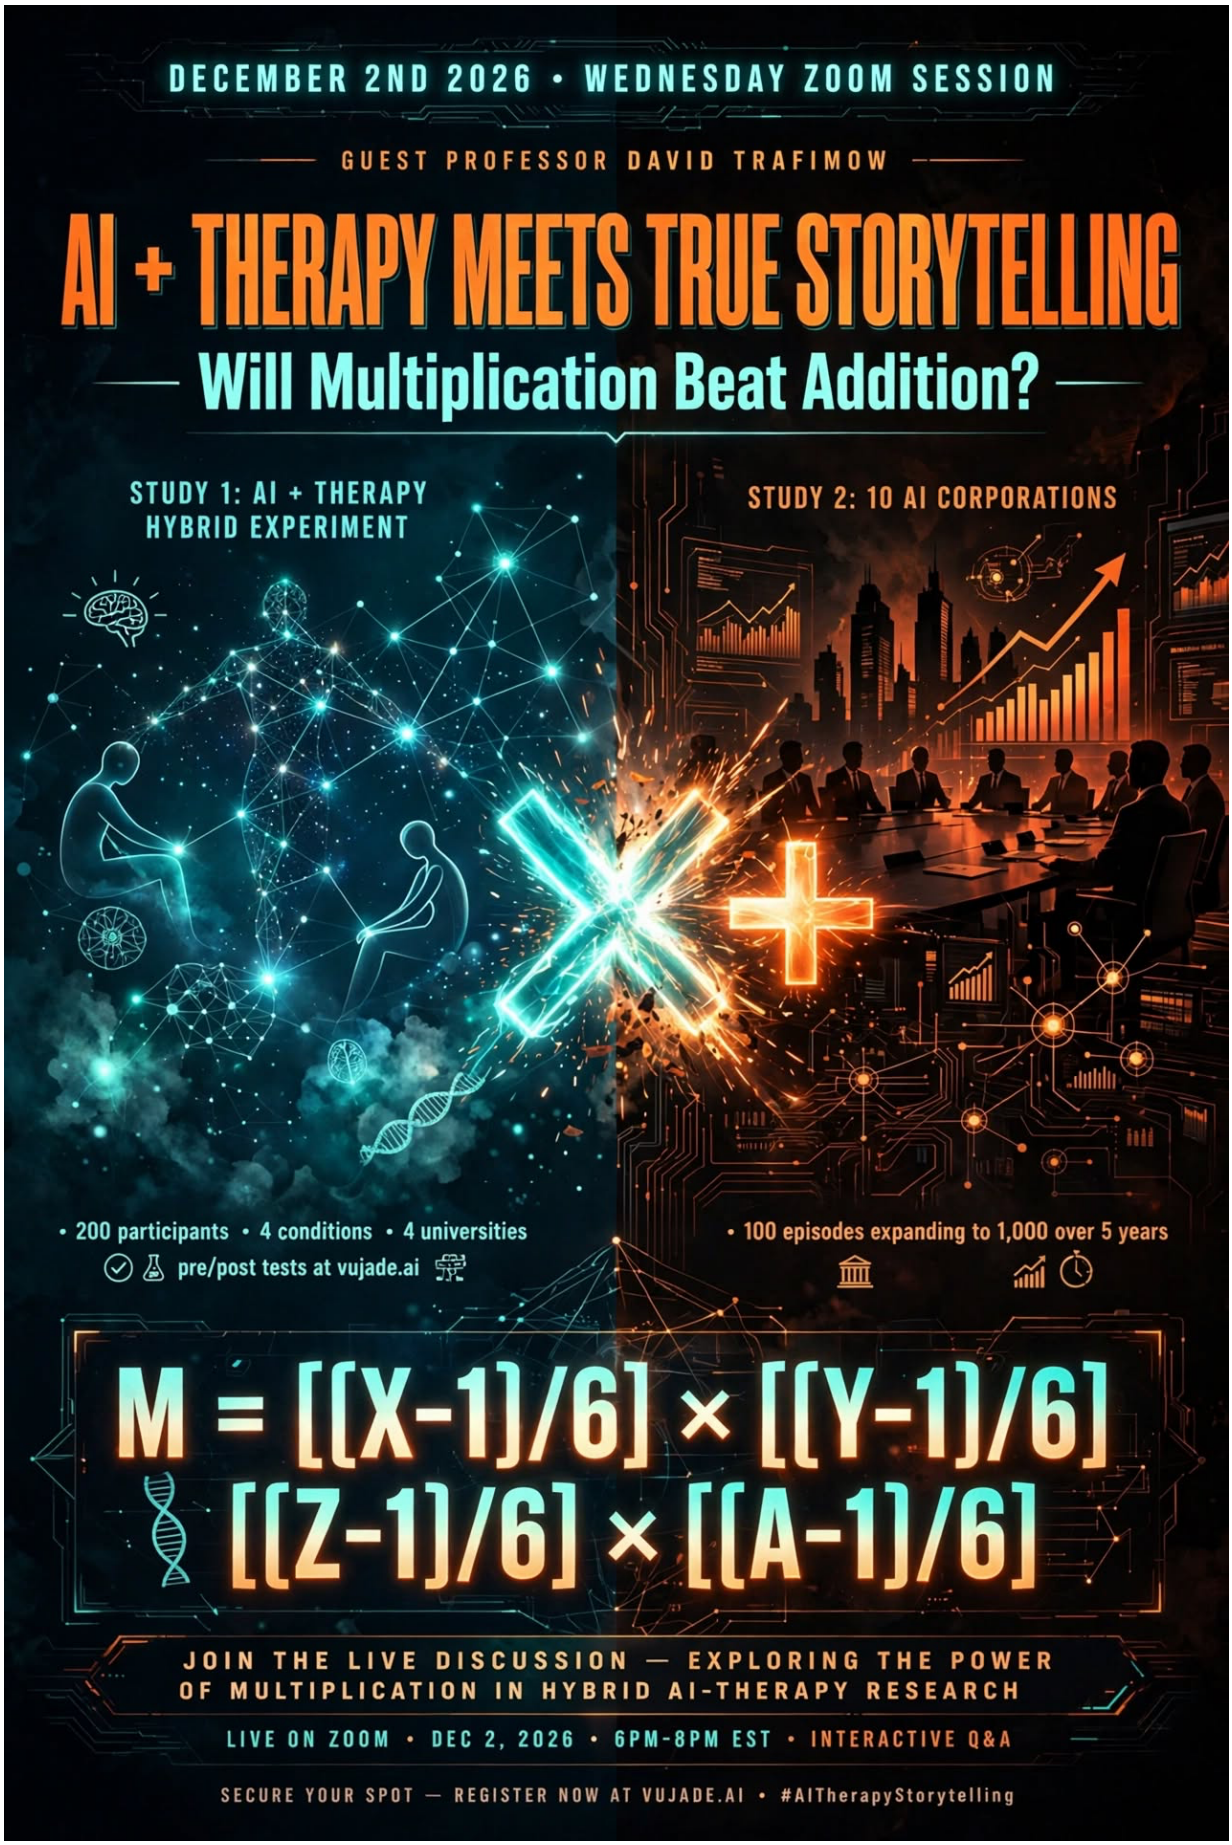
# Methodology Revision Report to David Trafimow, Monty Miller, and Doug Breckenridge

## The Concord Leadership Tesseract™ (XYZA) and VuJaDé™ Public Discourse Layer

David Michael Boje | Nov. 17 2026

A revised, auditable protocol for event identification, verification, behaviorally anchored scoring, and inter-rater assessment


DECEMBER 2ND 2026 • WEDNESDAY ZOOM SESSION
GUEST PROFESSOR DAVID TRAFIMOW
AI + THERAPY MEETS TRUE STORYTELLING
Will Multiplication Beat Addition?
STUDY 1: AI + THERAPY HYBRID EXPERIMENT
STUDY 2: 10 AI CORPORATIONS
• 200 participants • 4 conditions • 4 universities
pre/post tests at vujade.ai
• 100 episodes expanding to 1,000 over 5 years
M = [(X-1)/6] × [(Y-1)/6]
[(Z-1)/6] × [(A-1)/6]
JOIN THE LIVE DISCUSSION — EXPLORING THE POWER OF MULTIPLICATION IN HYBRID AI-THERAPY RESEARCH
LIVE ON ZOOM • DEC 2, 2026 • 6PM-8PM EST • INTERACTIVE Q&A
SECURE YOUR SPOT — REGISTER NOW AT VUJADE.AI • #AITherapyStorytelling

## Abstract

Leadership theory has repeatedly widened its field of view: from the leader's traits or style, to task and people behaviors, to the context in which behaviors work, and then to nonlinear networks that evolve over time. This chapter presents that progression for organizational development and change (ODC) practitioners and introduces David M. Boje's Concord Tesseract™ XYZA framework as an ethical, event-based extension. X tracks transactional-to-transformational practice; Y tracks power-over to serving others; Z tracks monophonic control to polyphonic participation; and A tracks alibi to Bakhtinian moral answerability. The chapter explains how the four dimensions can guide diagnosis, conversation, and intervention without requiring advanced mathematics. A brief NVIDIA transition illustrates how documented events, attributed reports, and market reactions can be read as a network shock under an explicit auxiliary assumption. The chapter concludes with a practitioner protocol and limitations concerning scoring, evidence, geometric representation, and validation in organizations facing rapid change.

*Keywords:* leadership theory, organizational development and change, complexity leadership, Concord Tesseract, XYZA model, moral answerability, True Storytelling

### Introduction to the Methodology

This report documents the methodological reset underway for the study of AI corporate leadership. The revised architecture pairs the Concord Leadership Tesseract™ (XYZA), which codes four dimensions of leadership conduct and answerability in dated events, with the VuJaDé™ Public Discourse Layer (V1–V4), which codes observable features of public communication. Both use behaviorally anchored seven-point ordinal ratings and a normalized multiplicative event index. The purpose is to preserve the multidimensional profile of each event and prevent a high score on one dimension from being treated as a substitute for a weak score on another.

## Purpose, Audience, and Boundaries

This is an update intended to explain how the research method is changing, why a separate methods chapter is now needed, what the attached workbook currently contains, and how the next version can be made auditable. It is written as a methodological report suitable for circulation among scholarly colleagues who will examine coding, evidence, and mathematical procedures.

The report is bounded to the source material supplied in this conversation: the earlier methodology draft, the November 17 revision-history notes, and the workbook named AI_Leadership_XYZA_Vujade_Expanded_v03_NV_OAI_Nov17_2026.xlsx.

It does not claim to have independently verified the workbook's citations, live URLs, news claims, or every quotation against external originals. It describes workbook contents as an inspected working snapshot and marks unresolved verification as unresolved.

Names of prospective reviewers and plans for contacting them are outside the scope. The report does not present a finished book chapter or a final empirical analysis of all ten leaders. It sets out a

repeatable protocol that can support the dedicated methodology chapter and later case chapters. Numerical summaries are descriptive of the workbook rows only; they are not inferential findings about leaders, corporations, or the AI industry.

Where the narrative notes make claims that differ from workbook structure or fields, the workbook is used to describe the present file and the notes are treated as evidence of the intended workflow and revision history. This separation protects the report from turning planned procedures into completed results.

