Chapter 2. The Chokepoint Leader: Jensen Huang, NVIDIA, and Moral Answerability at the Infrastructure Layer

David Michael Boje Revised Nov. 14, 2026

Abstract

I trace NVIDIA's 33-year development from a 1993 gaming-graphics venture to a central supplier of accelerated computing and AI infrastructure, while recognizing that gaming remains an important business. I examine Jensen Huang's leadership through the company's technical history, its financing and supply relationships, U.S.–China export-control policy, and the November 2026 leadership crisis. The reporting on Huang's relationships and the espionage suspicions surrounding a business executive connected to an NVIDIA customer must be distinguished from official corporate actions and from findings that have not been made. I use the WAVE event archive and behaviorally anchored XYZA scales to code dated episodes and calculate the multiplicative M index. This chapter keeps leader behavior, company and network events, financial indicators, and political-economic context visible as related but distinct forms of evidence. The analysis is provisional and open to falsification.

Keywords: NVIDIA; Jensen Huang; Colette Kress; AI infrastructure; export controls; moral answerability; XYZA; WAVE qualimetrics; Fourth Wave Grounded Theory; Peirce.

1. Why I begin at the infrastructure layer

In 1978, as I was completing my University of Illinois dissertation on centrality in interorganizational networks, I encountered the university's computing infrastructure in a very immediate way. I used its computer system to generate the printed dissertation rather than rely on a conventional typewriter. The dissertation was formally filed in 1979 (Boje, 1979). David A. Whetten, my dissertation chair, and I subsequently published research on organizational strategies, contextual constraints, centrality, and attributed influence (Boje & Whetten, 1981). The experience of producing a dissertation through a computing infrastructure and studying interorganizational networks at the same time has stayed with me.

I return to that research when I examine NVIDIA. Semiconductor design, fabrication, lithography, advanced packaging, systems, data centers, and model training involve different organizations and different kinds of power. NVIDIA's position in these networks raises a question I first studied in another technological era: how does network centrality become attributed influence? Centrality alone cannot establish responsibility for every downstream act; the causal links, contractual relations, and opportunities for corrective action must be investigated. This chapter asks where answerability travels through the network—and where it stops.

I begin behind the apparently immaterial AI interface. I see silicon, lithography, semiconductor fabrication, energy contracts, data-center buildings, workers, water, local

government, and communities. Each interface answer depends on an organization of material life that is not visible in the answer itself. This is why I begin with a leader whose products help make other leaders' AI products possible. Jensen Huang is not simply another name in a list of technology executives. NVIDIA's place in the computing stack makes the company a strategic site for asking how technological capability and answerability travel—or fail to travel—through organizational boundaries.

I call Huang the Chokepoint Leader and, in my earlier vocabulary, the Crystal Palace Architect. These are interpretive metaphors, not charges of unlawful conduct. A chokepoint is a narrow passage through which many projects must pass; the metaphor directs attention to dependency and bargaining power. The Crystal Palace names the beauty and coherence of an engineered system whose very elegance can make its excluded consequences harder to see. I ask what happens when a technically accurate account of performance per watt becomes the only available answer to questions about absolute energy use, community consent, or the distribution of benefits.

My thesis is intentionally risky: infrastructure suppliers may help produce the conditions of downstream harm while formal descriptions of responsibility place those consequences outside their own organizational frame. The competing thesis is just as important. Supplier and operator responsibilities may genuinely differ, and responsible suppliers may improve efficiency, share information, and cooperate with remedy within their control. The existence of NVIDIA chips at a facility does not by itself establish NVIDIA's control over a gas turbine, a water permit, or a utility board. My argument must survive that distinction.

I use first person because I am not hiding behind an impersonal research machine. I choose the questions, draw the initial map, decide which counter-stories deserve attention, and remain answerable for my interpretations. AI collaborators can propose other readings, but their fluency does not convert conjecture into observation.[1]

[1] Source protocol for this biographical section. I treat the Computer History Museum oral history reportedly recorded on August 9, 2024, as a priority primary-source lead, but its exact catalog record, interview title, and transcript location still undetermined. Where firsthand testimony conflicts with later accounts, I choose to compare dates and context rather than automatically treating recollection as conclusive. NVIDIA's executive biographies, historical timeline, and 2002 Form 10-K provide a separately attributable corporate record. This source hierarchy is part of my self-correcting research practice.

1.1 The making of an infrastructure leader: biography, apprenticeship, and the three-founder enterprise

I want to begin the NVIDIA story before the corporation existed, because a theory of leadership that starts with market capitalization can mistake the outcome for the cause. Huang's biography is not a decorative preface to the analysis. It is a sequence of formative circumstances, learned technical practices, and narrated interpretations that may illuminate how an individual came to exercise authority within a system that no individual built alone. My analytical obligation is to separate what is documented, what Huang later says those experiences meant, and what I infer from the resulting story.

Huang was born in Taiwan in 1963 and spent part of his childhood in Thailand before his family's move toward the United States. Accounts of his early schooling describe an extraordinary transition: he and his brother were sent ahead of their parents and attended the Oneida Baptist Institute in Kentucky. His retrospective stories about cleaning bathrooms, adapting to unfamiliar surroundings, and encountering hardship have become familiar elements of his public leadership narrative. I cannot infer a stable adult moral disposition from a childhood anecdote. I can ask why the story of hardship has become such a portable explanation for later corporate endurance. The precise birthplace is disputed in the biographical leads supplied for this chapter (Taipei versus Tainan); I retain Taiwan here rather than choose between them without examining the primary oral-history transcript.

In Oregon, Huang attended Aloha High School, developed his competitive table-tennis skills, and worked at Denny's. These episodes connect the mundane disciplines of practice, repetition, service work, and failure to the public vocabulary of resilience that he now deploys as a chief executive. Yet the same episodes resist a simple heroic reading. An immigrant childhood is lived within families, schools, labor markets, and institutions. Even a narrative organized around individual effort depends on relationships with teachers, parents, coworkers, competitors, and fellow students. My Z question therefore begins here: whose voices are audible when the success story is retold?

