sides of the gap — about what was promised and about what occurred — rather than language that increases epistemic humility about either. And Section 4.1's negativity-bias addition gives entrepreneurs a specific, well-documented rhetorical lever: messaging that frames outcomes in loss terms (what citizens have lost or been denied) should be more electorally potent than equivalent messaging in gain terms (what could be achieved), independent of the objective accuracy of either framing — simply because losses are weighted roughly twice as heavily as gains in political memory.

6.1 What Was Missing: The Industrialization of Phase 1, and a Quantified Paradox

PEFT's Phase 1 (expectation construction, Section 4) describes political actors setting the scripts against which performance will later be judged. The author's own subsequent analysis of political

advertising supplies the missing account of how thoroughly this phase has been industrialized by techniques developed for commercial marketing, and a quantified empirical signature of what successful low-coherence exploitation actually looks like.

The Cambridge Analytica case is the clearest documented instance: psychographic profiling derived from up to 87 million Facebook users, combined with Big Five personality scoring, allowed expectation-construction messaging to be tailored not merely to demographic segments but to individual psychological vulnerability — high-neuroticism voters receiving fear-based messaging on immigration and economic insecurity, low-openness voters receiving nostalgic appeals to lost national identity, in the same campaign, advancing the same political goal through deliberately incompatible emotional frames. This is Phase 1 construction operating with a level of audience-specific calibration Schank's original script theory (Section 2) did not anticipate: the same institutional script can be presented to different citizens as different scripts entirely, each calibrated to that citizen's specific susceptibility to expectation construction. Shoshana Zuboff's account of surveillance capitalism supplies the structural explanation for why this infrastructure migrated so easily from commercial to political use: the same data, the same prediction models, and the same behavioral-influence techniques developed to shape consumer purchasing decisions required no fundamental modification to shape political ones, since both are, in this framework, instances of the same underlying problem — predicting and influencing an agent's expectations.

A separate, quantitative analysis of the Dutch party system supplies a measurable signature of this dynamic in action. Scoring parties along four dimensions — organizational structure, practical policy effect, normative value-consistency, and overall convergence between vision and implementation — the author's analysis finds that the PVV combines the highest measured electoral effect (0.95) with the lowest measured value-consistency (0.20) of any major Dutch party, while a niche, single-issue party (the PvdD) shows the opposite pattern: comparatively low effect (0.60) but the highest value-consistency (0.90). This pattern, read through PEFT, is not coincidental: Ernesto Laclau's account of populist logic — that political effectiveness frequently derives from constructing antagonistic relationships rather than elaborating detailed, internally consistent policy — predicts precisely this trade-off. A political entrepreneur who must remain consistent with a comprehensive policy program faces a constraint a pure failure-exploiter does not: consistency limits how freely the expectation gap can be amplified, since some claims would contradict others already made. The measured inverse relationship between electoral effectiveness and value-consistency is, in PEFT's vocabulary, the quantitative fingerprint of unconstrained Phase 3 exploitation — and it generates a direct prediction for future research: political actors whose electoral effect substantially exceeds their measured value-consistency should show, on later analysis, disproportionate reliance on the certainty-raising and loss-framing rhetorical mechanisms identified above, while actors with high value-consistency but lower effect should not.

7. Implications: Reframing Democracy and the Path Forward

PEFT carries significant implications for democratic theory. First, it suggests that the standard focus on *preferences* as the fundamental unit of analysis is misplaced. Preferences are downstream of expectations. A theory operating only at the level of preference aggregation will always fail to explain how those preferences are constructed or why they systematically diverge from institutional capacity. Second, it challenges the Schumpeterian model of democracy as competitive elite

selection. If elites function primarily as failure exploiters, elite competition does not drive democratic improvement; it deepens expectation failure as a structural condition.

The theory thus implies a fundamental shift in the unit of democratic analysis: from the institution to the *cognitive-collective*. A democracy is not merely an institutional arrangement but a shared cognitive architecture — a distributed network of expectations, reference frames, and scripts — instantiated in institutions but not reducible to them.

If the core problem is the absence of closure — the inability to move from failure through explanation to script revision — then the design challenge for democratic theory becomes clear: the development of **citizen-level reflective instruments**. Existing civic technologies (voting aids, deliberation platforms) address preference expression and aggregation but do not address the prior cognitive question: what does the citizen actually expect, where did it come from, is it realistic, and how should it be revised in light of experience?

7.1 QUO?: A Working Instrument, Formally Grounded, Not Only a Proposal

The original essay described citizen-level reflective instruments as a frontier yet to be built. Since then, a first working version has moved from proposal to platform: **QUO?**, built on SWARP, addresses one specific failure mode the original essay did not anticipate in its design implications — the failure mode of existing municipal information systems themselves.