## Ten-Corporation Scope and Coverage

The Research Monograph I am writing identifies ten corporations: NVIDIA, OpenAI, Anthropic, xAI, Oracle, Meta, Microsoft, Amazon, Google/Alphabet, and DeepSeek. The current event ledger is not evenly distributed across them. It contains 100 NVIDIA rows, 95 OpenAI rows, and 40 rows across the other eight organizations. Microsoft has six; Anthropic, xAI, Oracle, Meta, Google/Alphabet, and DeepSeek each have five; Amazon has four.

This distribution is appropriate for an evolving archive and detailed case development, but it is not a balanced comparative sample. NVIDIA and OpenAI together account for 195 of the 235 rows, or approximately 83 percent of the ledger, thus far. Accordingly, any overall averages are dominated by those two firms and should not be described as estimates of all ten corporations. That serves chapter 2 on NVIDIA, and then the database is expanded day by duy until all ten AI corporations and their leaders are analyzed.  Therefore, the book has an arc of development of the method, theory, and possible practices of ODC.

I want define a coverage plan before adding large batches and get your input. One option is to set a minimum evidence target by organization and event family, then expand beyond the minimum when consequential events warrant. Another is to keep a purposive archive but explicitly limit conclusions to the selected episodes. Either choice can be methodologically defensible if it is declared and applied consistently.

Coverage should also be checked over time: years, source types, event families, leader transitions, and corporate roles. A table of counts is a descriptive diagnostic, not a quality score. The method should explain why an organization has few rows, not conceal that imbalance.

The change responds to a moving empirical environment. Governance events, leadership transitions, political shifts, and technology-sector shocks can alter the evidence base and the interpretation of earlier snapshots. The design therefore moves from a relatively small pilot toward a dated, source-linked event ledger that can be refreshed, re-audited, and extended across ten firms and a rolling five-year period. The planned book structure requires a stable methodology chapter before successive company chapters are revised.

The attached workbook contains 235 unique event rows, 100 labeled NVIDIA, 95 OpenAI, and 40 across the other eight firms. It spans June 16, 1977, through November 17, 2026; 20 rows predate

the five-year window ending November 17, 2026. Its formulas cover the event-level multiplicative indices. However, the workbook does not yet document completed independent double coding: the second-rater field is blank in all 235 rows, and there is no inter-rater agreement worksheet in the supplied file. Consequently, this report distinguishes a large coded working ledger from an independently verified and reliability-tested corpus.

The immediate methodological priority is to resolve that distinction, confirm the rolling-window cohort, and set explicit source, context, coding, adjudication, and agreement gates. The target of at least 100 verified episodes is treated as a validation milestone, not as already achieved by a row count.

### Table 1: Ten corporations and nonredundant analytical themes

| Corporation | Leadership focus | Distinctive chapter contribution |
|---|---|---|
| NVIDIA | Jensen Huang & Successor | Infrastructure chokepoints |
| OpenAI | Sam Altman | Governance and accountability diffusion |
| Anthropic | Dario Amodei | Constitutional commitments and reflexivity |
| xAI | Elon Musk | Adversarial public reframing |
| Oracle | Larry Ellison | Enterprise data concentration |
| Meta | Mark Zuckerberg | Scale and epistemic neutrality |
| Microsoft | Satya Nadella | Management discourse and Copilot |
| Amazon | Andy Jassy | Logistics, labor and refusal |
| Google | Sundar Pichai | Search mediation and answer authority |
| DeepSeek | Liang Wenfeng | Language, jurisdiction and boundary conditions |

For every corporation I propose three distinct evidence streams: dated executive texts and corporate documents; public social media and independent reporting; and product transcripts collected under a reproducible prompt protocol where an equivalent interface exists.

### Table 2: Chapter-by-chapter ODC approaches and rationale

| Case / chapter | ODC approach and status | Why it may fit |
|---|---|---|
| Ch. 2 — Jensen Huang / NVIDIA | Fourth Wave Grounded Theory (AID) — used in the chapter | Supports an evidence-led, self-correcting analysis: proposed explanations are tested against dated records and revised when evidence contradicts them. |

| Case / chapter | ODC approach and status | Why it may fit |
|---|---|---|
| Ch. 3 — Sam Altman / OpenAI | Dialogic Organization Development — selected | Makes room for competing stakeholder voices and dialogue that can inform organizational decisions. World Café and Future Search can structure dialogue; SEAM can examine dysfunction costs and negotiated change. |
| Ch. 4 — Dario Amodei / Anthropic | True Storytelling, with a proposed seven-principle / seven-B crosswalk — assigned | Tests whether safety commitments extend beneath public stories and lead to reflection, answerability, and changed practice. |
| Ch. 5 — Elon Musk / xAI | Restorying — candidate | Could compare founder and company accounts with counter-stories and stakeholder consequences. Competing narratives and whose account becomes authoritative are central questions. |
| Ch. 6 — Larry Ellison / Oracle | Systemic Organizational Constellation Method — exploratory candidate | Could map relationships and tensions across Oracle's enterprise, cloud, health, and acquired-product systems. Because the evidence base is limited, use should remain exploratory and be evaluated transparently. |
| Ch. 7 — Mark Zuckerberg / Meta | SEAM — assigned | Socio-economic inquiry can examine hidden dysfunction costs and how they are distributed across work and stakeholder relationships. |
| Ch. 8 — Satya Nadella / Microsoft | Double-loop and triple-loop learning — assigned | Could examine whether learning changes operating assumptions and governance principles, as well as procedures. |
| Ch. 9 — Andy Jassy / Amazon | High Reliability Organizing (HRO) — selected | Could focus on operational safety: noticing weak signals, learning from near misses, resisting overly simple explanations, and giving weight to frontline expertise. Amazon's nationwide OSHA ergonomics settlement makes safety systems and follow-through a concrete case focus. |
| Ch. 10 — Sundar Pichai / Google | Quantum Storytelling intervention — developmental candidate | Could examine how organizational stories and material conditions interact across AI research, infrastructure, and product groups. Google's integration of Brain and DeepMind makes cross-boundary inquiry relevant; the intervention design needs further specification. |
| Ch. 11 — Liang Wenfeng / DeepSeek | Transorganizational Networking Development — assigned | Could examine collaboration and accountability across the organization's wider AI and infrastructure network. |
| Ch. 12 — Conclusion of ODC | Concord Tesseract™ Leadership (CTL) — integrative ODC proposal | Synthesizes the cases into a movement from the dark side of leadership toward transmutative enlightenment. Concord joins contemplation with action to create new possibilities for enlightenment, unity, peace, justice, and equanimity in a prospective future. |

| Dimension | Low pole (1) | High pole (7) | What the coding examines |
|---|---|---|---|
| X | Transactional | Transformational | Observable leadership practices, from exchange/task control toward developmental change and transformation. |
| Y | Power-seeking and self-oriented | Serving others | How authority and organizational resources are oriented toward the leader's own power or toward serving others. |
| Z | Monophonic | Polyphonic | Whether other voices can speak, remain distinct, and affect organizational understanding and action. |
| A | Alibi or evasion | Moral answerability | Whether actors acknowledge their participation, answer to affected others, explain decisions, and respond to consequences. |

The rubric keeps X, Y, Z, and A analytically distinct: transactional/transformational practice; power/service orientation; monophony/polyphony; and alibi/moral answerability. Polyphony is not merely counting how many stakeholders appear in a document; it asks whether their voices remain unmerged and have standing in decisions.

The demonstration's multiplicative procedure is applied after the case evidence has been coded using the seven-point behavioral anchors. For dimension scores X, Y, Z, and A, the single composite is the normalized product:

$$\mathbf{M = [(X - 1)/6] \times [(Y - 1)/6] \times [(Z - 1)/6] \times [(A - 1)/6]}$$

The transformation maps each 1–7 placement onto a 0–1 scale and multiplies the four components. The result is a theoretical summary for the demonstration, not a probability. It does not replace the four component scores, the evidence record, or the interpretation of a case. The book consistently calls it a qualimetric/multiplicative demonstration informed by Trafimow's theoretical argument for multiplication; it does not claim Trafimow supplied or validated this particular XYZA specification.

This is a technical specification of the chosen method.

**Figure I: The Concord Tesseract™ operational illustration as the provisional APP starting map (first success Nov 13, 2026) – It will be recalculated once the full data is assembled for WorkdSheet.**

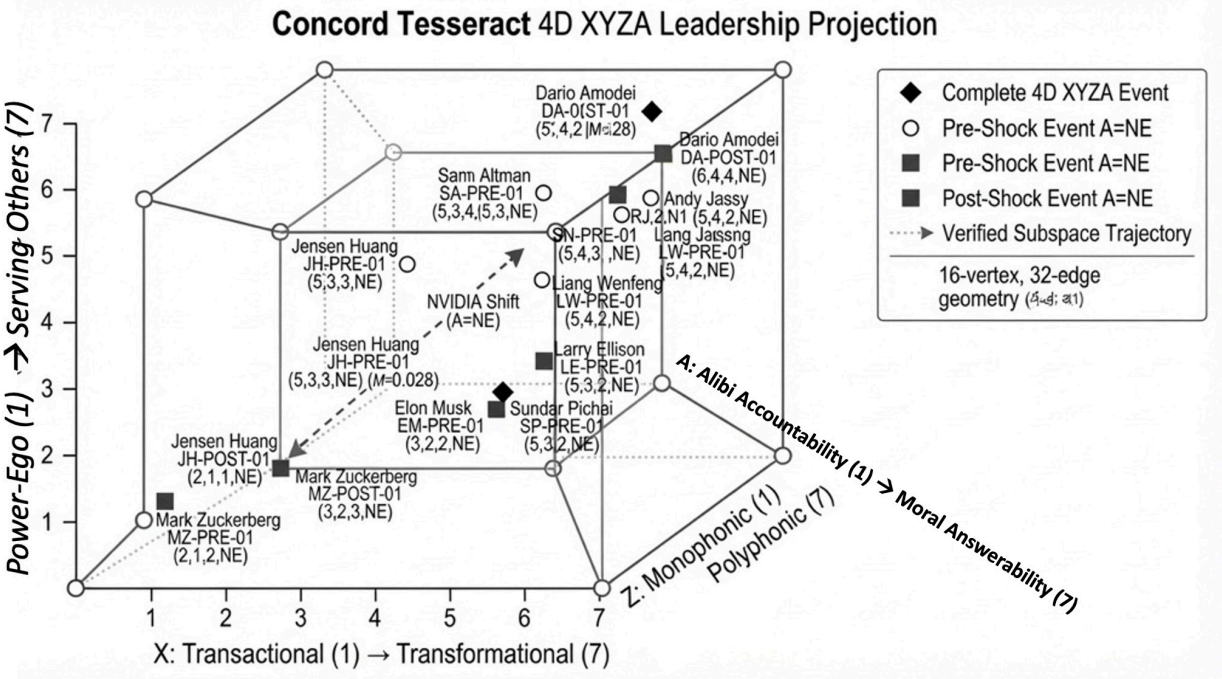



Figure 1 is the operational map readers will encounter again in the leader chapters. The large cube presents provisional X–Y–Z starting coordinates: X ranges from transactional to transformational practice, Y from power or egocentric orientation to serving others, and Z from monophonic to polyphonic participation. Sundar Pichai and Liang Wenfeng share the displayed coordinate (4, 3, 4), so the figure identifies them together at that point. The separate small A cube represents event-specific moral answerability through the three drawing dimensions defined in Appendix A; those dimensions are not extra scores.