Huang earned his electrical engineering degree at Oregon State University and subsequently completed a master's degree at Stanford while working. NVIDIA's corporate biography confirms both degrees and his earlier employment at Advanced Micro Devices and LSI Logic. A historical NVIDIA filing describes him as a microprocessor designer at AMD and later director of the Coreware unit at LSI Logic. I read this sequence not as a sudden leap from adolescent hardship to visionary CEO, but as an apprenticeship in semiconductor engineering and

in the commercial organization of complex technical systems. (NVIDIA, n.d.-b; NVIDIA Corporation, 2002).

The LSI Logic period matters especially. A chip is never only a chip: its value depends on the surrounding design process, customer requirements, manufacturing arrangements, reusable intellectual property, and the organizational capability to turn an architecture into a product. The Coreware responsibility makes it plausible that Huang learned to think in terms of platforms and dependencies rather than isolated devices. This is an inference about the significance of a documented position, not an authenticated quotation about his private intentions. It helps explain why I treat the later NVIDIA ecosystem as an organizational achievement whose reach exceeds the physical sale of a processor.

I must correct a recurrent simplification in founder storytelling. NVIDIA was not founded by Huang alone. NVIDIA's own corporate history identifies three founders in April 1993: Jensen Huang, Chris Malachowsky, and Curtis Priem. The familiar Denny's meeting belongs to their collective origin story. Malachowsky brought experience in engineering and graphics systems, and Priem was another important technical founder. Huang became the president and CEO, but the company's technical birth involved complementary expertise, disagreement, negotiation, and collective risk. The absence of the other founders from a CEO-centered account is itself a useful warning about monophonic corporate history. (NVIDIA, n.d.-c).

Chris Malachowsky's continuing role also interrupts the myth that the founding partners vanished once the public-facing CEO became famous. NVIDIA's executive biography credits Malachowsky with leading work in product engineering, operations, and research. That record does not prove equal decision power, but it documents enduring technical agency beyond the chief executive. A rigorous Y analysis must therefore distinguish Huang's formal authority from the distribution of expertise, initiative, and decision rights within NVIDIA. I cannot code an organization as wholly monophonic simply because the press often speaks in the voice of one founder. (NVIDIA, n.d.-a).

Curtis Priem belongs in this inquiry for a related reason. The founding partnership was a meeting of technical imaginations rather than the realization of one person's solitary plan. The available historical accounts establish his cofounder status, but a full account of his contributions, later departure, and interpretation of NVIDIA's direction requires additional primary records and, ideally, oral history. I resist filling that gap with a polished legend. Fourth Wave Grounded Theory

gives the absent voice a place in the archive: not as evidence already collected, but as a counter-story that may revise my reading.

One reason I retain this detailed career history is to avoid judging the man through a single late-career interview. The early technical and partnership record supplies counterevidence to an overly narrow portrait of a leader concerned only with short-term transactions. It also opens the possibility that a highly collaborative technical enterprise can coexist with concentrated executive storytelling. Those two observations pull my provisional XYZA coordinates in different directions, and the tension should remain visible rather than being resolved by rhetorical confidence.

1.2 From gaming graphics to AI infrastructure

NVIDIA's corporate history begins on April 5, 1993, when Jensen Huang, Chris Malachowsky, and Curtis Priem founded the company to develop 3-D graphics for gaming and multimedia. The company's commercial starting point was video-game graphics. In 1999, NVIDIA introduced the GPU as a distinct computing architecture. In 2006, CUDA opened GPU parallel processing to scientific and research applications. In 2012, the AlexNet result demonstrated the usefulness of GPU computation for deep neural networks, helping connect graphics processors to modern AI development (Krizhevsky et al., 2012; NVIDIA, n.d.-c).

The transition did not occur in a single leap from gaming chips to AI chips. NVIDIA's programmable GPU architecture could perform many operations in parallel; CUDA supplied a software environment through which researchers and developers could use that capability; later systems combined accelerators, networking, software libraries, and data-center-scale configurations. As models and training workloads grew, these complementary assets made NVIDIA's existing technical path increasingly useful to AI firms. That account is a historically grounded explanation of the transition, not a claim that GPU hardware alone created the AI industry.

The financial record confirms that the original gaming business continues. For fiscal 2026, NVIDIA reported $16.042 billion in Gaming revenue and $193.737 billion in Data Center revenue. Gaming is not a vanished origin story; it remains part of the company's portfolio. Data Center, however, had become the much larger revenue segment. This change helps explain why NVIDIA now occupies a central position in AI infrastructure and why a shock to its leadership can affect suppliers, customers, investors, and firms that depend on NVIDIA compute (NVIDIA Corporation, 2026a).

The history supports a strong X trajectory in product and platform development, while leaving Y, Z, and A as separate questions. Technical innovation does not by itself establish service to affected others, plural decision-making, or moral answerability. I therefore use the three-decade chronology to define the case and identify organizational consequences; I do not treat commercial success as a proxy for all four behavioral dimensions.

The 33-year history is a selective chronology, not a complete census of NVIDIA events. The linked AI Leadership Event Archive is a source-led seed archive that begins this chronology and preserves dates, quotations, source identifiers, event scope, four numeric scores, rationales, and revision status. The historical narrative and scores remain open to expansion as primary records and counter-stories are added.

2. My starting narrative and its evidentiary limits

In an October 31, 2026, working account of this NVIDIA case, I provisionally described a Tesseract assessment and a possible negative post-data gain. I also proposed an Infrastructure Wall, a Regulatory Arbitrage Loop, and a Recursive Self-Reference Loop. I preserve that account as an initial, dated conjecture in my research archive, not as a published empirical result. Because its component ratings and calculation cannot be independently reconstructed, I do not carry its composite forward as a score in this chapter.

In that first account I examined two rooms. The first concerned a reported June 2026 Senate testimony invitation and the language of export-control compliance. The second concerned the claim that a supplier's responsibility ends at the chip while downstream operators make power and siting decisions. The model-like passages quoted in that account are not authenticated transcripts. I need their original prompts, model versions, dates, complete outputs, and the original committee or company documents. A generated response that says it is running on an H100 is not evidence of the machine on which it actually ran.