The diagnostic starting point is empirical: recent research into municipal chatbots found that only roughly one in ten citizen questions is answered correctly, and nearly two-thirds are answered wrongly — typically with high apparent confidence, based on a single recognized keyword. This is not a technical defect but a category error: such systems have exactly one procedure for any question — look up an answer in the institution's own documents — when the actual space of citizen questions is structurally richer. Some questions concern the future and are not answerable by anyone today. Some embed an assumption that must first be tested against the record. Some belong to a different authority entirely. And some are not information requests at all but an expression of distress. A system with only one tool for all of these cases must guess, and guessing with apparent confidence is precisely what further erodes trust in government — a direct, small-scale instance of PEFT's Phase 2/Phase 3 dynamic (Section 4) playing out at the level of a single citizen interaction.

This diagnostic claim is not only descriptive; the author has subsequently given it a formal proof. A *Totality Theorem*, developed by combining classical question logic (the observation, going back to Boole, that any well-formed question presupposes a bounded universe of discourse) with Genrich Altshuller's TRIZ and Roni Horowitz's closed-world reduction of it (ASIT) — a method distilled from analysis of roughly 200,000 patents to two conditions: an inventive solution adds no new kind of object to the system, and it uses the cause of the problem as the solution — proves that while a system that can correctly *answer* every possible question is impossible in principle (the space of possible questions is not closeable, since a new question can always be asked about any existing catalogue of questions), a system that can correctly *treat* every possible question is achievable. Four simple tests — is this an information question; do its presuppositions hold; is it decidable from fact; from which source — sort all possible questions into exactly eight treatment classes, four of which are principled non-answers: refer to the responsible party, correct a false presupposition, hold the question open until the world decides, or display competing positions fairly without resolving them by fiat. Applied to QUO?, this is the closed-world solution in the ASIT sense: nothing is added to the system (no larger model, no extra data source, no escalation desk) that was not already present

in the question itself, its presuppositions, the institution's own records, the landscape of other authorities, and time; and the cause of the original failure — unanswerability — becomes, class by class, the mechanism of the honest response, exactly as the frozen antenna in Horowitz's classic example uses the very ice that threatened to break it as the structural reinforcement that holds it up.

QUO? is designed to do what the word originally meant before it narrowed, in literate bureaucratic societies, to a grammatical form requiring an answer: to investigate. It checks council records, permits, and decision logs directly, and — critically — where official communication and the documentary record diverge, it shows both: *the municipality says X; the records show Y.* This is a direct, working instance of Section 2.1's closure principle and Section 5.3's (i)-organ design criterion (citizen falsification with institutional weight) operating at the scale of an individual citizen's question, rather than only at the scale of aggregate survey data. For genuinely unanswerable questions about the future, the system does not guess; it holds the question open and notifies the citizen only when the underlying record actually changes — a direct, working operationalization of the predictive-processing principle (Section 2.1) that closure requires updating only on a genuine change in evidence, not on the mere passage of time. And for value questions that no fact can resolve, the system follows Section 2.3's resolution of the Mouffe challenge directly: it displays competing positions and connects citizens who hold the same open question to one another, rather than treating disagreement as an unrevised script awaiting correction.

This matters for PEFT's broader empirical pillar because the relevant comparison is not QUO? against an ideal, but QUO? against the existing landscape of civic technology, which the author's own work surveys directly. Opinion-clustering platforms such as Polis summarize a population's vote vectors but hold no model of the individual citizen's underlying script, so participation revises the population-level topology but nothing about the citizen's own expectations. Narrative-sensemaking platforms such as SenseMaker collect richer, structured stories but end the loop at the point an analyst reads the resulting landscape, with no per-citizen update mechanism. Deliberation platforms such as Decidim and Loomio treat the citizen as a role-bearer moving through a process rather than as an agent whose profile updates on evidence. Data-sovereignty platforms such as Solid relocate where a citizen's data lives without offering any account of what that data means or how beliefs built on it should change. And voting-advice applications produce a static ideological snapshot with no mechanism by which subsequent political behavior feeds back into the model. Each of these, in Schank's terms (Section 2), is a script-application engine without script revision; none of them weaves expectation, evidence, and revision into a single operating loop. QUO?'s contribution, on the available description, is specifically that missing loop — though it is essential to state plainly, in the spirit of Section 2.3's commitment to epistemic honesty over overclaiming, that an architecture being *instrumented* to test a reduction in expectation failure is not the same claim as a *demonstrated* reduction in any measured population, and the author's own account of the platform's first deployment period treats this distinction as the central open question for the months ahead rather than as a settled result.

7.2 What the Five External Additions Demand of Such an Instrument

The integrations above are not only theoretical; they constrain how any citizen-level reflective instrument should actually be built.

From **predictive processing** (Section 2.1): the instrument must address *precision*, not merely content. Simply presenting citizens with disconfirming information will not produce revision if the

underlying prior is held with very high confidence; an effective instrument likely needs to lower the perceived certainty of the prior before disconfirming evidence can be integrated at all.

From **negativity bias** (Section 4.1): the instrument needs explicit counterweighting. Users will register and recall instances of expectation failure roughly twice as readily as instances of fulfillment, so a naively designed reflective tool will tend to reinforce, rather than correct, an already pessimistic prior — exactly the opposite of its intended purpose.

From **valence politics** (Section 3.1): the instrument does not need to invent a new performance-layer measure from scratch. Thirty years of validated competence-perception survey batteries already exist and can be adapted directly.