## Full XYZA behavioral anchors

Score observable, attributable behavior in a dated event. Use the highest anchor fully supported by the evidence. Preserve the source, counter-reading, actor attribution, and coder confidence. Do not infer motive or private character.

**Table 3: X behavioral anchors: transactional to transformational**

| Score | Anchor | Behavioral definition |
|---|---|---|
| 1 | Transactional exchange | Directions, output targets, compliance, or reward/punishment exchanges dominate; decisions are bounded to immediate task or return. |
| 2 | Compliance-centered | Efficiency and control remain primary; incentives or monitoring are refined, with little change to work design or stakeholder purpose. |
| 3 | Limited development | Task exchange remains central, while some training, consultation, or development is added without changing core processes. |
| 4 | Mixed transaction/development | Performance commitments coexist with a developmental purpose; at least one practice or process is revised and reviewed. |
| 5 | Purposeful redesign | Purpose and development guide material changes to products, work, or governance; stakeholder consequences enter evaluation. |

| Score | Anchor | Behavioral definition |
|---|---|---|
| 6 | Systemic transformation | Operating model or accountability structures are redesigned to build capability and address system-level consequences across boundaries. |
| 7 | Sustained shared transformation | Learning and long-horizon stakeholder consequences shape revisable decisions; shared productive value is built into ongoing practice. |

## Table 4: Behavioral anchors: power/egocentric to serving others

| Score | Anchor | Behavioral definition |
|---|---|---|
| 1 | Power concentrated | Decision rights and benefits are concentrated in the leader or organization; observable costs or risks are shifted to others. |
| 2 | Solidarity without service | Service language appears, but decisions, resources, access, and benefits remain concentrated; no material service action is evidenced. |
| 3 | Consulted, not empowered | Stakeholders are heard but have little effect on decisions and receive no meaningful transfer of access, resources, or authority. |
| 4 | Partial shared service | Some stakeholder input changes a decision and limited resources, access, or protections are provided; important asymmetries remain. |
| 5 | Material service | A documented action materially serves affected others through access, resources, protections, or benefit, with results monitored. |
| 6 | Shared power and burden | Rights or resources are shared; the organization reduces burdens or repairs harms while accepting measurable costs or constraints. |
| 7 | Institutionalized service | Governance sustains stakeholder influence; benefits and burdens are tracked, and inequities are addressed over time. |

## Table 5: Z behavioral anchors: monophonic to polyphonic

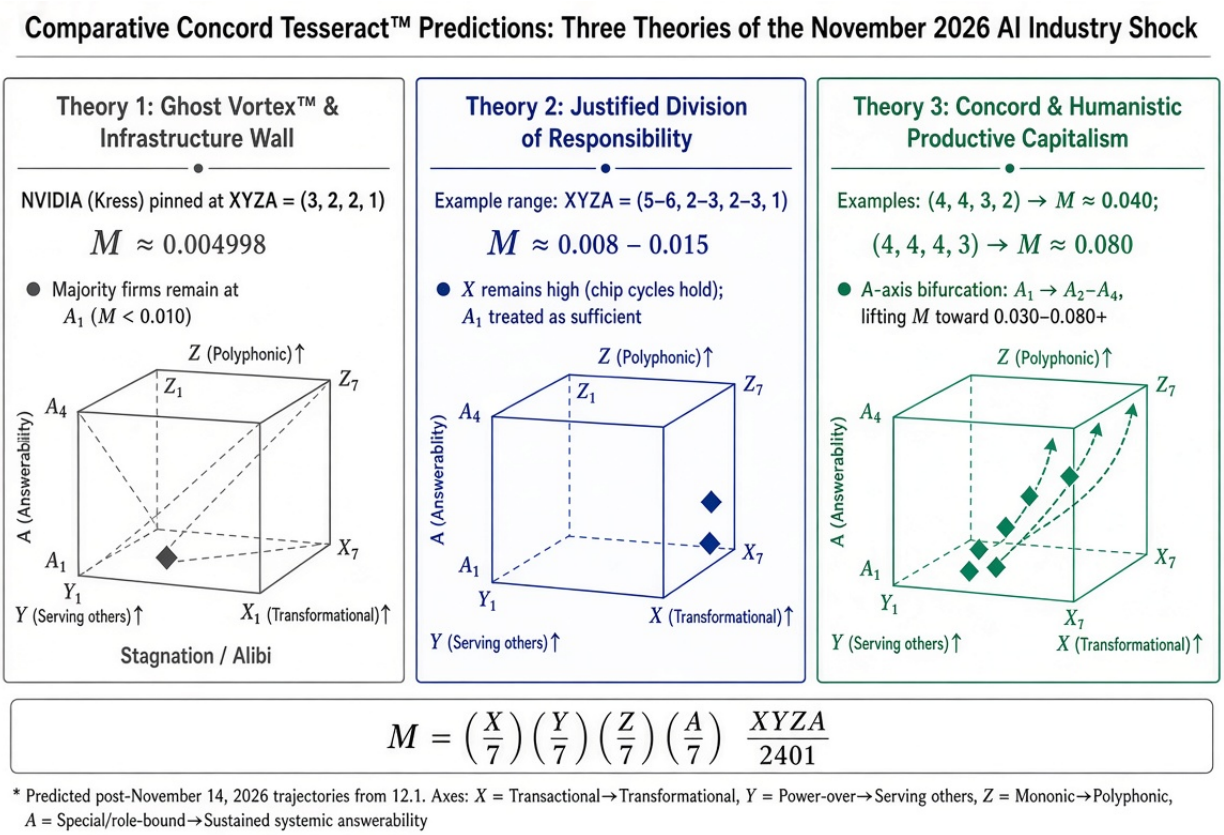
| Score | Anchor | Behavioral definition |
|---|---|---|
| 1 | Single authorized voice | One official account governs; competing accounts are absent from the decision record or actively excluded. |
| 2 | Dissent excluded | Alternative voices may be acknowledged symbolically but do not affect deliberation or decisions. |
| 3 | Filtered plurality | Some perspectives are solicited, then narrowed or filtered; they have little influence on interpretation or decision. |
| 4 | Voices considered | Multiple perspectives are represented and discussed, and some influence is visible, although the dominant account remains controlling. |
| 5 | Voices shape options | Affected people's accounts alter options, priorities, or explanations, and disagreement remains visible. |
| 6 | Dialogic revision | Competing perspectives interact, revise interpretation or decision, and receive follow-up; dissent is preserved rather than erased. |
| 7 | Institutionalized polyphony | Ongoing structures give diverse stakeholders meaningful influence; unresolved tensions remain accessible and can reopen decisions. |

## Table 6: A behavioral anchors: special to moral answerability

| Score | Anchor | Behavioral definition |
|---|---|---|
| 1 | Special / role-bound answerability | Answers for the content of a specific act, decision, duty, or procedure. This is not yet moral answerability for unique participation. A role is an alibi when used to avoid that personal answer. |
| 2 | Non-alibi in Being | Owns unique participation and choice in the event rather than relying on office, market pressure, procedure, or board action as an excuse. |
| 3 | Affected others and consequences | Names particular affected others and concrete consequences; accepts that these consequences require an answer. |
| 4 | Dialogic answerability | Engages affected others' accounts and allows their testimony to challenge or revise the leader's understanding. |
| 5 | Answerable commitment and repair | Commits to a specific response or repair tied to the event, includes affected others in shaping it, and accepts material cost or constraint. |
| 6 | System intervention and verified repair | Implements and verifies a corrective intervention in the relevant organizational or interorganizational system, addressing linked causes rather than only visible symptoms. |
| 7 | Sustained systemic answerability | Sustains answerability over time: affected others can review results, system learning changes practice, and leaders remain open to correction across connected layers. |

A1 is special or role-bound answerability (See Bakhtin references below) for the content of a specific act. It is neither a zero nor a finding of moral answerability. A2 is the first moral anchor: non-alibi ownership of unique participation. A3–A7 require progressively stronger evidence of consequences to affected others, dialogic engagement, commitment and repair, verified system intervention, and sustained correction. The November apology is coded A1 only for its narrowly defined admission/apology event. Do not generalize that event score to Huang's whole post-window profile.

**Figure 2: Graphic Summary of three Bets on the Future Scenarios**




**Interpretation of the Comparative Boje's Concord Leadership Tesseract™ Predictions**

The figure above visualizes the sharply divergent futures that the three rival theories forecast for leadership positioning on the Concord Leadership Tesseract after the November 2026 NVIDIA crisis.

***Theory 1 (Ghost Vortex™ & Infrastructure Wall)***→ predicts structural stasis. Leadership remains trapped in the low-A, low-Y, low-Z region of the tesseract. Colette Kress / NVIDIA is expected to stay pinned near the coordinates (3, 2, 2, 1), producing an extremely low multiplicative profile of \( M \approx 0.005 \). The broader sector is forecast to remain largely stuck at A1 (special/role-bound answerability), with most firms recording \( M < 0.010 \). In visual terms, the single dark marker sits near the bottom-front corner of the cube — a position of organizational alibi, circular capital loops, and

boundary excuses that displace systemic costs onto external stakeholders. Theory 1 treats the crisis as confirmation of a deep governance pathology rather than a temporary personal failure.

***Theory 2 (Justified Division of Responsibility)***→ anticipates continuity of technical excellence inside existing legal and operational boundaries. The X-axis (Transactional → Transformational) remains elevated (X = 5–6) because chip architecture and product cycles are assumed to be robust. However, the A-axis stays at 1; role-bound accountability is regarded as sufficient for a hardware supplier. The resulting M values stay modest (roughly 0.008–0.015). Visually, the markers shift rightward along the X dimension but remain low on the vertical A dimension. Theory 2 frames the November events as an isolated executive breach plus normal market repricing — problems that ordinary board succession can resolve without requiring polyphony, community co-governance, or expanded moral answerability.

***Theory 3 (Concord & Humanistic Productive Capitalism)***→ forecasts a genuine upward bifurcation. External democratic pressure and investor scrutiny are expected to push firms out of the A1 "alibi" zone into the A2–A4 range of non-alibi ownership, dialogic answerability, and verified repair. Supporting rises in Y (serving others) and Z (polyphony) accompany this movement, lifting multiplicative profiles into the 0.030–0.080+ range — a three- to fifteen-fold increase relative to the low-A1 cluster. The green markers and upward arrows in the right panel illustrate this expansion of the decision space. Theory 3 treats the crisis as a *kairos* moment capable of catalyzing True Storytelling organizational development and socio-economic accountability.

**Overall interpretive contrasts**

The above figure is quite preliminary, a demonstration of what can be attemped, and makes the theoretical stakes concrete and falsifiable.

- Theory 1 expects the tesseract to remain compressed near its origin.

- Theory 2 expects a rightward (high-X) but still low-A configuration.

- Theory 3 expects a clear vertical and outward expansion along the non-compensatory A-axis (and supporting Y/Z dimensions).

Because the multiplicative index \( M \) is non-compensatory, even strong performance on X cannot offset persistent weakness on A. Consequently, the three theories generate not merely different narratives but measurably different locations in four-dimensional leadership space. Subsequent earnings calls, SEC filings, board actions, and local siting contracts can therefore be mapped directly onto these predicted trajectories, allowing empirical discrimination among the three accounts of the November 2026 inflection point.

## The Methodological Change in Brief

My revised design has three linked changes.

First, the event episode becomes the core unit of evidence. Each row is tied to a date, a focal person or organization, an event family, a short quotation or documented action, contextual material, and a traceable source.