I do not delete those rooms. They are valuable clues to how a defensive organizational story might be generated. But I change their epistemic status. They are objects of inquiry: what did a model say when prompted about accountability, what assumptions were built into the question, and what independent sources would confirm or refute the resulting claim? This is a methodological correction that keeps the original intellectual risk alive.

2.1 The unfinished multiplicative proposal

My early Chokepoint Capture Vector proposed combining geohistorical located sourcing, regulatory inoculation, and wall evasion. The source manuscript contains blank equation placeholders, so I cannot honestly report a numeric vector value. I can, however, explain its proposed logic. Geohistorical sourcing asks how precisely a supplier can identify global production nodes; regulatory inoculation asks whether compliance talk deflects wider accountability questions; wall evasion asks whether the supplier/operator boundary is used to stop inquiry into shared consequences. Multiplying these proposed factors would require defined scales, justified functional form, and observations. Until those exist, the vector is an abductive design sketch, not a validated index.

3. Fourth Wave Grounded Theory: how I move between ground and theory

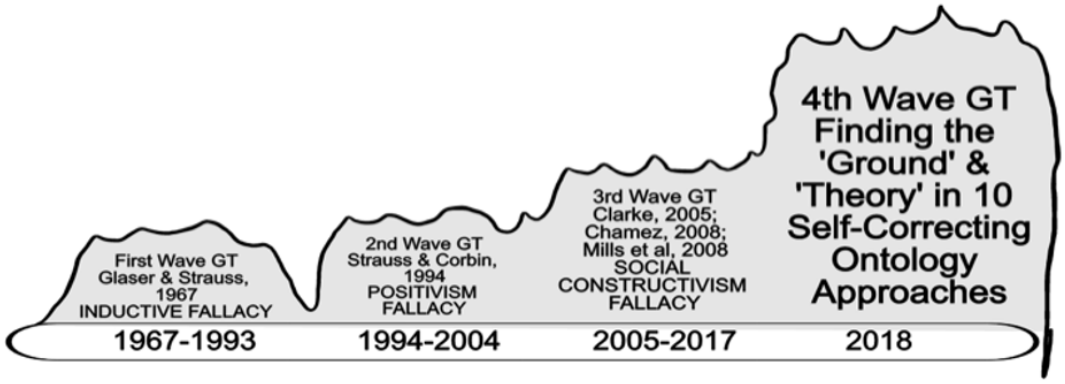


Figure 2.1. Four waves of grounded theory. Author-supplied conceptual history; the fourth wave is my programmatic proposal.

In the historical schematic I distinguish four methodological waves. The first, associated with Glaser and Strauss (1967), establishes the ambition to develop theory through systematic engagement with field material. Its possible fallacy is to mistake an inductive orientation for the absence of prior concepts. The second, associated here with Strauss and Corbin (1990), emphasizes coding procedures and systematic comparison. Its danger is procedural certainty: a method can look rigorous while its categories become insulated from surprise.

The third wave, associated in my diagram with social constructionist and constructivist developments including Charmaz (2006) and Clarke (2005), takes situated meanings, positionality, and the politics of interpretation seriously. Its possible vulnerability is not that constructionism must be relativistic, but that an analyst can invoke multiplicity without specifying what observation would force revision. These are critical characterizations for discussion, not claims that every scholar in a wave commits the same error.

The fourth wave is my and Grace Ann Rosile's programmatic movement toward finding both ground and theory in self-correcting ontological approaches. The source figure identifies ten such approaches, but this chapter does not reproduce an authenticated list from our 2020 book. I therefore do not invent the ten names. Instead, I demonstrate the central methodological commitment: no appealing theory, no sophisticated code, and no AI-generated consensus is exempt from correction by the concrete record and the people it concerns.

The 2020 conversational storytelling interviews project matters here because I do not want a corporate archive to become a substitute for lived stories. A data-center hearing, a worker's account, and a community's experience may resist the executive chronology. The fourth-wave task is to make these tensions generative rather than to file them away as noise. This is also why I resist treating the figure as a simple progress ladder: each wave carries resources I still need.

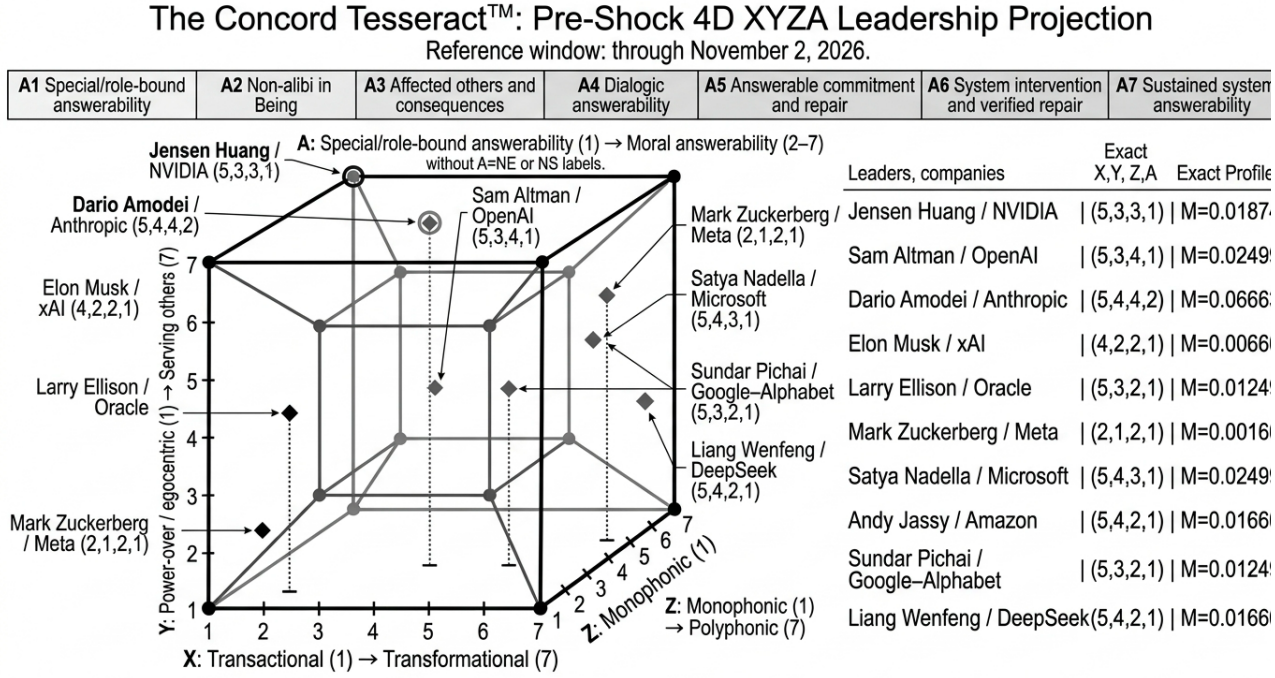


Figure 2.2. The Concord Tesseract™: Pre-Shock 4D Leadership Projection (through Nov 2, 2026).