From the **thermostatic model** and **punctuated equilibrium theory** (Sections 4.1 and 5.1): the instrument's diagnostic layer can predict, in advance, *where* a given citizen's expectations are likely to be revisable through ordinary feedback (high-salience, clearly-attributable domains) and where they are likely to require the instrument's active support (low-salience, diffusely-attributable, or relational domains) — allowing the tool to focus its scarce engagement effort where unaided learning is least likely to occur on its own.

Developing tools that help individuals and communities make their political scripts explicit, compare them with reality, and engage in supported revision represents the next frontier. Such instruments would operate at the level of individual cognition while producing collective effects. As citizens develop more reflexive and accurate expectation structures, the political market for unachievable expectation construction would shrink, and the structural incentive for failure exploitation would begin to diminish.

8. Conclusion

The convergent crisis of democracy is, at its cognitive core, a crisis of expectation without reflection. Political Expectation Failure Theory provides a unified framework for understanding how modern democratic systems are structurally predisposed to generate, violate, and exploit citizen expectations without ever achieving the closure necessary for learning and adaptation. By integrating cognitive science, empirical political research, and complexity theory, PEFT reveals that the problem is not a series of discrete failures but a self-reproducing structural condition: institutionalized expectation failure without closure.

This edition has not changed that diagnosis. It has strengthened it along two independent tracks that turned out, on close reading, to converge. Five external research traditions give each of PEFT's three pillars a load-bearing connection to a body of work that already has data, measurement instruments, and a track record of surviving empirical testing: Wlezien's thermostatic model bounds the non-learning claim precisely, with a further demographic boundary supplied by the author's own Leiden turnout data; predictive processing gives script revision — and its failure — a formal mechanism; punctuated equilibrium theory replaces an ecological metaphor with a tested model of institutional attention, now paired with candidate real-time monitoring instruments (self-organized criticality, Kuramoto coherence) that the author's platform work has built but not yet validated; valence politics theory and state valence failure connect PEFT directly into thirty years of European electoral behavior data, complemented by cross-national evidence on nostalgic deprivation as the

specifically relational-layer driver of populist support; and negativity bias research supplies the asymmetry that explains why the cycle accumulates discontent rather than merely oscillating.

The second track is the product of reading the author's full subsequent corpus rather than only the passages that name PEFT — a deliberate methodological correction, since the richest extensions turned out to be written in other vocabularies entirely. The Fractal Karma model, read alongside its Enneagram layer, explains not only that different citizens, organizations, and political cultures fail in characteristically different ways, but the specific defensive mechanism — rule-stacking, narrative reframing, failure-externalization, forceful pre-emption — by which each worldview avoids registering its own failure. Argyris and Schön's single-loop/double-loop distinction gives "closure" an operational test. Ashby's requisite variety, formalized via the Bronze Mean threshold, gives PEFT's claim of institutional non-learning a numerical lower bound rather than only a qualitative description, and the quaternion state — four constitutionally independent organs, one of which (the coherence council) no existing democracy currently contains — gives that threshold a constitutional design response. Two twenty-year Dutch natural experiments, the VVD's 2005 manifesto and the documented real-estate revolving door, show the diagnosis operating in named, checkable cases rather than only in the abstract: one shows a correctly diagnosed script that was never revised, the other a script that was actively protected from ever generating the disconfirming evidence revision would require. And Chantal Mouffe's agonistic challenge, read fully rather than filtered out, is left standing rather than resolved: a theory whose only recommendation is more closure would itself become the kind of antagonism-suppressing technocracy its own diagnosis warns against, which is why this edition restricts PEFT's closure mechanism to performance and procedural scripts and treats relational and value disagreement as legitimate rather than as an unrevised script awaiting correction.

The path forward does not lie in tinkering with institutional design alone. It requires a fundamental shift in perspective — from the institution to the cognitive-collective — and the development of entirely new tools for reflection that empower citizens to become active agents in revising their own political scripts. QUO?, described in Section 7.1, is a first working instance of such a tool, now resting on a proven totality theorem rather than only a design intuition, built to do what the word "question" originally meant before literate, bureaucratic societies narrowed it to a form requiring an answer: to investigate. Whether it, or any instrument like it, measurably reduces political expectation failure in any actual population is a claim the architecture is now built to test, not a claim this edition makes on its behalf. The crisis of democracy demands not just better institutions, but better cognitive instruments for the citizens who must inhabit them, built on mechanisms — precision-weighting, bounded attention, loss aversion, competence perception, thermostatic feedback, worldview-specific failure topology, variety thresholds — that we can now actually measure, and, in at least one case, actually build and actually test.

Annotated Reference List

Altshuller, G. S. (1984/1996). *Creativity as an exact science*; *And suddenly the inventor appeared: TRIZ, the creative problem solving.* Technical Innovation Center. The source of TRIZ, distilled from analysis of roughly 200,000 patents into forty inventive principles and a contradiction matrix. Horowitz's (1999) further reduction of TRIZ to the two closed-world conditions underlying Section 7.1's account of QUO?'s design is the more direct source for this edition, but Altshuller's original corpus is the evidentiary base both rest on.