Second, the evaluation remains dual-layered: XYZA and VuJaDé are related but not merged conceptually.

Third, the coding and verification pipeline becomes staged, so that finding an event, resolving its source, checking context, rating its dimensions, and independently adjudicating its score are distinguishable tasks.

I have in the Spreadsheet the basic XYZA, a proposed VuJaDé sensory-cognitive scheme, and into a broad architecture. It operationalizes VuJaDé as four behavioral dimensions—sensory integration, interactional pacing, process versus outcome, and relational stretch. Therefore, this report centers the current V1–V4 layer and treats the earlier V-KADE text as a superseded working formulation. The labels should not be used interchangeably in chapters or figures.

The index construction is also clarified. For a complete episode with four observed scores, the normalized product is calculated as the product of each score divided by seven. NE is a substantive code for insufficient evidence; it is not a zero and cannot be inserted as a number. Missing cells, unreviewed values, and NE must remain distinct states. An index is reported only when all four coordinates are evaluable.

Finally, the language of verification is tightened. A unique event row is not automatically a verified episode; a score is not independently reliable merely because it is present; and a workbook formula does not validate its inputs. This distinction is central to the reset.

## Why a New Methods Chapter Is Needed

The Routledge Research Monograph I am developing, that original case-by-case approach needs a methodological chapter that explains how the evidence base changes over time. A stable chapter can state the rules once, while each company chapter applies them to a distinct event chronology. Without a shared methods chapter, readers may encounter unexplained changes in time boundaries, source standards, scoring anchors, and the meaning of a pre/post comparison from one corporate case to the next.

Rapidly changing events create a methodological problem as well as a narrative one. A single company profile captured at one point can become obsolete when an executive role changes, a governance process opens, a regulatory or market event alters the stakes, or new primary records become available. The response is not to treat each new event as a complete rewrite of the theory. It is to preserve a versioned event ledger and make the chapter's claims conditional on the data cut-off and verification status.

The methods chapter can also separate the theory contribution from the data-management procedure. The Tesseract is the four-dimensional conceptual projection; XYZA provides the behavioral dimensions; the ledger records event-level evidence; the multiplicative formula summarizes a fully scored episode; and inter-rater procedures examine whether trained readers apply the anchors consistently. Each is necessary, but none is a substitute for the others.

The methods chapter should thus serve as a reproducibility contract: a reader can identify the unit of analysis, inspect the source trail, understand how NE is used, recalculate an index, and see which conclusions are provisional.

## Revision History and Version Logic

The history represented in the attached notes moves from an initial 56-episode pilot to a larger 235-row v0.3 working ledger. The narrative records an intended expansion of NVIDIA and OpenAI materials, consolidation of source references, and an aspiration to scale toward 500 or more episodes. It also records a user request to stop and assess the method before chapters are revised. That is a sound sequence: stabilize the evidence and coding rules first, then revise book prose against a versioned data cut.

The current workbook retains both earlier pilot-coded and later batch-coded status labels. It reports 56 rows as "Pilot Coded (v0.2)" and 179 as "Batch Coded (v0.3)." Those labels should be preserved as provenance labels, but they do not by themselves mean source-verified, context-verified, double-coded, or adjudicated. A future version should record those states in separate fields rather than compressing them into one status.

Use immutable event IDs and versioned workbook releases. Revisions to source URLs, contextual paraphrases, coding anchors, or ratings should be logged with the date, editor, reason, old value, and new value. If a source is invalidated or an event falls outside the window, retain the audit history and mark the record's status; do not silently delete it or reuse its ID.

An "as of" date must appear on the title page, workbook, figures, and chapter tables. The current cut is November 17, 2026. This date identifies the state of the archive, not the date on which all records were independently verified.

## Research Design at a Glance

The design is a longitudinal, documentary, comparative study of public leadership discourse and related governance events in ten AI corporations. It draws on dated public artifacts, including public statements, speeches, investor communications, formal filings, policy texts, and other records that can be cited and revisited. The design does not infer private motives or diagnose psychological traits. It codes observable actions, statements, disclosures, and omissions relative to explicit behavioral anchors.

The core flow is: define the five-year window; identify candidate events; resolve each event to an authoritative source; preserve the source locator and surrounding context; assign event and actor metadata; score XYZA and, where evidence permits, V1–V4; independently code a calibration sample; adjudicate disagreements without erasing defensible interpretations; and release a versioned dataset with a transparent status report.

The two coding layers answer different questions. XYZA concerns how the event is positioned on transformation, service orientation, polyphony, and answerability. VuJaDé concerns what can be observed in the public communicative act: sensory integration, pacing, process transparency, and relational stretch. Some source types support one layer better than the other. A written filing may provide robust evidence about process disclosure but no evidence about vocal pacing.

Comparisons across firms are descriptive until event selection, coverage, and coding reliability are sufficiently balanced. The current distribution is heavily concentrated in two organizations. A dashboard can reveal the ledger's shape, but it cannot by itself establish representative sampling or support unqualified rankings.

## The Unit of Analysis: Event Episode

An event episode is a bounded, dateable public act or decision that can be linked to a source artifact and a focal actor. Examples may include a particular statement in a speech, a defined exchange in a hearing, a specific policy announcement, a governance decision, or a formal filing. The episode should be narrow enough to support a clear coding rationale and broad enough to preserve its meaning in context.

The focal quotation is not the entire episode. It is the short target passage on which the code is centered. The surrounding material records what came immediately before and after, using direct quotation only when the source is available and paraphrase when the report summarizes it. A coding memo explains how the target and context support each score. This prevents a vivid one-line excerpt from carrying claims that belong to another part of the source.

An episode may contain more than one speech act, but multiple claims should not be collapsed simply to increase row counts. If one source contains several distinct acts, create separate rows only when each act has a distinct analytic focus, can be independently scored, and can be traced to its own locator. Record the common source cluster so those rows are not mistaken for independent documents.

The analysis unit remains the event episode even when the book later aggregates evidence into a leader or corporation trajectory. Aggregation is a second-level interpretation and should not be allowed to obscure event-level variation.

## Time Window and Rolling Cut-Off

The current reporting cut-off is November 17, 2026. A literal five-year rolling window ending on that date begins November 17, 2021, inclusive. If the project adopts a different operational rule—for example, beginning January 1, 2021—that choice would create a longer-than-five-year interval and must be described accurately. This report uses the literal rolling-window interpretation for its audit counts.

In the supplied ledger, 215 of 235 dated rows fall on or after November 17, 2021; 20 rows precede that boundary. The earliest date is June 16, 1977, while the ledger title describes a span beginning in 1993. Those pre-window rows may be useful historical background, but they should be flagged as contextual rather than included in the rolling five-year analytic corpus.

Each record needs at least three dates kept conceptually separate: the date of the event or utterance, the date the source was published, and the date the research team accessed or verified the source. The workbook currently includes an Event Date column and a Source ID, but the audit did not find a distinct access/verification date field in the event ledger. Add fields for publication date and last-checked date before the next release.

When an event's exact date is uncertain, do not assign a false precision. Record the best-supported date, indicate its basis, and apply the inclusion rule consistently. Rolling-window eligibility is determined by the event date, not by when an old document is discovered.

## Inclusion and Exclusion Rules

Include an episode when it is publicly attributable, relevant to leadership conduct or organizational governance, bounded to a date or a defensible date range, linked to a retrievable source, and sufficiently contextualized for a reader to understand the target act. The episode should contribute evidence to at least one stated analytic question. An item may be eligible for XYZA while remaining NE on one or more VuJaDé dimensions.

Exclude or quarantine records that cannot be traced to a source, contain a quotation that cannot be located in the source, rely only on an unattributed paraphrase, duplicate another event without adding an analytically distinct act, or fall outside the time window when the analytic corpus is defined as five years. Historical cases can remain in a separate context table.

Do not make inclusion depend on whether a speech act supports a positive or negative interpretation. Candidate selection needs a written rationale and should include constructive, contested, routine, and adverse evidence when it meets the same relevance criteria. Otherwise, the event archive risks becoming a collection of confirming examples.

Each excluded or deferred candidate should have a brief disposition: duplicate, out-of-window, source unavailable, context insufficient, not leadership-relevant, or other specified reason. A disposition log makes selection visible and permits later reconsideration if evidence changes.

## Source Hierarchy and Evidence Types

Primary records should anchor the coding whenever available: official filings, formal statements, transcripts, public speeches, hearing records, policy documents, and direct recordings. A secondary news report can identify an event or contextualize it, but the ledger should state whether the quotation is directly verified in a primary source or is reported by another source. Evidence classes should be visible rather than left to inference.

The workbook has a separate APA 7 Source Registry with nonempty source rows and direct URL fields. The current event rows refer to 194 distinct Source IDs, while the registry contains 196 source entries. These are useful traceability features, but the presence of a URL does not establish that it resolves, that it supports the quotation, or that its metadata are correct. Those are verification tasks.

For each source, preserve author or organization, publication date, title, outlet or document type, URL, access date, and evidence class. If a document is revised online, capture a stable copy or archived locator when permissible, and note the version. If a source is removed or a URL changes, retain the former URL and record the replacement.

Every event row should be traceable from its short quote to an exact page, paragraph, timestamp, transcript line, filing section, or equivalent locator. APA 7 citation and the direct locator serve different functions: one identifies the source; the other helps a reviewer find the relevant evidence.

## Candidate Identification and Sampling

The research team should identify candidate events through a documented search strategy rather than relying solely on model-generated suggestions or memory. Search logs can record the firm, leader, date interval, event family, query terms, search date, and screening decision. The purpose is not to make the archive exhaustive at any cost; it is to make the process sufficiently reproducible and to reduce selective retrieval.

For the immediate reset, the milestone is at least 100 episodes that pass the agreed verification gates. This is not a claim that the workbook's 100 NVIDIA entries, or its 235 total rows, have already passed those gates. It is a minimum verified corpus milestone before the revised empirical claims are treated as stable. Expansion toward 500 or more can continue in batches after the schema and controls are established.

Batch collection is useful when several researchers or tools contribute records. Each batch needs a defined firm, date range, source or event type, and unique ID range. Before merge, compare event IDs, source URLs, quotation text, dates, and source clusters. Cross-batch duplicates should be resolved before totals are reported.

Candidate discovery and verification should remain separate. Discovery is broad and may include promising leads; verification is a conservative process that confirms the event against its source and context. A candidate may remain in a staging sheet without being counted in the verified analytic corpus.

## Context Windows and Quotation Practice

The target excerpt should be accompanied by enough context to establish what it responds to and what follows from it. The revision notes refer to several sentences before and after the target; the event ledger currently stores a short quotation and a context paraphrase rather than consistent verbatim before/after blocks. The next schema should distinguish these forms explicitly.

For each episode, record the exact target passage and its source locator. Then include either (a) two or more sentences before and after when the source format permits and copyright limits allow, or (b) a concise, clearly labeled paraphrase of the surrounding content when reproducing the text is not appropriate. For audiovisual events, include a timestamp and describe whether the transcript is official, machine-generated, or researcher-transcribed.

Context should not be used to smuggle in unverified claims. The focal passage, prior context, following context, and analyst interpretation should be visually or structurally separate. A reader must be able to tell which words belong to the speaker, which are paraphrased, and which are coder commentary.