The pre-shock Huang reference is (X=5, Y=3, Z=3, A=1), with M=0.018742. The post-window records separate Huang's departure episode from Kress's interim-leader episode. Figure 2.2 must be redrawn from Tables 2.7 and 2.8, with the former CEO event point identified separately from the current NVIDIA leader. The tables, rather than the superseded illustration, govern the coordinates.

How I read the axes. WAVE applies the behaviorally anchored 1–7 scales to dated records. X, Y, Z, and A are all scored numerically for every coded episode; A1 is not zero and is not "no answerability." The three A-cube drawing dimensions sometimes used to help render an event are graphic dimensions, not additional measured scores. Appendix A supplies the common anchors and Appendix D the updateable matrix and M operation.

For each record, M is calculated from the four numeric coordinates using $M=(X/7)\times(Y/7)\times(Z/7)\times(A/7)$. The pre-shock reference M values appear in Table 2.7. The event examples in Table 2.5 show how the same operation applies to a narrowly defined apology episode and to the broader departure episode; those event units are kept separate.

Leader	XYZA	M
Jensen Huang / NVIDIA	(5,3,3,1)	0.018742
Sam Altman / OpenAI	(5,3,4,1)	0.024990
Dario Amodei / Anthropic	(5,4,4,2)	0.066639
Elon Musk / xAI	(4,2,2,1)	0.006664
Larry Ellison / Oracle	(5,3,2,1)	0.012495
Mark Zuckerberg / Meta	(2,1,2,1)	0.001666
Satya Nadella / Microsoft	(5,4,3,1)	0.024990
Andy Jassy / Amazon	(5,4,2,1)	0.016660
Sundar Pichai / Google–Alphabet	(5,3,2,1)	0.012495
Liang Wenfeng / DeepSeek	(5,4,2,1)	0.016660

For the November event-window comparison, Table 2.8 lists numeric firm-and-leadership episode codes for all ten companies. NVIDIA's current post-window leader is Colette Kress; Huang's departure episode is shown separately in Table 2.5 and in the figure note. A1 is used where the episode contains no documented system intervention for affected others. It is a scored anchor, not a missing value. The M values are calculated for every four-component record.

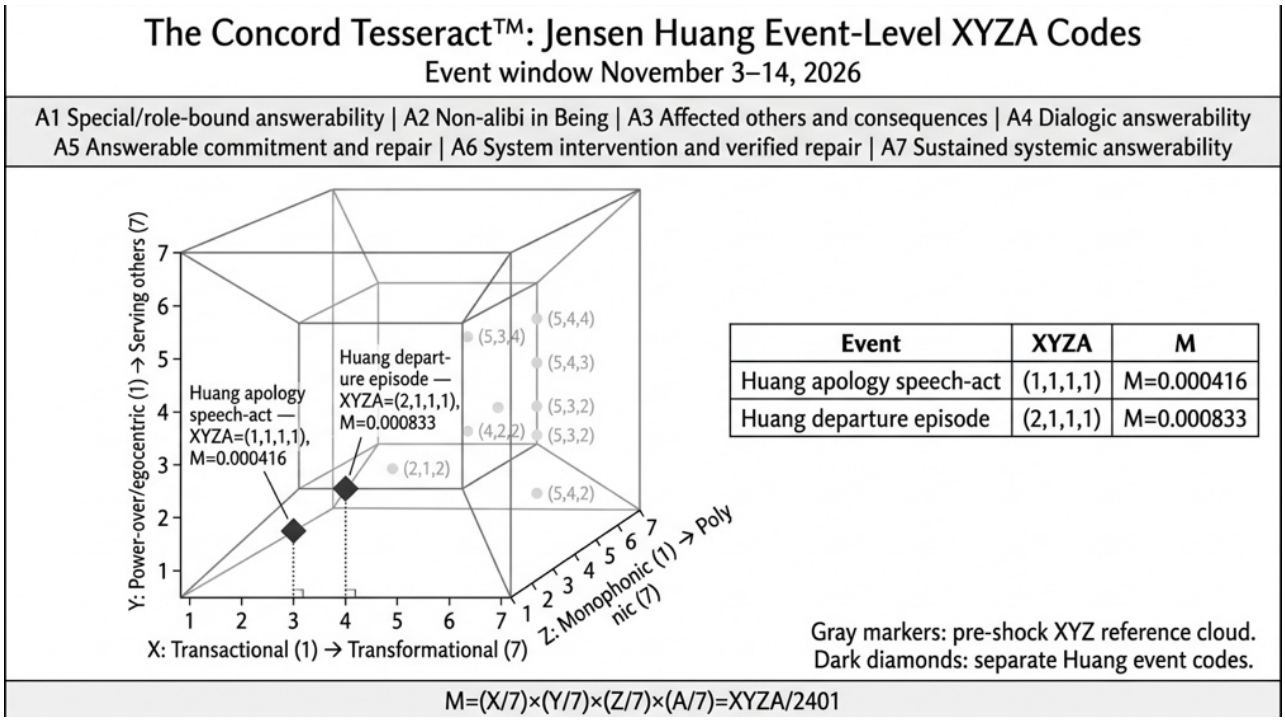


Figure 2.3. The Concord Tesseract™: Pre-Shock 4D Leadership Projection (through Nov 23-12, 2026).