Alkemade, D. G. A. (2021, 2024, 2025). "Ragebol en Sabel: Patriotse burgerbewapening, revolutie en terreur in Holland, 1786–1787"; "Why was slavery not abolished in 1798?"; *Radicale democratie: Pieter Vreede en de Nederlandse Revolutie* (dissertation, Leiden University, forthcoming). Documents "constitutional domestication" — Willem I's post-1813 strategy of absorbing radical-democratic actors and institutions into a monarchical settlement that retained democratic form while keeping power concentrated. Section 5.4 uses this as the historical mechanism explaining why a (k)-organ-equivalent has not previously formed in Dutch governance: co-optation, not only suppression, is the documented risk any such body's design must resist.

Argyris, C., & Schön, D. A. (1978). *Organizational learning: A theory of action perspective.* Addison-Wesley. The distinction between single-loop learning (adjusting rules within unexamined governing assumptions) and double-loop learning (revising the assumptions themselves) supplies the formal vocabulary for what the original essay called, only qualitatively, the difference between script defense and script revision. Section 5.2's case studies of the VVD's 2005 manifesto and the documented real-estate revolving door are, in these terms, instances of sustained single-loop response to a correctly diagnosed double-loop problem, and of active prevention of the evidence double-loop revision would require.

Ashby, W. R. (1956). *An introduction to cybernetics.* Chapman & Hall. The law of requisite variety — only variety can absorb variety — underlies Section 5.3's Bronze Mean threshold and the broader cybernetic account, developed elsewhere in the author's corpus, of why systems with strong self-confirmation and weak self-correction mechanisms achieve short-term stability at the cost of long-term adaptive capacity. Gives PEFT's claim that institutions are "designed for stability, not learning" a formal rather than only descriptive basis.

Barber, B. (1984). *Strong democracy: Participatory politics for a new age.* University of California Press. A foundational text in participatory democratic theory. Barber argues for a more engaged citizenry beyond the thin model of representative democracy. In the context of PEFT, his work provides a precedent for moving beyond institutionalist views, though PEFT argues that participation alone is insufficient without the cognitive tools for expectation revision.

Baumgartner, F. R., & Jones, B. D. (1993/2009). *Agendas and instability in American politics.* University of Chicago Press. The foundational text of punctuated equilibrium theory, built on the "error accumulation model" of bounded institutional attention. This edition uses PET to replace panarchy theory as PEFT's primary account of systemic non-learning, because PET offers a directly tested mechanism — rather than an ecological analogy — for why institutions attend to few issues at a time, letting failure accumulate silently elsewhere until a punctuation event forces correction.

Cambridge Analytica files; Zuboff, S. (2019); Laclau, E. (2005). *The Guardian* investigative series; *The Age of Surveillance Capitalism* (PublicAffairs); *On Populist Reason* (Verso). Sources for Section 6.1's account of industrialized expectation construction: Cambridge Analytica's psychographic microtargeting of up to 87 million Facebook users; Zuboff's structural explanation for why commercial behavioral-influence infrastructure transfers so readily to political use; Laclau's account of populist effectiveness as deriving from constructed antagonism rather than policy elaboration, used to interpret this edition's measured inverse relationship between electoral effect and value-consistency in the Dutch party system.

Clarke, H. D., Sanders, D., Stewart, M. C., & Whiteley, P. F. (2004, 2009, 2011). *Political choice in Britain*; *Performance Politics and the British Voter*; "Valence Politics and Electoral Choice in

Britain, 2010." Oxford/Cambridge University Press; *Journal of Elections, Public Opinion and Parties*. The founding works of valence politics theory, which argues that voters choose between parties on perceived competence to deliver consensus goals rather than disagreement over the goals themselves. This edition identifies valence theory as PEFT's most direct theoretical neighbor and its empirical pillar's most mature available measurement tradition.

Clarkson, C., et al. (2017). "Human occupation of northern Australia by 65,000 years ago." *Nature*, 547, 306–310. The archaeological dating evidence for the migration carrying San-derived coherence practice to Australia roughly 65,000 years ago, underlying Section 1.1's timeline.

Cremaschi, S., Inglehart, R., & De Vries, C. E. (2024). "Geographies of discontent: Regional inequality and political dissatisfaction in Europe." *European Journal of Political Research*, advance online publication. A key empirical contribution from the third phase of De Vries's program. It documents the spatial and social clustering of political dissatisfaction, demonstrating that expectation failure is not a uniform phenomenon but is geographically embedded, creating persistent zones of institutional alienation.

De Vries, C. E. (2007). "Sleeping giant: Fact or fairytale? How European integration affects national elections." *European Union Politics*, 8(3), 363–385. An early work establishing the concept of latent political expectations. It shows that electorates hold dormant scripts regarding European integration that can be activated, shaping national electoral outcomes.

De Vries, C. E. (2018a). *Euroscepticism and the future of European integration.* Oxford University Press. A pivotal work that established the benchmark theory of EU public opinion. It provides a core empirical foundation for PEFT by demonstrating that dissatisfaction with the EU is a function of script failure — a deviation from nationally-derived reference frames — rather than a response to objective institutional performance.