The quote-length limit used in the workbook is a practical safeguard, not a guarantee of representativeness. A short quote can be coded only when surrounding context is sufficient. If the target is an action rather than speech, store a concise action description and cite the primary record documenting it.

## Source Registry and Traceability

The workbook's Source Registry is a strong structural component of the new design. Event rows carry Source IDs, while the registry holds APA-style references, direct URLs, and evidence classifications. This relational pattern is preferable to repeating full references in every row, provided that the IDs remain stable and every event-to-source link is valid.

The next release should add explicit verification fields: source located, target quote matched, context checked, metadata checked, URL checked, checked by, checked on, and verification note. A single "verified" label is too coarse because it hides which parts of the evidence were actually checked. These fields can be represented as yes/no/not applicable plus a date and initials, or as a controlled staged status.

One source may support multiple episodes. In that case, multiple rows should point to the same source ID and include distinct locators. The source registry should not be duplicated simply because a second event is extracted from the same document. Conversely, a single episode may require multiple sources; the schema should allow a primary source and supporting sources rather than forcing a one-source assumption.

At each release, the workbook should report the number of event rows, unique event IDs, unique source IDs, sources with live URLs checked, records with primary evidence, records with exact locators, and records with verification gaps. These measures make the evidence base more legible than the word "verified" alone.

## Concord Tesseract™: Conceptual Role

The Concord Leadership Tesseract™ is the study's four-dimensional conceptual frame. XYZA gives the behavioral coordinates used to describe each event: X, movement from transactional toward transformational conduct; Y, movement from power-seeking or egocentric orientation toward serving others; Z, movement from monophonic control toward polyphonic inclusion of voices; and A, movement from alibi or evasion toward moral answerability.

Each event receives its own four-coordinate profile. A point in the Tesseract is therefore an episode-level representation, not an intrinsic or permanent score assigned to a person. A leader's trajectory is built from dated events and interpreted in context. A pre/post profile summarizes a defined evidence window and must be presented with the underlying time boundaries and event count.

The Tesseract visualization can show a four-dimensional relation through a three-dimensional projection, with the fourth coordinate represented through a nested cube, color, line style, or other clear visual convention. The figure is a representational aid; the table of coordinates and the formula remain the authoritative data. The geometry should not imply precision beyond the ordinal anchors.

Chapter language should use terms such as "coded episode," "provisional profile," and "observed trajectory" until verification and inter-rater procedures are complete. It should avoid describing ratings as objective measurements of a leader's character.

## XYZA: Behavioral Anchors and Coding

The workbook stores integer scores from 1 to 7 for each of X, Y, Z, and A. The earlier methodology draft describes representative anchors at 1, 4, and 7. For inter-rater work, intermediate scores need explicit behavioral descriptions, not an assumption that coders will interpolate identically. The codebook should specify what evidence distinguishes adjacent ratings across the full seven points.

An anchor describes observable evidence in the episode and its context. It should not rely on an evaluator's general impression of whether a leader is "good," "ethical," "transformational," or "inclusive." For example, a high Z score should require evidence of multiple perspectives being represented or incorporated, not simply an invitation to comment. A high X score should require evidence of consequential change rather than ambitious language alone.

Each rating should have a short rationale linked to the evidence. The rationale should identify the relevant passage or decision, explain why the selected anchor fits, and note plausible alternative readings. A score without a memo is difficult to audit, particularly when the event occurs during a contested transition.

Before formal reliability coding, the team should freeze a codebook version and train coders on examples that are not part of the agreement sample. When an anchor is changed, record the change and recode affected episodes if the change could alter their scores.

## X and Y: Transformation and Service Orientation

X ranges from a largely transactional or administrative response at the low end toward transformational action at the high end. Low scores can describe compliance, routine execution, or limited adjustments. Higher scores require evidence of a substantial shift in organizational practice, capabilities, or institutional arrangements. A large promise does not by itself establish a high score; the episode must show the nature and status of the change.

Y ranges from power-seeking, self-protective, or narrowly extractive conduct toward service to others and stewardship. The score concerns observable commitments, resource choices, accountability practices, or consequences for affected parties. A public statement of concern may be relevant but is not sufficient evidence of service if no action, process, or responsibility is shown.

Both dimensions need context. A policy that benefits one group may impose costs on another; a transformation may increase capability while concentrating control. Coders should explain which evidence is focal and avoid treating growth, scale, or technical achievement as inherently transformational or service-oriented.

For the next full codebook, provide behavioral examples for ratings 1 through 7, decision rules for mixed evidence, and examples of non-evidence. These examples can be drawn from public records after source review and should be anonymized or rotated in training so agreement is not inflated by memorizing the pilot cases.

## Z and A: Polyphony and Answerability

Z ranges from monophonic communication, where one authorized voice dominates and dissent is absent or silenced, toward polyphony, where multiple stakeholder perspectives enter the decision or meaning-making process. Mere consultation is not equivalent to shared influence. Coders should look for who speaks, whose account is represented, how disagreements are handled, and whether input changes the decision or its explanation.

A concerns answerability. The prior draft links the low end with evasion or an alibi for being and the high end with moral answerability, ownership of consequences, situated responsibility, and remedial action. Because the scale carries a specific Bakhtinian theoretical meaning, the codebook should distinguish ordinary legal or administrative accountability from the stronger claim of answerability in a concrete event.

A score must be based on the act and its context. A leader may acknowledge risk without accepting responsibility; accept responsibility without describing remedy; or disclose a remedy without making affected parties part of the process. The codebook should specify how these combinations map to intermediate values rather than collapsing them into a general virtue rating.

The workbook currently assigns A values across all 235 rows, with a large concentration at the low end. This is a feature to review, not a finding to celebrate: it calls for anchor calibration, checks for coder drift, and verification that ratings follow the same evidence rules as the other dimensions.

## A: Evidence, NE, and Theoretical Precision

NE means “not evaluable” because the source or context does not provide adequate evidence for a particular dimension. It is not a rating below 1 and must never be converted to zero. A blank cell means the score has not been entered or is missing; it should not be silently treated as NE. These statuses need separate validation rules in the workbook.

An episode can be evaluable on some XYZA dimensions but not others. For example, a formal statement may show an operational change relevant to X while providing insufficient evidence of stakeholder voice relevant to Z. The coding protocol should permit dimension-specific NE values and explain the evidence threshold for each dimension.

The use of A as an answerability axis must not imply that every episode gives equal opportunity to observe moral responsibility. Some events are technical disclosures; others are explicit responses to harm or challenge. Coders should rate what the event demonstrates, not infer that absence of a moral statement in an unrelated artifact proves a general moral deficiency.

If the study decides that every XYZA dimension is required for a composite, the event-level M_XYZA should be NE whenever any coordinate is NE. The four-dimensional point may still show evaluable coordinates, but the composite should not be calculated from a partial vector unless a separate, pre-registered rule is adopted.

## VuJaDé™: Purpose and Four Dimensions

The attached workbook operationalizes VuJaDé as a public discourse layer with four seven-point behavioral dimensions. V1 is Sensory Integration; V2 is Interactional Pacing; V3 is Process versus Outcome; V4 is Relational Stretch. Its purpose is to code observable communicative performance in an event, not hidden psychological traits and not the same construct as XYZA.

VuJaDé is paired with XYZA because structural conduct and public discourse can move together or diverge. An executive may use clear process language while governance evidence remains incomplete; a formal policy may specify duties even when a speech is rhetorically narrow. A paired event ledger allows such divergence to be examined rather than forced into one score.

The earlier draft's "V-KADE" expansion should be treated as superseded for the current model. It combines modality labels and personality-like categories in a way that does not match the attached workbook's four operational dimensions. Reusing that label would confuse the scale, its formula, and its evidence requirements. Use VuJaDé V1–V4 consistently in the new chapter, worksheets, and figures.

The spelling and trademark convention should be fixed in the manuscript style sheet. Once chosen, use the same form in headings, captions, workbook tab names, and figure labels. The mathematical variable can be written M_VuJaDé or M_Vujade, but one notation should be selected and applied consistently.

## V1 and V2: Sensory Integration and Pacing

V1 ranges from single-channel abstraction or hype at the low end toward a full sensory-relational synthesis at the high end. Its anchors examine how the communication integrates sensory or material frames and whether those frames are grounded in lived reality. A written document can sometimes support coding of sensory integration through its content, but it cannot establish vocal tone, visual staging, or embodied delivery unless the relevant medium is available.

V2 ranges from closed unilateral broadcast or directive at the low end toward sustained deliberative reciprocity at the high end. It concerns interactional pacing and responsiveness: whether questions are controlled or deflected, whether uncertainty is engaged, and whether dialogue continues across turns. A written essay normally does not show live pacing. A transcript may capture turn sequence but may not capture pauses, overlap, or tone; audiovisual evidence may be required for those features.

The present workbook marks V1 as NE in 168 rows and V2 as NE in 169 rows. This pattern is consistent with the notes' caution that static text cannot establish sensory presentation or pacing. It also demonstrates why NE must be treated as a substantive data state and why V1/V2 coverage cannot be assumed from the presence of a quotation.

For each record, the memo should state what medium was reviewed—text, transcript, audio, video, or a combination. It should identify which observable features support the score and which remain unavailable. This makes the score reproducible and prevents modality from being inferred from topic alone.

## V3 and V4: Process and Relational Stretch

V3 ranges from pure outcome hype or inevitability claims toward end-to-end process answerability. The intermediate anchors distinguish outcome-dominant claims, selective process disclosure, balanced explanation, transparent mechanisms, auditable process, and explanation of limits or past errors. A high V3 score should be supported by concrete mechanisms and evidence, not simply detailed technical language.

V4 ranges from an insular or shareholder-only echo chamber toward institutionalized multi-stakeholder attunement. The scale asks which groups are named or considered, whether their burdens and perspectives are translated into the discussion, and whether relational consideration changes the arrangement. Ceremonial mention of "society" or "humanity" is not equivalent to substantive stretch.

The current workbook records V3 and V4 as NE in 89 rows each. This is still substantial missing evaluability. The distinction between NE and a low score matters: an episode may not discuss process or stakeholders at all, or the source may be incomplete. Coders should use the anchor definitions to decide whether silence in a relevant public act is evidence of a low score or whether the dimension is genuinely not assessable.

V3 and V4 can often be scored from written policy or disclosure texts, but each score remains episode-specific. A public statement does not prove that a disclosed process operates in practice. Memos should separate stated process, observed process, and verified implementation.

## The Modality Gate: What Can Be Observed

The new methodology needs an explicit modality gate before V1 and V2 coding. For each event, reviewers should know whether they have only a published text, an official transcript, a third-party transcript, audio, video, or a recording with interaction. Each medium supports a different set of observations. The source type should be stored as metadata, not inferred from the event family.

For example, an official text can support analysis of stated sensory metaphors, process disclosures, or audience references. It cannot establish the speaker's vocal pacing, audience response, or visual environment. A transcript can reveal turn-taking if it faithfully captures interaction but may omit nonverbal and acoustic features. A video can support richer observations, provided the coder records timestamps and limits judgments to observable behavior.

The workbook's V_Sensory_Mode field contains categorical values and descriptive notes. It should be renamed or split so that the code does not conflate a modal category with a description of evidence availability. Recommended fields are: source medium; observed modality tags; V1 evaluability; V1 score; V2 evaluability; V2 score; and a brief evidence memo.

This gate improves construct validity by preventing "not observed" from becoming "low." It also clarifies why a large documentary archive can still produce a smaller evaluable sample for communication features that depend on performance.