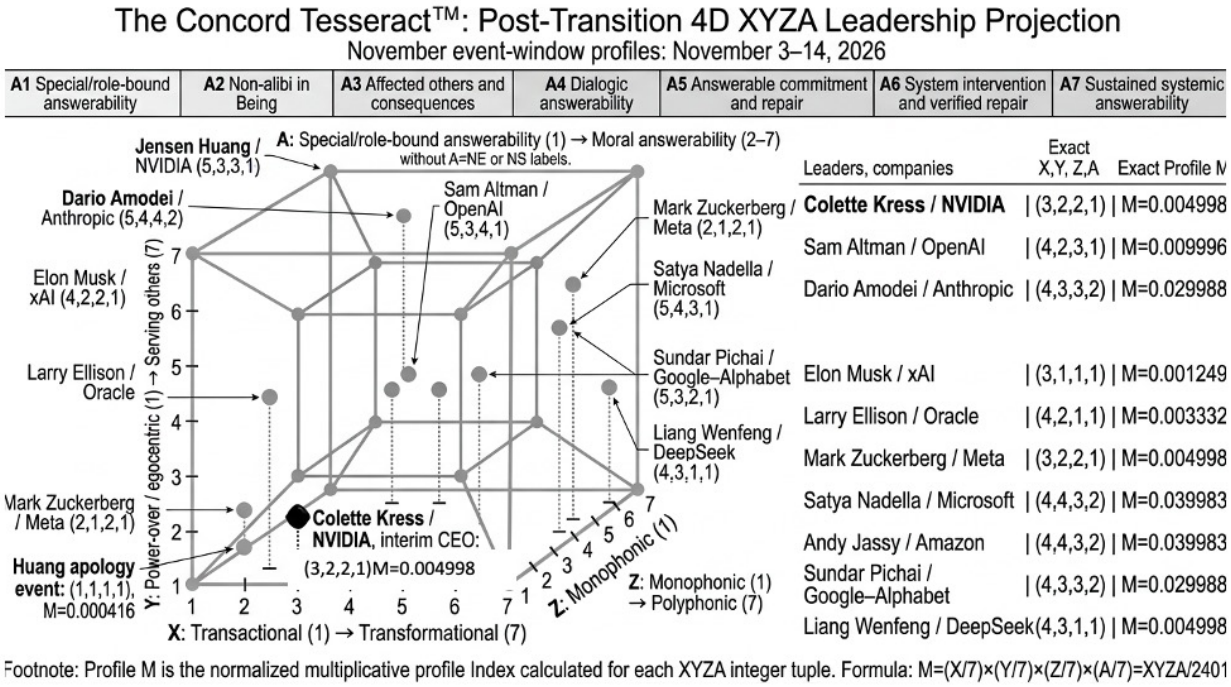


Figure 2.4. The Concord Tesseract™: Pre-Shock 4D Leadership Projection (through Nov 3-14, 2026).

3.1 Abduction is the imaginative guess before a worked-out hypothesis

Peircean abduction is not yet a neatly operationalized hypothesis. It is the imaginative leap that asks whether an initially surprising fact could make sense under a new possible explanation. My guess is that infrastructural transparency can coexist with answerability displacement. I did not derive this guess by calculating a probability. I arrived at it by noticing a tension between precise descriptions of hardware and diffuse accounts of responsibility.

Deduction comes afterward. I transform the guess into rival expectations about what documents, actions, and counter-stories I ought to encounter. Induction then compares those expectations with dated evidence, including negative cases. I return to abduction when the record surprises me again. In this way ADI is a recursive navigation practice, not a three-step ritual with a final destination.

The distinction is crucial to the monograph. A provisional 1–7 placement can express where my inquiry begins. It does not certify that the leader is a particular kind of person. I want the next researcher to tell me which bearing was wrong, why it was wrong, and how the chart should change.

4. The XYZA Navigation Figure: a four-dimensional research compass

I call this the Navigation Figure because sailors do not confuse a chart with the sea. A chart gives bearings and known hazards; the voyage produces observations that may require a corrected chart. My ten-leader map similarly provides a standing reference across chapters. The leader examined in a chapter is highlighted, but the other nine remain visible so that the comparative argument does not silently change. The origin and stems show a conceptual reference foundation, not a proven historical movement.

X runs from transactional (1) to transformational (7); Y from power-over or egocentric orientation (1) to serving others (7); Z from monophonic (1) to polyphonic (7); and A from special or role-bound answerability (1) toward increasingly demanding moral answerability (2–7). A is always a numeric 1–7 score. A1 identifies role-bound accountability, acknowledgment, or

explanation without demonstrated intervention in the system affecting others. A2 begins the non-alibi in being: the leader owns a unique participation and starts a concrete intervention. The later A anchors require progressively more dialogic, reparative, verified, and sustained action. The figure may show the XYZ projection with the A scale above it, but each record remains one four-dimensional XYZA vector.

The chapter compares pre-shock reference profiles through November 2 with November event-window codes through November 14. Tables 2.7 and 2.8 supply the numerical coordinates and M values. For NVIDIA, the post-window leadership profile is Colette Kress as interim CEO (3,2,2,1; M=0.004998); Jensen Huang's departure episode is recorded separately. The coordinates are working archival codes, not psychological diagnoses or causal estimates. Market movement and political context are recorded alongside the scores and do not mechanically determine them.

4.1 My first abductive bearing for Huang

The pre-shock reference vector for Huang is X=5, Y=3, Z=3, A=1, with profile M=0.018742. These codes are carried from the current workbook rather than newly inferred from the November events. The post-window records are in Tables 2.7–2.8. Appendix C holds the chapter-by-chapter score tables, Appendix D the updateable matrix and multiplicative rule, and Appendix E the image-generation prompts, which will be revised after the tables are finalized. Every code remains open to challenge against dated evidence.

Table 2.1. Provisional Huang XYZA abductive scoring and strongest counterevidence

Axis	Pre-shock code	Basis in the reviewed coding record	Evidence that could raise the code
X	5/7	Pre-shock workbook reference; transformational design and development are represented in the coded profile.	Evidence that would support a higher anchor: sustained redesign with verified system-wide stakeholder consequences.
Y	3/7	Pre-shock workbook reference; stakeholder service is partial and decision power remains uneven in the reviewed record.	Evidence that affected others share decision rights, resources, or protections and that outcomes are monitored.
Z	3/7	Pre-shock workbook reference; plural perspectives are filtered and their decision influence is not established.	A documented case in which dissent or community testimony changes a consequential decision.
A	1/7	A1 is role-bound accountability, acknowledgment, or explanation without demonstrated system intervention. A2 begins the non-alibi in being: personal ownership followed by a concrete first intervention.	Evidence for higher anchors must show intervention for affected others, dialogic decision change, funded repair, verified systemic effects, and sustained answerability, as detailed in Appendix A.