De Vries, C. E. (2018b). "What causes economic perceptions? The role of media, partisanship and economic experience." *Electoral Studies*, 51, 82–93. This article reinforces the cognitive nature of political evaluation by showing that political behavior responds to perceived, rather than objective, economic conditions. This confirms the PEFT premise that it is the subjective discrepancy between expectation and experienced reality that drives political behavior.

De Vries, C. E., & Hobolt, S. B. (2020). *Political entrepreneurs: The rise of challenger parties in Europe.* Princeton University Press. A cornerstone work for PEFT. It introduces and systematically analyzes the role of political entrepreneurs as the key agents who identify, amplify, and exploit expectation failures for electoral gain, transforming them from accidental byproducts into a cultivated political resource.

De Vries, C. E. (2026, forthcoming). *Symfonie van onvrede.* Amsterdam University Press. The culminating work in De Vries's program. It conceptualizes the rise of the radical right as the cumulative political expression of decades of unresolved expectation failure, positioning this phenomenon as a systemic immune response to a chronic condition.

Dryzek, J. S. (2000). *Deliberative democracy and beyond.* Oxford University Press. A key text in deliberative democratic theory. While PEFT shares its focus on the quality of citizen engagement, it argues that deliberation is often hampered by unexamined expectations, and that it must be preceded by cognitive reflection on the scripts participants bring to the table.

Ferwerda, J., et al. (2024/2025). Comparative survey research on nostalgic deprivation across nineteen European countries. Identifies nostalgic deprivation — simultaneous economic, social, and political status loss — as a stronger predictor of populist support than economic anxiety alone, across both left- and right-populist parties. Supplies Section 3.2's cross-national mechanism for the relational layer of PEFT and grounds the original essay's claim of global PEF diffusion in specific, comparative survey evidence rather than only structural analogy.

Fiske, A. P. (1992). "The four elementary forms of sociality: Framework for a unified theory of social relations." *Psychological Review*, 99(4), 689–723. Identifies four elementary relational modes — communal sharing, authority ranking, equality matching, and market pricing — from which the author's subsequent work derives a "Fiske vector" for representing the relational content of political scripts. Supplies PEFT's cognitive pillar with a natural language for what a citizen's relational script (Section 2) actually consists of, and for comparing it directly against the relational signature of specific governing decisions or party platforms.

Friston, K. (2010). "The free-energy principle: A unified brain theory?" *Nature Reviews Neuroscience*, 11(2), 127–138. The source text for the free energy principle, on which predictive processing's account of precision-weighted belief updating is built. This edition uses it as a formal heuristic for political belief revision — and its failure — rather than as a settled neural mechanism. Section 2.2 extends this further via the author's own work, treating the citizen as a Bayesian agent whose chronically elevated free energy — the citizen's experience of political expectation failure — is itself shaped by the specific quaternion-weighted worldview through which political reality is interpreted.

Frissen, P. H. A. (2026). *De neutrale staat: Pleidooi voor een conservatief pluralisme.* Boom. Argues that the state should confine itself to protecting disagreement and abandon ambitions of social engineering. Section 5.3's quaternion state is positioned as occupying the middle ground between Frissen's retrenchment and the 2005 *Liberaal Manifest*'s redesign ambition: PEFT agrees that neutrality alone does not solve the cognitive problem, since a state that merely protects conflict without helping the parties to it recognize their own blind spots leaves the structural causes of escalation untouched.

Gunderson, L. H., & Holling, C. S. (Eds.). (2002). *Panarchy: Understanding transformations in human and natural systems.* Island Press. The source text for panarchy theory, used in the original essay to frame democratic instability as a pathological adaptive cycle. This edition retains panarchy as a useful descriptive image but replaces it as the theory's primary mechanism with punctuated equilibrium theory (Baumgartner & Jones, above), which offers a directly tested political analogue rather than an ecological one.

Heaviside, O. (1884–1889). Reduction of Maxwell's quaternion-based equations to vector form for electrical engineering use. The historical and physical anchor for Section 1.1's outside-story/inside-story distinction: Heaviside's removal of the scalar component from Maxwell's original quaternion formulation of electromagnetism is treated, in the author's broader corpus, as the technical prototype for a more general pattern of discarding a coherence-and-relation register in favor of a force-and-causation register — a pattern this edition argues PEFT's three pillars are independently working to reverse in political theory.

Konstapel, H. (2006). *De Geschiedenis van het Cyclische Denken*, version 7. Constable Research. An unpublished working paper, predating PEFT, SWARP, and the 19LQVM by two decades,

independently developing the coherence-versus-control distinction from comparative anthropology (the Medicine Wheel's "preserving the centre" principle), Chinese five-phase cosmology (the Sheng/Ko/Wu cycles), Vedic cosmology (the Tattvas), and Edward Dewey's Foundation for the Study of Cycles. Identifies Aristotle, rather than Heaviside, as the historical figure who reduced cyclical, multi-sphere cosmology to linear causality and a single material sphere, and identifies a recurring "Discovery followed by Exploitation" two-phase pattern across Hellenistic, Roman, and colonial-era history. Cited in Section 1.1 as evidence that this edition's central distinction is the most recent formal statement of an argument the author has refined continuously since at least 2006, not a theory invented retroactively to fit older material.