## Multiplicative Index: Formula and Meaning

For a complete four-dimensional vector with scores from 1 to 7, the event-level XYZA index is M_XYZA = (X/7) × (Y/7) × (Z/7) × (A/7). The VuJaDé index is M/Vn= (V1/7) × (V2/7) × (V3/7) × (V4/7). The attached ledger implements these products with Excel formulas and returns “NE” when the calculation encounters nonnumeric inputs.

Because the scale starts at 1 rather than 0, the product has a minimum of (1/7)^4, or 1/2401, and a maximum of 1. A single rating of 1 with the other three ratings at 7 yields 1/7, not zero. Therefore, wording such as “a score of 1 collapses the index to zero” is mathematically incorrect. The defensible interpretation is that low values exert a strong multiplicative penalty and cannot be offset by addition.

The raw product is sensitive to the number of dimensions and has a compressed range. That is part of the chosen index definition, but the report should display the formula, bounds, and interpretation plainly. If a geometric mean or another transformation is considered later, it is a methodological change and should be versioned rather than substituted silently.

The index is a summary of the coded coordinates, not a probability, percentage of ethical leadership, or ratio-scale measure unless further assumptions are justified. The ordinal anchors should remain visible alongside the composite in tables and figures.

## NE Handling and Composite Rules

When any required coordinate is NE, the corresponding four-factor multiplicative index is NE. Do not set NE to zero, one, the scale midpoint, or the mean of other dimensions. Such substitutions would create information that the source does not contain and could distort comparisons across text and audiovisual records.

The workbook uses IFERROR formulas to avoid Excel errors when an entry is not numeric. This is a useful technical safeguard, but it does not distinguish all states. A robust formula should return a visible code for NE, a separate code for missing input, and a value only when all four inputs are valid integers from 1 to 7. Out-of-range scores should trigger a validation alert rather than disappear inside IFERROR.

Report coverage alongside the index: number of candidate episodes, number with four evaluable coordinates, number with NE on each dimension, and number excluded from the composite. This prevents readers from comparing a mean based on broad coverage with another based on a narrow, selected subset.

NE rates may themselves reveal source limitations or systematic asymmetry. For instance, written records may support process transparency but not interactional pacing. That pattern should be reported as a property of the evidence base, not repaired by imputing performance scores.

## Aggregation: Event Scores and Leader Trajectories

The event-level product should remain the principal multiplicative score. A leader or corporation summary can include the distribution of event-level M values, dimension-specific distributions, evaluable counts, and a clearly defined time window. If the workbook shows means of X, Y, Z, A, or M across events, these are descriptive summaries of coded observations; they are not a new multiplicative index and should not be labeled as one.

Unequal numbers of events create unequal influence. The present ledger includes 100 NVIDIA episodes and 95 OpenAI episodes but only four to six for several other firms. A corpus-wide average therefore reflects the sampling distribution, not an equal-weight comparison of ten corporations. The report should show organization-specific counts and avoid ranking leaders from raw pooled averages.

Pre/post comparisons require fixed window definitions and comparable evidence rules. The current matrix provides pre-shock and post-window coordinates for ten leaders, but these are summary inputs distinct from the event ledger and should be tied to the rows and evidence windows that justify them. Any manually entered trajectory score should include provenance, calculation method, and reviewer status.

When multiple events are extracted from one source, the analysis should consider source clustering. Several statements from one essay or hearing are not independent documents. For descriptive summaries, note the clustering; for inferential analyses, use appropriate clustered or hierarchical approaches only after the study has a defensible sampling frame.

## Verification Gates: From Candidate to Analytic Record

The revised protocol should use staged statuses. A practical sequence is: candidate identified; source located; event date and actor confirmed; focal quote or action matched; surrounding context captured; source metadata checked; coding memo completed; independent rating complete; disagreement reviewed; final disposition recorded. A record can pass some gates and remain unresolved at others.

The first gate is source authenticity: does the cited source exist and support the claimed event? The second is quotation and context fidelity: does the text match, and is the excerpt interpreted with its neighboring material? The third is analytic eligibility: is the event within scope and relevant to a stated dimension? The fourth is coding reliability: can independent coders apply the anchor consistently? These gates should not be conflated.

The supplied workbook's 235 unique IDs are a useful inventory. Its Adjudication_Status column currently contains only "Pilot Coded (v0.2)" or "Batch Coded (v0.3)," not source-verification outcomes. The notes describe rows as "verified," but the workbook does not provide row-level evidence that all four verification gates were completed. The report therefore does not count all 235 as verified episodes.

The at-least-100 milestone should be met by 100 records that pass predeclared source and context checks and are ready for independent scoring. The exact number can be reported once the gate fields are populated and reviewed.

## Independent Coding Workflow (Stage Two, once the pilot development of theory, and method is completed)

Independent coding begins only after the codebook and episode packets are stable. Each rater receives the same episode identifier, source locator, context window, medium description, and anchor definitions. The packet should not disclose another rater's scores or the aggregate profile being tested. Coders should score each dimension separately and record concise evidence rationales.

A calibration stage should use a small training set that is excluded from the formal reliability estimate. Coders discuss how they interpret anchors, revise ambiguous wording, and then freeze a codebook version. The formal pilot sample is then independently rated without consultation. If a material anchor change occurs, identify affected records and repeat the relevant coding rather than mixing versions.

For four coders, agreement can be assessed pairwise across the six unique coder pairs, or with a multi-rater statistic suited to the data and design. A sheet labeled "Rater_Agreement_Kappa" is not present in the supplied workbook. The event ledger does contain a Rater2_V1_V4 field, but it is blank on all 235 rows. Thus, the current workbook documents no completed second-rater comparison.

After the pilot, the team should report how many episodes were double-coded, how many were adjudicated, which dimensions had the most disagreement, and whether coding instructions changed. Reliability should be tied to a specific codebook version and sample.

## Weighted Cohen's Kappa: Proposed Use

Weighted Cohen's kappa is appropriate for two coders rating an ordinal scale when disagreements of different sizes should receive different penalties. For seven ordered categories, linear weights can be defined as $w_{ij} = 1 - |i - j|/(7 - 1)$. Agreement on the same category receives weight 1; disagreement loses credit in proportion to the distance between categories.

The weighted coefficient is $\kappa_w = (P_o^w - P_e^w)/(1 - P_e^w)$, where $P_o^w$ is the observed weighted agreement and $P_e^w$ is the expected weighted agreement under the marginal distributions. Compute it separately for each dimension and coder pair. With four raters there are six pairwise comparisons per dimension, and the results should be shown as a distribution or table rather than collapsed into one unexplained number.

Kappa is sensitive to prevalence and marginal distributions. A high proportion of low scores or a narrow use of categories can affect the coefficient. Report the contingency tables or marginal distributions, the number of paired observations, the weighting rule, and uncertainty intervals where feasible. A reliability coefficient describes consistency under a protocol; it does not demonstrate that the construct is valid or that the source evidence is true.

Prior notes mention thresholds around .80 or .82 and Landis–Koch interpretive labels. These should remain provisional discussion points until the research team approves a criterion. The coefficient is not a mechanical pass/fail substitute for examining where and why coders disagree.

## Adjudication and Retained Differences

Adjudication follows independent scoring; it does not replace it. When ratings diverge beyond a pre-agreed threshold, reviewers return to the source and the behavioral anchor, compare rationales, and record the basis for the final disposition. A threshold such as a difference greater than two points may be useful for triage, but it should be applied consistently and documented before the formal sample is examined.

Not every difference has a single indisputable answer. A speech act can support competing readings because it addresses multiple audiences, contains qualified claims, or omits relevant evidence. When disagreement remains defensible, retain both interpretations in an adjudication record. A consensus score may be reported for a specific analysis if the team defines how it was reached, but the rival ratings should not be erased.

The notes refer to preserving “Retained Differences” as an ethical and interpretive practice. The supplied workbook has an Adjudication_Status field, but current statuses are batch labels rather than adjudication outcomes. The next version should separate coding status, verification status, reliability status, and adjudication outcome into distinct columns.

An adjudication memo should include episode ID, affected dimension, each original score, coder rationale, source passages reviewed, discussion outcome, any anchor clarification, final score if applicable, and whether the disagreement is retained. This record is part of the evidence trail.

## Workbook Architecture: Current Design

The attached workbook contains five worksheets: 10_Leader_XYZA_Vujade_Matrix; Corpus_Analytics_v03; Event_Ledger_XYZA_Vujade; APA7_Source_Registry; and Vujade_1_7_Codebook. This is a useful relational architecture: a pre/post matrix, a descriptive dashboard, an episode ledger, a source registry, and a four-dimensional Vujade codebook are separated into distinct functions.

The Event Ledger has 24 columns, including event ID, event date, company, focal leader, event family, short quotation, context paraphrase, source ID, XYZA scores and product, Vujade modality and V1–V4 fields, the VuJaDé product, a layer gap, coupled index, coding memo, second-rater field, and status. The structure is considerably more refined than the initial short methods text.

The ledger contains formulas in all 235 rows for M_XYZA, M_Vujade, XYZA_Vujade_Gap, and Coupled_Index. Those formulas provide a useful computational scaffold. Their presence does not validate the inputs or establish agreement. The dashboard's live COUNTIF and AVERAGEIF formulas should also be checked against the exact company labels and intended data window before figures are exported.

The Source Registry has 196 populated source entries, while event rows point to 194 distinct source IDs. This difference may reflect unused or reserved sources; it should be reconciled and explained, not silently removed.

## Workbook Audit: Snapshot Findings

The workbook's event ledger contains 235 rows with unique event IDs. Company counts are: NVIDIA 100; OpenAI 95; Microsoft 6; Anthropic 5; xAI 5; Oracle 5; Meta 5; Google/Alphabet 5; DeepSeek 5; and Amazon 4. Dates range from 1977-06-16 to 2026-11-17. Of these, 215 fall on or after 2021-11-17 and 20 are earlier.

All 235 rows contain numeric XYZA scores and formulas for the XYZA product. The VuJaDé fields are populated with either an integer or NE; V1 is NE for 168 rows, V2 for 169, V3 for 89, and V4 for 89. These numbers describe the current coding file and should be recalculated if any row changes.

The status field contains 56 "Pilot Coded (v0.2)" and 179 "Batch Coded (v0.3)" labels. The Rater2_V1_V4 field is blank on all rows. The workbook has no separate inter-rater agreement worksheet. Therefore, the supplied workbook does not support a claim that 235 episodes have been verified through dual coding or that weighted kappa has been computed.

The workbook title and dashboard use "verified" language, and the revision notes describe 235 as verified. That wording should be revised until row-level checks and independent reliability data are documented. A safer current label is "expanded coded working ledger, v0.3."

## Audit Implications and Required Reconciliation

Four inconsistencies require resolution before the workbook is used as a final analytic dataset. First, the five-year-window specification conflicts with the ledger's historical span. Second, the narrative notes call the rows verified while the workbook records batch/pilot coding and blank second-rater fields. Third, the notes mention a kappa tab that is absent from the attached workbook. Fourth, the earlier V-KADE proposal differs from the actual V1–V4 VuJaDé codebook.

These are ordinary version-control issues in a changing project, not grounds for discarding the work. They do mean that chapter prose must not import status statements from an earlier assistant-generated summary without checking the current workbook. The methods chapter should state which file version governs each table and what its labels mean.