Reading this table: provisional Huang XYZA abductive scoring and strongest counterevidence provides a working comparison rather than an automatic verdict. The following discussion specifies what evidence could change the interpretation.

Table 2.1 is a navigational log for the pre-shock reference vector. The last column identifies evidence that could revise a code. It does not turn general aspirations into higher scores: a higher Z requires evidence that multiple voices changed a consequential decision, and a higher A requires evidence matching the next moral-answerability anchor.

4.2 How the fuller career history challenges my own XYZA bearings

The expanded biographical record is consistent with the workbook's X=5 pre-shock profile, which recognizes purposeful product and platform development. Further evidence would be needed

to support X6 or X7: systemic organizational redesign, verified consequences across boundaries, or sustained shared transformation over time.

The partnership record supports further examination of the workbook's Y=3 pre-shock code. Three founders created the enterprise, and NVIDIA acknowledges Malachowsky's sustained leadership in engineering and research. Those facts prevent me from equating CEO visibility with complete personal control. They do not establish shared authority with affected communities, workers, or customers; that requires evidence of who influenced consequential decisions.

For Z = 3, the most revealing question is whether the voices of Malachowsky, Priem, engineers, developers, customers, employees, and affected residents are permitted to transform the official account. A three-founder origin may demonstrate plurality at the moment of creation without proving consequential polyphony today. Conversely, a strong CEO narrative may coexist with substantial technical disagreement and learning inside the organization. I would need meeting records, oral histories, governance accounts, and concrete cases where dissent changed a decision before I could move this coordinate confidently.

For A, neither childhood hardship nor later financial success establishes answerability. The pre-shock profile is A1 because the reviewed record supports role-bound or special answerability without a demonstrated system intervention for affected others. A2 begins when a leader owns a unique participation without using office, market pressure, procedure, or board action as an alibi and starts concrete intervention. Huang's apology episode and departure episode receive separate A1 codes in this working archive because neither episode, as presently documented, demonstrates system repair for affected others. This code does not determine private character or legal liability.

The scoring implications are asymmetrical. The career narrative supplies substantial reasons to revisit a simplistic X judgment and meaningful reasons to investigate Y and Z further. It supplies less direct evidence about A, because an admirable professional trajectory is not the same as an observed response to a particular harm. This asymmetry is a methodological gain: I can identify precisely which new records would change each coordinate instead of moving all four together because a biography is inspiring or troubling.

4.3 November 2026: leadership transition, export policy, and market context

The November record changes the case sharply. Huang acknowledged infidelity in a November 8 statement. Reuters subsequently reported relationships involving Alice Huang, a business executive associated with Megaspeed International, an NVIDIA customer and partner, and NVIDIA employees, including senior vice president Kimberly Powell; Reuters also reported a

relationship with another employee whose identity it did not publish. Those are attributed reports, not all independently adjudicated findings. Reuters reported that Huang had described interactions with Alice Huang to senior NVIDIA leaders as routine business dealings, based on sources and messages reviewed by the news agency. NVIDIA's board announced an independent investigation. On November 12, NVIDIA's filing reported that the board had accepted Huang's resignation and appointed CFO Colette Kress interim CEO (Cherney & Dastin, 2026; NVIDIA Corporation, 2026b; Reuters, 2026a, 2026b).

The espionage claim requires a separate evidentiary label. U.S. authorities reportedly scrutinized Alice Huang over suspicions of intelligence activity for China, but Reuters stated that it could not verify whether she worked for the Chinese government; Beijing denied that she was a spy, and Reuters reported no public finding that Megaspeed violated export rules. I therefore describe this as an unverified espionage suspicion and an active investigative question. I do not write that an espionage operation has been proven or that the reported relationship established a chip diversion (Reuters, 2026b).

The export-control chronology also complicates a simple claim that every NVIDIA chip sold to China violated a U.S. directive. The United States imposed licensing restrictions on advanced computing chips, including successive restrictions on A100/H100-class products and later H20 and H200 products. NVIDIA's 2026 quarterly filing reports that the U.S. government issued licenses for limited sales to specified Chinese customers, while Chinese government restrictions prevented or limited some shipments. The filing says licensed H200 shipments were only a fraction of those allowed and represented less than one percent of Data Center revenue in the most recent quarter. It also describes tariffs and an inspection process. These records show that U.S. policy included both restrictions and conditional licensing. They do not establish that Huang or NVIDIA violated export law in the alleged relationship, or that U.S. officials knowingly approved any illicit diversion (NVIDIA Corporation, 2026c).

The Trump administration's role must be stated with the same care. The administration changed export-license policy for specified advanced chips, including H200, to case-by-case review when conditions are met; NVIDIA's filing separately describes limited licenses and unresolved restrictions. This is evidence of a policy decision to permit some licensed sales, not evidence that the president or cabinet knowingly approved an illicit transfer or a spy operation. The public financial-disclosure record also shows that Trump bought as much as $1.6 million and sold

as much as $1 million of NVIDIA stock in August 2026. The transactions were disclosed before the November scandal reporting; they warrant examination as a conflict-of-interest context, but do not by themselves establish advance knowledge, insider trading, or a causal connection to the later events. Reuters reported that Trump said Huang had told him he did not know Alice Huang was suspected of spying; that statement is attributed reporting, not an independent finding about what either man knew. The chapter can examine the tension among national-security policy, commercial access to China, political access, and AI-industry influence without collapsing these separate claims into one conclusion (Bureau of Industry and Security, 2026; Office of Government Ethics, 2026; NVIDIA Corporation, 2026c; Reuters, 2026b, 2026f).