Konstapel, J. (2026). *Understanding Language and Evolution Through the Sphenoid.* constable.blog. The source for Section 1.1's foundational distinction between the outside story (selection pressure, incentive, external force) and the inside story (a system actively maintaining its own coherence and crossing internal thresholds). Connects Dambricourt Malassé's sphenoid research, Cotterill's "probe-by-movement" account of cognition, Maxwell's original quaternion electromagnetism, and the documented colonial-era suppression of coherence-based knowledge traditions into a single argument that does not mention PEFT, expectation failure, or democracy anywhere in its text — and is, on the reading defended in Section 1.2, one of the most consequential sources for this edition precisely because it does not.

Konstapel, J. (2026). *De Indusbeschaving: Een Geschiedenis zonder Macht*; *Het Jainisme en de Oudste Beschaving op Aarde*; *200.000 jaar Kennisvergaring van de Mensheid in Kaart.* constable.blog. The sources for Section 1.1's precise dating of the pre-3117-BCE coherence-based period, the Piora oscillation as the documented historical trigger for the shift to control-based social organization (the "first epistemocide"), and the Indus Valley Civilization as the corpus's primary empirical case for sustained, large-scale, low-power-concentration social organization. Introduces *aparigraha* as a constraint independently derivable from quaternion algebra applied to social eigenstates.

Konstapel, J. (2026). *65,000 Jaar Wetenschapsgeschiedenis* (*Coherence Intelligence Across 65,000 Years*). constable.blog. The source for the eleven-tradition cross-civilizational survey (Aboriginal Australian, Hopi, Kabbalah, Ifá, Dogon, Zoroastrian, Taoist, Vedic, Mesopotamian, Maya, Egyptian) underlying Section 1.1's claim that coherence-based knowledge systems constitute parallel, operationally effective methodologies rather than pre-scientific superstition, and the documentary basis for dating the second epistemocide (the colonial-era suppression of these traditions, 1500–1980 CE) referenced in the same section.

Maxwell, J. C. (1865). "A dynamical theory of the electromagnetic field." *Philosophical Transactions of the Royal Society*, 155, 459–512. The original quaternion-algebra formulation of electromagnetism, including the scalar component later dropped by Heaviside (above). Functions in this edition as the physical case study underlying the outside-story/inside-story distinction of Section 1.1, not as a claim about political science in its own right.

Horowitz, R. (1999). *Creative problem solving in engineering design.* Doctoral dissertation, Tel-Aviv University. Reduces Altshuller's TRIZ to two closed-world conditions and develops ASIT (Advanced Systematic Inventive Thinking). Horowitz's own dissertation flags, without resolving, the open problem of how to define a "closed world" outside engineering. Section 7.1's Totality Theorem answers that open problem for the specific domain of question-answering by identifying

the closed world of any communicative problem with its universe of discourse, giving QUO?'s design a formal proof rather than only an engineering intuition.

Jones, B. D., & Baumgartner, F. R. (2005). *The politics of attention: How government prioritizes problems.* University of Chicago Press. Extends punctuated equilibrium theory with comparative cross-national evidence linking institutional design — particularly the number of independent venues processing an issue — to the severity of policy punctuation. Supplies PEFT with a falsifiable institutional variable in place of the original's general claim that institutions are "designed for stability, not learning."

Kahneman, D., & Tversky, A. (1979). "Prospect theory: An analysis of decision under risk." *Econometrica*, 47(2), 263–291. The founding statement of loss aversion, on which Soroka's negativity bias research (below) is built. Supplies the psychological basis for the asymmetry between expectation failure and expectation fulfillment in the PEF cycle.

Kenoyer, J. M. (1998). *Ancient cities of the Indus Valley civilisation.* Oxford University Press. The standard archaeological synthesis documenting the Indus Valley Civilization's absence of fortifications, weapons caches, conquest-palaces, and domination iconography across two millennia of urban settlement. Section 1.1 uses this as the empirical anchor for treating low-power-concentration social organization as a documented historical achievement rather than only a theoretical possibility.

Jaini, P. S. (1979); Matilal, B. K. (1981). *The Jaina path of purification* (University of California Press); *The central philosophy of Jainism* (L. D. Institute of Indology). Standard academic treatments of Jain ontology and the *anekāntavāda* (many-sidedness) doctrine. Section 1.1 draws on the *aparigraha* (non-possession) principle from this tradition as an architectural rather than merely ethical constraint on institutional power concentration, paralleling the quaternion state's design goals in Section 5.3.

Konstapel, J. (2026). *Van Stem naar Profiel: Het Personal Political Profile.* constable.blog. Translates PEFT, the Fractal Karma model, Friston's free energy principle, and Fiske's relational models theory into a single applied architecture for citizen-political matching (the PPP). The source for this edition's Section 2.2 account of worldview-specific failure topology and for treating the citizen as a Bayesian agent whose generative model of "what good governance feels like" is updated — or fails to update — according to the same quaternion that governs script revision more broadly.