The next workbook release should include an "Audit and Data Dictionary" tab listing version, cut-off date, time-window rule, status definitions, formula definitions, known limitations, and change log. A "Verification" tab should track the staged source and context checks. A "Rater Agreement" tab should be added only when actual independent ratings are available; it should not contain sample or illustrative values that could be mistaken for results.

Before producing public figures, reconcile company name spellings, date formats, source IDs, dashboard ranges, formulas, and pre/post window definitions. Save a frozen version used for each chapter analysis.

## From 235 Rows to a Verified 100-Episode Milestone

The project has a large enough candidate archive to support a disciplined verification sprint, but the count of 235 should not be mistaken for the target of 100 verified episodes. The first step is to apply the rolling-window rule and mark the 20 earlier records as historical context or exclude them from the analytic count. The remaining 215 are candidates for source and context review, not automatically verified records.

Prioritize breadth and auditability rather than simply confirming the largest batches first. The 100-episode milestone should include defined representation across organizations, years, and event families to the extent the evidence permits. If the distribution remains uneven, report it and limit comparative conclusions accordingly. Do not fill quotas with weak sources merely to reach a count.

For each candidate, confirm the source, exact event date, focal actor, target passage, context, APA metadata, direct locator, and eligibility. Then assign separate verification fields and a verified date. A second reviewer can audit a defined subset or the full milestone, depending on the agreed reliability plan. The criteria for calling a record “verified” should be written before review begins.

Once the first 100 pass the gates, freeze a release and use it as the basis for the methodology chapter’s example tables. Continue to add events in controlled batches toward the larger 500-plus ambition, with every later release preserving IDs and a change log.

## Implications for Chapter 2 and Later Cases

Chapter 2 can demonstrate the method in depth because NVIDIA is the largest current company sub-ledger. It should show the event sampling rule, a short evidence table, several fully contextualized examples, the XYZA and VuJaDé code rationales, and a transparent distinction between the exploratory profile and any independently verified result. A dense row count alone should not determine the chapter's conclusions.

Chapter 3 and subsequent case chapters can reuse the shared methods chapter and apply the same event packet, codebook version, verification gates, and cut-off rule. They should not introduce new scale definitions without recording a version change. Any firm-specific exception—such as a different source medium or a leadership transition—should be described as a case condition rather than a hidden change to the general method.

The manuscript sequence described in the notes runs from Chapter 2 through Chapters 3–11 and the conclusion. The methods chapter can support that sequence by defining a common table template: episode ID and date; source type; context; XYZA vector; VuJaDé vector; NE fields; index; verification state; and analytic memo.

When new data alter an earlier trajectory, revise the evidence window and report what changed. Do not merely overwrite the previous figure. Preserve the earlier state as a dated version so the book can explain why a case analysis changed.

## Tables, Figures, and the Tesseract Projection

The core figure should show that each event is represented by a four-coordinate XYZA vector and that the vector is projected into a Tesseract-style visual form. The underlying score table must display the four coordinates and the formula result. A three-dimensional rendering should identify how the fourth coordinate is encoded and include a legend; no perspective drawing should imply that a viewer can directly see four independent spatial axes.

A second figure can compare XYZA with VuJaDé while preserving their separate constructs. A paired panel or two aligned plots is clearer than placing eight dimensions into one crowded cube. A contrast line or gap measure should be used only after its calculation and interpretation are defined. The workbook currently has a `XYZA_Vujade_Gap` formula and a `Coupled_Index` formula; these require a conceptual explanation before publication.

Pre/post trajectory figures should show the dates or event windows behind each point and identify the evidence count. If the matrix coordinates are judgment-based summaries rather than computed aggregates from ledger rows, label them accordingly. If they are calculated from data, show the algorithm and treatment of NE values.

Figures should be generated from frozen workbook values and cross-checked against the table. Every visual should include the dataset version and cut-off date. The codebook and data table, not the illustration prompt, govern the score.

## Quality Assurance and Data Governance

Quality assurance has three layers: evidence integrity, coding integrity, and computational integrity. Evidence integrity asks whether a record is real, attributable, dated, and supported by its source. Coding integrity asks whether anchors are explicit, ratings are reproducible, and disagreement is retained or adjudicated transparently. Computational integrity asks whether formulas reference the correct cells and handle NE and missing values as specified.

Basic workbook checks should flag duplicate event IDs, duplicate event/locator combinations, missing source IDs, invalid dates, out-of-range scores, unsupported company labels, formulas replaced by values, unexpected blanks, and mismatches between the ledger and dashboard. The checks should be rerun for each release and their outcomes logged.

Protect the source trail. Keep the original source URL, any archived version, access date, and local citation metadata. Do not store only a pasted quote. If a transcript is corrected or updated, note the source version used for coding. For audiovisual evidence, record the timestamp and whether a transcript was checked against the recording.

Use a change log with version number, date, responsible editor, affected IDs, and explanation. A correction to a source locator is different from a change to a rating anchor; both should be visible. This governance makes the workbook suitable for incremental expansion.

## Limitations and Interpretive Safeguards

The archive is purposive and currently highly imbalanced across companies. The report does not treat it as a representative sample of all AI leadership discourse. Conclusions should be limited to the records and windows actually coded, with the selection process and missing coverage reported.

Ordinal ratings summarize judgments against behavioral anchors. Multiplicative indices impose a specific mathematical model and do not turn those judgments into objective physical quantities. The product is sensitive to low scores and to NE coverage; interpretation must stay close to the underlying dimensions and evidence.

Public communication is not the same as organizational practice. VuJaDé measures observable discourse, while XYZA may code speech acts and governance records relevant to institutional conduct. Neither alone proves implementation or internal decision processes. Triangulation across different public sources can improve the evidence base but cannot eliminate all ambiguity.

Historical material and events outside the five-year window may provide context but should not be mixed into the rolling analytic corpus without a declared rule. Rapidly changing events also mean that “current profile” is always an archive cut, not a permanent judgment.

## Claims the Report Can and Cannot Make Now

The workbook supports these present-tense descriptive claims: it contains 235 unique event rows; it has the stated company distribution and date span; it includes event-level XYZA and VuJaDé formula columns; it includes a Source Registry and a seven-point VuJaDé codebook; and it currently has no populated second-rater field or separate kappa worksheet. These claims are directly checkable from the file.

The workbook does not yet support claims that all 235 events are independently verified, that all references and URLs have been checked, that weighted Cohen's kappa meets a threshold, that four-rater agreement is established, or that the ten-company comparison is balanced. It also does not substantiate a causal claim that a leader's score changed because of a particular shock.

Once the verification and reliability gates are complete, the report can state the actual number of verified episodes, the double-coded sample size, the agreement results by dimension, and the number of episodes adjudicated. Each claim should cite the workbook release and codebook version.

This distinction is not a retreat from the theory. It makes the theory testable and the empirical record credible by identifying what has been done, what remains provisional, and what would count as stronger evidence.

## Implementation Sequence for the Next Release

1. Freeze the current v0.3 workbook and record its checksum or archived copy. 2. Confirm the exact five-year window and label historical records. 3. Reconcile source IDs, company labels, event dates, and dashboard ranges. 4. Add fields for medium, source locator, publication date, access date, verification gates, and separate coder scores.

5. Finalize the current XYZA and VuJaDé codebooks, including all seven behavioral anchors and rules for mixed evidence. 6. Select a calibration set, train coders on non-study examples, and then conduct blind independent coding on the formal pilot. 7. Calculate agreement using declared pairwise weighted kappa procedures and inspect disagreements dimension by dimension. 8. Adjudicate while retaining defensible differences.

9. Mark at least 100 episodes as verified only after the stated source and context checks pass. 10. Release a new workbook version with corrected titles, formulas, coverage summaries, and a change log. 11. Draft the standalone methodology chapter from the frozen version. 12. Revise Chapter 2 and subsequent cases using the same rules and clearly identified cut-offs.

This is a staged sequence. It avoids simultaneously changing the corpus, codebook, formulas, and chapter interpretation without a record of which change produced which result.

## Decisions to Resolve Before Reliability Coding

The analysis and coding team should settle the exact operational window; whether out-of-window records remain in a separate historical table; which source types qualify as primary evidence; what minimum context packet is required; and which verification states must be complete before an event counts toward the 100-episode milestone.

For scoring, confirm the final wording and trademark style for VuJaDé, whether V1 can be evaluated from written content as well as observed modality, which source media are necessary for V2, and whether NE on one axis makes the full composite NE. Confirm whether A's upper anchors require evidence of remedy, answerability, or both, and how mixed evidence is resolved.

For reliability, decide the number of raters, the formal pilot sample, the pairwise statistic, the weighting matrix, how confidence intervals will be estimated, what disagreement threshold triggers adjudication, and whether a numerical threshold is a goal or a gate. Avoid importing a threshold from a prior draft without contextual justification.

For analysis, decide how the pre/post matrix is derived, how source clustering is handled, whether the goal is a purposive archive or a balanced comparison, and which aggregate summaries are descriptive only. Recording these decisions before full coding reduces retrospective flexibility.

## Appendix A: Compact XYZA Codebook Framework

X — Transactional to Transformational. Score the observable scale of change in the event, from routine compliance or limited adjustment toward consequential redesign or transformation. Require evidence of action, status, or mechanism; do not score aspiration as achievement.

Y — Power-seeking to Serving Others. Score the observable orientation toward beneficiaries, burdens, resource choices, and stewardship. A high score requires more than benevolent language; identify whose interests are served and what conduct supports that interpretation.

Z — Monophonic to Polyphonic. Score the range and influence of voices represented in the event. Distinguish unilateral announcement, selective consultation, and meaningful integration of multiple perspectives. Note whether dissent is heard and whether input affects the process.

A — Alibi/Evasion to Moral Answerability. Score the event's ownership of consequences, situated responsibility, and response or remedy. Distinguish formal compliance from answerability. Do not infer a global trait from one episode.

For every dimension, maintain anchors at 1–7, exemplar passages, boundary cases, non-evidence examples, and an NE rule. The earlier methodology draft supplied only selected anchor points for XYZA. Full inter-rater use requires all seven anchors to be behaviorally described and version-controlled.

## Appendix B: Compact VuJaDé Codebook Framework

V1 — Sensory Integration. 1: single-channel abstraction or metrics-only hype. 2: narrow modal framing with boilerplate acknowledgment. 3: paired registers without material action. 4: balanced but scripted modal presence. 5: coherent cross-modal framing. 6: embodied and material grounding in lived reality. 7: full sensory-relational synthesis.

V2 — Interactional Pacing. 1: closed unilateral broadcast. 2: controlled Q&A or rapid deflection. 3: selective response to easy questions. 4: structured formal deliberation. 5: reflective engagement with uncertainty. 6: dialogic responsiveness to critique. 7: sustained deliberative reciprocity.

V3 — Process versus Outcome. 1: pure outcome hype or inevitability. 2: outcome dominance with procedural alibi. 3: selective process disclosure. 4: balanced process and outcome account. 5: transparent mechanism and limits. 6: auditable process and tracing of past errors. 7: end-to-end process answerability.

V4 — Relational Stretch. 1: insular or shareholder-only framing. 2: ceremonial mention of broad groups. 3: segmented stakeholder talk. 4: partial cross-boundary translation. 5: active perspective-taking about burdens. 6: relational stretch connected to contractual change. 7: institutionalized multi-stakeholder attunement. These anchors are transcribed from the supplied workbook codebook; refinements should be versioned.