NVIDIA's centrality also has a financial dimension. The company announced financing platforms intended to mobilize more than $500 billion in third-party capital for AI infrastructure, and its filing described conditional guarantees capped at $105 billion for an OpenAI-linked data-center project. These figures are announced or conditional commitments, not realized cash losses. They support the description of NVIDIA as a major financing participant in the AI buildout; calling it the industry's "bank" is a metaphor that must not obscure the actual contractual terms (NVIDIA, 2026; NVIDIA Corporation, 2026c).

The market shock should be reported with its time window and attribution limits. Reuters reported that NVIDIA shares had fallen more than 30 percent from their October 6 peak by November 14, erasing more than $1.8 trillion in market capitalization. Market capitalization is not cash removed from the company, and Reuters also described funding and regulatory concerns affecting AI-linked stocks. The election added a political signal: data-center opposition gained visibility in the November 3 results, with debate over electricity, water, land use, and local consent. Neither the stock decline nor the Democratic gains assign lower XYZA scores automatically. They alter the political-economic context and create new events to code when leaders or organizations act (Reuters, 2026d, 2026e).

This context gives the chapter a place for anti-narrative bets on the future. I distinguish a bet from an observed finding. One bet is that AI infrastructure may increase productivity and access while concentrating returns among owners of compute and finance, with energy, water, and labor costs shifted toward other stakeholders. A rival bet is that competition, public policy, and community bargaining will distribute more of the gains. Each bet requires a mechanism, a time

horizon, indicators, and evidence that could falsify it. The WAVE scores describe coded events; they do not forecast which bet will prevail.

5. The evidence trail: quotations, rival readings, and source status

The next table distinguishes company records, regulatory filings, and attributed reporting. The exact quotation fragments are short; surrounding context is paraphrased. An attributed news report is not converted into an official finding, and company claims are identified as company claims. Each row retains a counter-reading that could change my interpretation.

Table 2.2. Selected NVIDIA records, reported claims, and rival readings

ID / date	Record or topic	Axis	Source status	Rival reading / scoring implication
NV-HIST / 1993–2012	Gaming and multimedia origin; GPU milestone in 1999; CUDA in 2006; GPU-enabled deep learning in 2012.	X	Company history and peer-reviewed conference paper.	A cumulative technology path supports X transformation; it does not alone establish service, polyphony, or A.
NV-FY26	Data Center revenue was $193.737 billion; Gaming revenue was $16.042 billion.	X / context	NVIDIA Form 10-K.	Data Center dominance is real, while gaming continues. Revenue does not measure social distribution.
NV-EXPORT / 2025–2026	NVIDIA reports U.S. licenses for specified products and customers, with Chinese restrictions limiting some shipments.	A / context	NVIDIA Form 10-Q; regulatory filing.	Licensed trade is distinct from an unauthorized diversion; a public finding of a specific violation is not established here.
NV-REL / November 2026	Huang acknowledged infidelity; news reports describe relationships with a customer-side executive and NVIDIA employees.	A / Z	Huang statement reported by Reuters; details attributed to sources.	Use the report to define an event record, not to infer a total personal profile without further evidence.
NV-INTEL / November 2026	Authorities reportedly scrutinized Alice Huang; Reuters said it could not verify government employment or spying.	A / context	Attributed reporting; investigation unresolved in the public record reviewed.	Do not call the allegation proven or equate it with evidence of an export-control violation.
NV-CEO / November 12	NVIDIA's SEC filing records Huang's resignation and Kress's interim appointment.	X / A	Primary SEC Form 8-K.	Board action and appointment are organization-level events; they do not by themselves demonstrate moral repair.

NV-CAPITAL / 2026	NVIDIA announced a third-party capital-mobilization target and disclosed a conditional guarantee cap tied to AI infrastructure.	X / Y	Company announcement and SEC filing.	These are target and contingent commitment figures, not realized financing or losses.
POLITICAL ECONOMY / November 2026	Reuters reported a large NVIDIA market-cap decline and greater visibility for data-center opponents after the election.	Y / Z / context	Independent reporting; market and political effects have multiple causes.	Treat as political-economic context and a source of future events, not as a direct behavioral score.

Reading this table: quotation leads, opposing interpretations, and archival verification provides a working comparison rather than an automatic verdict. The following discussion specifies what evidence could change the interpretation.

The first three records show why the history and the coding must remain distinct. The GPU/CUDA/AlexNet sequence supports an account of technical transformation, and the annual report confirms that Data Center revenue now greatly exceeds Gaming revenue. Neither fact tells me who benefits, whose voice shapes decisions, or whether NVIDIA has intervened to repair consequences for others. Those questions require separate Y, Z, and A evidence.

The November records also have different evidentiary status. A board filing can establish a resignation and appointment. Reuters can report what sources and messages indicate about private conduct, while still leaving intelligence claims unverified. I preserve those distinctions because a scandal narrative can become monophonic when allegation, fact, and inference are collapsed into one account.

6. My abductive–deductive–inductive inquiry in practice

6.1 Abduction: the Infrastructure Wall as a wild possibility

I imagine a wall that is made not of secrecy but of accurate distinctions. The supplier makes chips, the customer builds a data center, the utility sells electricity, and the public agency grants permits. Each statement may be true. My abductive surprise is that the sum of true statements may still leave affected people unable to locate a responsible interlocutor. I call that possible mechanism the Infrastructure Wall. I do not yet know whether it describes NVIDIA's conduct, a wider infrastructure regime, or the language my own prompts elicited from AI systems.

6.2 Deduction: rival hypotheses without confusing H0 with a null hypothesis

I use H1, H2, H3, and R1 as labels for substantive rival explanations, not as statistical null-hypothesis tests. Trafimow's methodological work emphasizes making assumptions and levels of

theorizing explicit; I apply that discipline here by specifying what records could discriminate among explanations. The XYZA product used in this monograph is my operationalization. I do not attribute this exact formula to Trafimow unless a specific source passage establishes that attribution (Trafimow, 2023).