Konstapel, J. (2026). *Waarom D66 in 60 Jaar Niets Heeft Bereikt.* constable.blog. The sister case to Section 5.2's VVD analysis, applying the same single-loop/Blue-dominant diagnostic to a different Dutch governing party, confirming that the failure pattern is worldview-bound rather than party-specific.

Konstapel, J. (2026). *De Vastgoedstaat.* constable.blog. The source for Section 5.2's second natural experiment: documents the revolving door between Dutch governing parties and the real-estate development sector, commissioned research with circular legitimization, participation processes restricted to cosmetic detail, and the deliberate non-integration of public records that would otherwise make the gap between official communication and documented decision visible.

Konstapel, J. (2026). *De Gapende Kloof tussen Pro en Zijn Kiezers in Leiden.* constable.blog. The source for Section 4.1's demographic boundary on the thermostatic mechanism: documents

structural non-participation by income, education, and tenure status in Leiden municipal elections, and the resulting narrow effective electoral base of the governing coalition.

Konstapel, J. (2026). *Awen Grid* (companion essay integrating the Recursive Harmonic Codex with the 19-Layer Quaternion Vacuum Model). constable.blog. The source for the formal derivation of Section 5.3's Bronze Mean threshold: the three imaginary units of quaternion algebra force a three-fold recursive branching whose natural fixed point is $\beta = (3+\sqrt{13})/2$, generating the four named coherence phases (minimal coherence, integration, autopoietic capacity, recursive self-reference) this edition maps onto PEFT's distinction between institutional self-maintenance and institutional self-revision. Also the source for this edition's explicit falsifiability standard: threshold-clustering, rather than continuous variation, in an empirical proxy for coherence capacity is the stated condition whose absence would falsify the selection rule itself.

Konstapel, J. (2026). *Waarom Nederland Nooit Verandert*; *De Mislukte Drempel: Nederlandse Institutionele Geschiedenis als Coherentie-Faseovergang.* constable.blog. The source for the historical application of the Bronze Mean threshold to the 1619 Synod of Dordrecht and subsequent Dutch institutional forms (VOC, pillarization, the polder model), and for the four-organ quaternion state constitutional design. This edition treats the historical application as an interpretation fitted to known outcomes rather than a confirmed prediction, per the falsifiability standard stated in the Awen Grid source above.

Konstapel, J. (2026). *Progress Report on the Development of SWarp.* constable.blog. The source for Section 7.1's comparison of QUO?/SWARP against existing civic technology (Polis, SenseMaker, Decidim, Loomio, Solid, voting-advice applications) and for the explicit epistemic caveat that the architecture is instrumented to test a reduction in expectation failure, not yet shown to achieve one in any measured population.

Konstapel, J. (2026). *De Vraag als Gesloten Wereld: ASIT, het Universum van Discourse, en de Erotetische Grondslagen van SWARP-3.* constable.blog. The source for Section 7.1's Totality Theorem: proves that a system can treat every possible question honestly across eight exhaustive treatment classes even though no system can correctly answer every possible question, and connects this result to Horowitz's (1999) closed-world reduction of TRIZ.

Konstapel, J. (2026). *Het Vergeten Manifest: Waarom de VVD Haar Eigen Diagnose Twintig Jaar Heeft Genegeerd.* constable.blog. The source for this edition's Section 5.2 case study: a detailed twenty-year tracing of a political party's correct diagnosis of relational script failure and its subsequent, almost perfectly selective, single-loop response. Introduces the fractal state as a constructive alternative to both the original manifesto's procedural remedies and Frissen's (2026) neutral-state retrenchment.

Konstapel, J. (2026). *Quo? — Het Vraagteken Krijgt Zijn Werk Terug.* constable.blog. Describes QUO?, the first working citizen-level reflective instrument built on the SWARP platform, discussed in this edition's Section 7.1. Documents the empirical failure rate of existing municipal chatbot systems and specifies the four classes of honest non-answer the system gives in place of a guessed one.

McWhinney, W. (1997). *Paths of change: Strategic choices for organizations and society.* Sage. The four-worldview theory (analytic, sensory, social, mythic) underlying the Fractal Karma quaternion introduced in Section 2.2. Explains why a relational political problem, approached only

with analytic instruments, tends to receive only procedural answers — a category mismatch this edition identifies as a recurring signature across multiple Dutch political case studies.

Parpola, A. (1994). *Deciphering the Indus script.* Cambridge University Press. Four decades of work identifying the Indus script's most frequent sign as a star/fish homophone in Proto-Dravidian, supporting an astronomical rather than purely commercial interpretation of Indus seals. Section 1.1 uses this to support the claim that Indus urban planning, astronomically aligned to cardinal directions, expresses the same "as above, so below" coherence principle documented independently in Aboriginal songlines.