## Appendix C: Recommended Event-Ledger Fields

Identification: stable event ID; event date; publication date; company; focal leader(s); role at event; event family; within-window flag; batch ID; source-cluster ID. Evidence: exact target passage or action; before-context; after-context; transcript status; medium; timestamp or page/paragraph locator; primary source ID; supporting source IDs; APA 7 reference; URL; access date.

Coding: codebook version; XYZA X, Y, Z, A; rationale by dimension; VuJaDé V1, V2, V3, V4; rationale by dimension; NE reason by dimension; M_XYZA; M_VuJaDé; calculation status; optional descriptive contrast field with definition. Do not use one general memo to obscure which evidence supports which dimension.

Quality: source located; quote matched; date confirmed; context captured; reference checked; URL checked; verifier and date; coder 1 ratings; coder 2 ratings; coder 3/4 ratings if applicable; agreement statistic and sample identifier; adjudication status; retained-difference note; final disposition; version-change note.

Controlled values should be validated with drop-down lists and formula checks. Missing, NE, pending, excluded, and not applicable should be distinct. The data dictionary should define every field, accepted values, and whether it is manually entered or calculated.

## Appendix D: Formula and Reporting Specifications

For a complete four-coordinate XYZA event: M_XYZA = (X/7)(Y/7)(Z/7)(A/7). For a complete four-coordinate VuJaDé event: M_VuJaDé = (V1/7)(V2/7)(V3/7)(V4/7). Inputs must be integers in the closed interval 1–7. If any required score is NE, report the index as NE; if any score is blank or invalid, report a data-quality status rather than a numerical index.

If `Gap = M_XYZA − M_VuJaDé` is retained, define direction explicitly: positive values mean the structural index exceeds the discourse index under the chosen coding and the reverse for negative values. Since each layer may have different evaluability, report only cases where both indices are calculable and provide the matched N.

If `Coupled = sqrt(M_XYZA × M_VuJaDé)` is retained, explain why this geometric coupling is useful and how it behaves under missingness. It is not interchangeable with either layer's index. Do not use it as a substitute for showing the XYZA and VuJaDé vectors.

Every published table should state N of candidate rows, N source-verified, N fully context-checked, N double-coded, N with a calculable index, and the time window. Report the codebook and workbook version in the caption.

## Selected References and Source Notes

The earlier methodology draft and workbook identify the following methodological works for the report's framing. Bibliographic details should be checked against the authoritative records before final book submission; this report has not independently verified the source registry or every URL.

Cohen, J. (1968). Weighted kappa: Nominal scale agreement with provision for scaled disagreement or partial credit. Psychological Bulletin, 70(4), 213–220. Landis, J. R., & Koch, G. G. (1977). The measurement of observer agreement for categorical data. Biometrics, 33(1), 159–174.

St Quinton, T., & Trafimow, D. (2025). The need to consider auxiliary assumptions in preregistration practices. Frontiers in Psychology, 16, Article 1633990. Trafimow, D. (2022). Generalizing across auxiliary, statistical, and inferential assumptions. Journal for the Theory of Social Behaviour, 52(1), 37–48.

Trafimow, D. (2023). A new way to think about internal and external validity. Perspectives on Psychological Science, 18(5), 1028–1046. Trafimow, D., St Quinton, T., & Weber, M. (2024). Crediting or blaming theories or auxiliary assumptions for empirical victories or defeats. Theory & Psychology, 34(4), 508–530. These citations inform the methodological discussion; they do not establish that the present dataset has passed reliability or validity tests.

## Closing: A Method That Can Be Revisited

The methodological reset is best understood as a move from a promising pilot architecture to a versioned research system. XYZA and VuJaDé provide complementary lenses; the episode ledger makes evidence dateable and traceable; multiplicative formulas preserve the non-compensatory logic; verification gates distinguish candidate material from confirmed records; and inter-rater procedures test whether the anchors are applied consistently.

The current workbook is a substantial working resource. It already includes a broad event ledger, a source registry, two scoring layers, formula columns, a pre/post matrix, and a seven-point VuJaDé codebook. The audit also shows why a methods chapter is needed now: the time-window boundary, verification language, coder agreement, version status, and earlier V-KADE terminology must be reconciled before the expanded data are presented as final.

The next defensible claim is not that all 235 events are verified. It is that the project has an expanded 235-row coded ledger and a clear, auditable path to a verified 100-episode milestone, followed by staged expansion. Once the first release passes those gates, the book can present its analyses with clearer boundaries and a stronger account of how the evidence was made.

This protocol leaves the research open to correction. New evidence can change an event's interpretation; a revised anchor can change a score; and an independent coder can reveal a disagreement worth preserving. Versioning those changes is part of the method and part of the contribution.

## References


Bakhtin, M. M. (1993). Toward a philosophy of the act (V. Liapunov, Trans.; V. Liapunov & M. Holquist, Eds.). University of Texas Press.

Boje, D. M. (1995). Stories of the storytelling organization: A postmodern analysis of Disney as "Tamara-land." Academy of Management Journal, 38(4), 997–1035.

Boje, D. M. (2001). Narrative methods for organizational and communication research. Sage.

Boje, D. M., & Fatemi, F. (2026, in press). Tamaraland meets tesseract leadership of AI corporations: Ghost Vortex, Consent Machine, and the dark side of leadership. Tamara Journal of Critical Organization Inquiry.

Boje, D. M., Hillon, Y. C., & Mele, T. M. (2017). 21st century university and the failure of business process reengineering. Organization Development Journal, Vol. 35(1), Spring: 91-106. (Spring 2017): 91-106. Pre-press version: https://davidboje.com/vita/paper_pdfs/21st%20Century%20University%20and%20BPR%20Failure_BHM.pdf

Boje, D. M., & Sanchez, M. (Eds.). (2018). The Emerald handbook of quantum storytelling consulting. Emerald Publishing. https://doi.org/10.1108/9781786356710

Boje, D. M., & Saylors, R. (2023). The management thought of Louis R. Pondy: Reclaiming the enthinkment path. Routledge. Confirm publication year and subtitle against the copy used in the manuscript.

Boje, D. M., & Rosile, G. A. (2027). The dark side of leadership storytelling: Betraying stakeholder trust. Advance-copy chapter, pp. 11–25. Complete parent-book and publisher metadata from the advance-copy title page.

Boje, D. M., Rosile, G. A., Dennehy, R., & Summers, D. J. (1997). Restorying reengineering: Some deconstructions and postmodern alternatives. Communication Research, 24(6), 631-668.

Boje, D. M., & Saylors, R. (2024). The management thought of Louis R. Pondy: Reclaiming the Enthinkment path. Routledge.

Çoban, C. (2022). The dark side of leadership: A conceptual assessment of toxic leadership. Business Economics and Management Research Journal, 5(1), 50–61.

Conger, J. A. (1990). The dark side of leadership. Organizational Dynamics, 19(2), 44–55.

DeepSeek. (2026). Transparency Center. https://deepseek.com/en/transparency/

Fatemi, F., & Boje, D. M. (2026a). Ghost Vortex and the non-alibi in being: Dark-side leadership, Bakhtinian moral answerability, and the existential stakes of artificial intelligence. Journal of Philosophical Investigations, 20(56), 19–34. https://doi.org/10.22034/jpiut.2026.75002.4784

Fatemi, F., & Boje, D. M. (2026b). Life as narrative: Time, answerability, and the unfinalizable self. Journal of Philosophical Investigations, 20(56), 783–798. https://doi.org/10.22034/jpiut.2026.75003.4785

NVIDIA. (n.d.). Chris A. Malachowsky: Founder and NVIDIA fellow. https://nvidianews.nvidia.com/bios/chris-a-malachowsky

NVIDIA. (n.d.). Jensen Huang: Co-founder, president and chief executive officer. https://www.nvidia.com/en-us/about-nvidia/board-of-directors/jensen-huang/

NVIDIA. (n.d.). Our history: Innovations over the years. https://www.nvidia.com/en-us/about-nvidia/corporate-timeline/

NVIDIA. (n.d.-a). Chris A. Malachowsky: Founder and NVIDIA fellow. Retrieved November 5, 2026, from https://nvidianews.nvidia.com/bios/chris-a-malachowsky

NVIDIA. (n.d.-b). Jensen Huang: Co-founder, president and chief executive officer. Retrieved November 5, 2026, from https://www.nvidia.com/en-us/about-nvidia/board-of-directors/jensen-huang/

NVIDIA. (n.d.-c). Our history: Innovations over the years. Retrieved November 5, 2026, from https://www.nvidia.com/en-us/about-nvidia/corporate-timeline/

NVIDIA Corporation. (2002). Annual report on Form 10-K [Fiscal-year identification to confirm against the filing]. https://media.corporate-ir.net/media_files/NSD/NVDA/reports/10k02.pdf

NVIDIA Corporation. (2002). Annual report on Form 10-K. https://media.corporate-ir.net/media_files/NSD/NVDA/reports/10k02.pdf

NVIDIA. (2026, November 12). NVIDIA announces CEO transition. https://nvidianews.nvidia.com/news/nvidia-announces-ceo-transition

Securities and Exchange Commission. (2026, November 12). NVIDIA Corporation Form 8-K (event date November 11, 2026). https://www.sec.gov/Archives/edgar/data/1045810/000104581026000082/nvda-20261111.htm

Trafimow, D. (2023). Methodological issues in psychology: Concept, method, and measurement. Routledge. https://www.routledge.com/Methodological-Issues-in-Psychology-Concept-Method-and-Measurement/Trafimow/p/book/9781032429786

Savall, H. (2018). Radical origins to economic crises. Springer International Publishing. ISBN 978-3-319-93365-8 ISBN 978-3-319-93366-5 (eBook) https://doi.org/10.1007/978-3-319-93366-5

Savall, H., Péron, M., Zardet, V., & Bonnet, M. (2018). Socially responsible capitalism and management. Routledge.

Savall, H., & Zardet, V. (2011). The qualimetrics approach: Observing the complex object. Information Age Publishing.

Savall, H., & Zardet, V. (2023). Socio-economic approach to management treatise: Theory and practices. Information Age Publishing.

Trafimow, D. (2012). The role of auxiliary assumptions for the validity of manipulations and measures. Theory & Psychology, 22(4). https://doi.org/10.1177/0959354311429996

Trafimow, D. (2022). Generalizing across auxiliary, statistical, and inferential assumptions. Journal for the Theory of Social Behaviour, 52(1), 37–48. https://doi.org/10.1111/jtsb.12296

Trafimow, D. (2023). A new way to think about internal and external validity. Perspectives on Psychological Science, 18(5), 1028–1046. https://doi.org/10.1177/17456916221136117

Trafimow, D. (2026). The unappreciated theoretical implications of multiplication versus addition. In-press chapter manuscript, World Scientific Encyclopedia Series 3, Volume 1. Advanced pre-press pdf https://davidboje.com/ws/Multiplication-AI.pdf

Trafimow, D. (2026, accepted). The unappreciated theoretical implications of multiplication versus addition. Chapter 1, Volume 2, 'Dark Side of Leadership in the AI Revolution.

Trafimow, D., St Quinton, T., & Weber, M. (2024). Crediting or blaming theories or auxiliary assumptions for empirical victories or defeats. Theory & Psychology, 34(4), 508–530. https://doi.org/10.1177/09593543241253548