Mouffe, C. (2000/2005). *The democratic paradox*; *On the political.* Verso. Argues that the political is irreducibly conflictual and that suppressing antagonism in favor of rational consensus produces a democratic deficit that creates the conditions for right-wing populism. Section 2.3 treats Mouffe's challenge as a serious, unresolved internal tension for PEFT rather than dismissing it: a theory whose only design recommendation is closure risks becoming exactly the kind of antagonism-suppressing technocracy Mouffe warns against, which this edition addresses by restricting PEFT's closure mechanism to performance and procedural scripts.

Pfeijffer, I. L. (2025). *Absolute Democratie.* De Bezige Bij. A collection of essays diagnosing democracy as increasingly dismantled from within, in democracy's own name, rather than only threatened from outside — majoritarian or winning factions eroding rule-of-law checks, press freedom, and minority protections under cover of electoral legitimacy. Section 5.4 reads this as the end state of PEFT's Phase 4 (non-resolution) when no (k)-organ-equivalent exists to interrupt the process, since the institution that would name the line being crossed is exactly the kind of independent, executive-resistant body absolute democracy's logic erodes first.

Schank, R. C., & Abelson, R. P. (1977). *Scripts, plans, goals and understanding: An inquiry into human knowledge structures.* Lawrence Erlbaum. The foundational text for script theory. It provides the cognitive architecture that underpins PEFT, defining scripts as the mental structures that organize human comprehension and action, and establishing the premise that failure is the engine of learning.

Schank, R. C. (1982). *Dynamic memory: A theory of reminding and learning in computers and people.* Cambridge University Press. Expands upon script theory to detail how cognitive systems process failure. It identifies the critical steps of explanation, reminding, and generalization that lead to script revision. PEFT's central concept of "non-learning" is defined by the absence of this sequence in the political domain.

Schumpeter, J. A. (1942). *Capitalism, socialism and democracy.* Harper & Brothers. The classic articulation of the elite-competitive model of democracy. PEFT directly challenges this model, arguing that if elite actors are incentivized to exploit, rather than resolve, expectation failure, then elite competition becomes a driver of instability rather than a mechanism for democratic improvement.

Schlebusch, C. M., et al. (2017). "Southern African ancient genomes." *Science*, 358(6363), 652–655. The genetic evidence dating the San lineage to roughly 200,000–300,000 years, the deepest surviving human population branch, used in Section 1.1 as the biological anchor for the longer timeline PEFT is read against.

Lewis-Williams, D. (2002). *The mind in the cave.* Thames & Hudson. The standard scholarly treatment of San rock art and trance-based consciousness practice, used in Section 1.1 to support treating San coherence practice as a documented, continuously practiced knowledge tradition rather than a speculative reconstruction.

Snowden, D. J., & Boone, M. E. (2007); Holling, C. S. (2001); Gunderson, L. H., & Holling, C. S. (2002). "A leader's framework for decision making" (*Harvard Business Review*); "Understanding the complexity of economic, ecological, and social systems" (*Ecosystems*); *Panarchy* (Island Press). Sources for Section 5.1's Cynefin/panarchy integration and the documented Dutch case material: structural waiting times across nearly every public domain (healthcare, asylum, housing, infrastructure) read as a "rigidity trap" in late-K, alongside the Dutch Delta Programme for water safety as a working counter-example of adaptive governance that has avoided this trap in one specific domain for years.

Soroka, S. N. (2014). *Negativity in democratic politics: Causes and consequences.* Cambridge University Press. Demonstrates that negative political and economic information is weighted roughly twice as heavily as positive information of equivalent magnitude. Supplies PEFT's PEF cycle with the asymmetry it previously lacked, explaining why the cycle accumulates discontent over successive iterations rather than oscillating around a stable equilibrium.

Soroka, S., & Wlezien, C. (2010). *Degrees of democracy: Politics, public opinion, and policy.* Cambridge University Press. Cross-national confirmation of the thermostatic model across the US, UK, and Canada, specifying the conditions (issue salience, centralized institutional responsibility) under which the public does learn from policy. Supplies PEFT's non-learning claim with a precise, falsifiable boundary.

Konstapel, J. (2025). *Fractale Democratie: Van Vertrouwenscrisis naar Wijkcirkels.* constable.blog. The source for Section 5.4's Judging/Perceiving refinement of Fiske's four relational modes, applied to the 2025 Dutch party system, and for the wijkcirkel (neighborhood-circle) sub-national implementation proposal, including its own seven-framework falsifiability review (Ostrom, deliberative democracy, social identity theory, complexity theory, public choice theory, the capabilities approach, historical institutionalism).

Stacey, R. D. (2010). *Complexity and organizational reality.* Routledge. A key text applying complexity theory to organizational dynamics. It provides the conceptual tools for understanding non-linear dynamics and emergent behavior, retained alongside punctuated equilibrium theory as a complementary description of how democratic systems, as complex adaptive systems, can become trapped in non-learning configurations.

Wlezien, C. (1995). "The public as thermostat: Dynamics of preferences for spending." *American Journal of Political Science*, 39(4), 981–1000. The founding statement of the thermostatic model of opinion-policy feedback. Supplies the tested counter-case against which PEFT's claim of structural democratic non-learning can be made precise rather than absolute.