Table 2.3. Rival explanatory claims and evidence that could defeat them

Claim	If true, I expect	Evidence that weakens it
H1 Infrastructure Wall	Repeated boundary rhetoric with avoidable non-cooperation	Sustained, traceable cooperation and remedy
R1 justified division of responsibility	Decisions and obligations match actual control	Documented influence combined with refusal to respond
H2 differentiated shared responsibility	Responsibility varies across design, siting, operation and repair	Uniform conduct across stages contradicts differentiation
H3 AI framing artifact	Wall appears under leading prompts but fades under neutral prompts	Independent records reveal same pattern without prompting

Table 2.3 lays out rival explanations and identifies evidence that could weaken each one. The discussion that follows specifies what would change my interpretation.

The four explanations are not four equally likely boxes. They organize my search. A committee record could disconfirm a dramatic claim about refused testimony. A utility filing could show that an operator, not a chip supplier, controlled a fuel choice. A supplier document could show early knowledge of consequences and meaningful influence over decisions. The evidence matters precisely because each observation can move the inquiry in a different direction.

6.3 Induction: an archive is a record, not a pile of AI summaries

For this monograph a research archive is a retrievable and dated chain connecting the original question, the original answer or document, the person or organization responsible for it, my coding decision, and every subsequent correction. A news summary is not the original recording. A generated quotation is not the corporate statement it purports to describe. A screenshot is not enough without provenance. I want another researcher to be able to find the same item and disagree with me about its meaning.

Table 2.4. Case-specific archival register and counter-story search

Issue	Primary record to retrieve	Alternative explanation	Current status
Senate testimony	Committee invitation, hearing record, response	Scheduling or jurisdiction	Unverified
Energy and emissions	Permits, utility filings, contracts, public comments	Operator rather than supplier control	To retrieve

Supply concentration	Filings, credible market data, contracts where public	Substitution and rival suppliers	To retrieve
AI runtime claims	Prompt, model/version, deployment evidence	Invented runtime knowledge	Unverified
Stakeholder voice	Hearings, worker and resident testimony, responses	Meaningful participation or remedy	To retrieve

Table 2.4 is a case-specific archival register for the next research pass. Its pending fields identify missing records, not findings.

Each row in Table 2.4 is an assignment for a future archival pass. I am not presenting an empty "pending" column as a result. I am showing exactly where the chapter remains vulnerable. The most important missing materials are not additional executive quotations alone. They are records of decisions and responses by affected parties. A research archive should preserve both a corporation's answer and the question that the answer failed—or succeeded—to address.

6.4 Correction: what has changed already

I no longer assert that responsibility never reaches the hardware layer. That phrase belonged to the strong October 31 conjecture; I now treat it as a claim to test. I also separate the historical profile from the November shock. The archive traces NVIDIA from its 1993 founding, but its current historical rows are a seed chronology, not a complete census of every event across 33 years. The pre-shock reference and event-window scores must remain labeled by their units and evidence status. A movement between two plotted points is a change in coded episode or profile; it is not, by itself, a causal estimate of the scandal's effect.

7. WAVE scoring, qualimetrics, and the multiplicative index

7.1 WAVE event coding and the multiplicative index

WAVE begins with a dated event unit. Each archive row identifies the actor, company or network level, event family, exact short excerpt or document passage, surrounding context, source, evidence class, and coding rationale. The qualitative record is then coded separately on X, Y, Z, and A against the common 1–7 behavioral anchors. The archive also records frequency counts, financial indicators, and relevant political-economic context. These measures remain visible beside the qualitative coding; they are not silently combined into M.

For every coded episode, I assign numeric X, Y, Z, and A scores from 1 to 7. A profile score is a labeled synthesis of the evidence within its stated time window, not a direct score of a stock movement or a network shock. Where an event occurs at company or industry level, I identify the actor level and score the documented organizational response or exposure; I do not attribute an event to a CEO merely because that person leads the company.

A is always numeric from 1 to 7. A1 codes role-bound accountability, acknowledgment, or explanation without a demonstrated system intervention for affected others. A2 begins the non-alibi in being through personal ownership and a concrete first intervention. A3–A7 require progressively stronger attention to affected others, dialogic decision change, funded repair, verified system intervention, and sustained answerability. Thin or disputed evidence is recorded in the confidence and source-status fields; it does not turn A into NE or NS.

For each event or explicitly labeled profile, the multiplicative index is $M=(X/7)\times(Y/7)\times(Z/7)\times(A/7)$, equivalently $XYZA/2401$. The minimum is 1/2401, approximately 0.000416, and the maximum is 1. M is a derived index, not a fifth coordinate, a probability, a measure of financial performance, or a verdict on private character. The product encodes a noncompensatory assumption: a low dimension constrains the index. Treating 1–7 behavioral anchors as equally spaced numeric factors is an explicit modeling assumption that future researchers should test.

I apply multiplication as this monograph's qualimetric operationalization, informed by Trafimow's demand that theoretical and auxiliary assumptions be made explicit. This distinction matters: the scale anchors are qualitative behavioral judgments; event frequencies are counts; financial measures retain their units and periods; and SEAM hidden-cost indicators such as absenteeism, accidents, turnover, non-quality, direct productivity, and lost opportunities require their own denominators and valuation methods. None is added to M unless the scoring model is explicitly revised and tested (Trafimow, 2023).

Table 2.5. WAVE XYZA scores and multiplicative calculations for NVIDIA episodes

Record / unit	X	Y	Z	A	M	Interpretation
Huang pre-shock reference profile	5	3	3	1	0.018742	Working profile through Nov. 2; A1 is role-bound answerability.

Record / unit	X	Y	Z	A	M	Interpretation
Huang apology event	1	1	1	1	0.000416	Narrow event code; statement alone does not demonstrate system repair.
Huang departure episode	2	1	1	1	0.000833	Attributed reporting plus resignation; event code, not a verdict on character or legal liability.
Kress interim-leader episode	3	2	2	1	0.004998	Provisional transition code; appointment alone shows no verified post-appointment system intervention. Low confidence.

Table 2.5 shows the same multiplication for the pre-shock Huang profile, the narrow apology episode, the departure episode, and Kress's provisional interim-leader episode. All four components are numeric in every row. These are different units and should not be read as a direct causal estimate of personal change.

The row label controls the interpretation. Huang's pre-shock profile summarizes the working reference used in this chapter. His apology and departure episodes are event-level codes. Kress's row is a provisional interim-leadership code based on the appointment and transition record; the appointment alone is not evidence of system repair. The provisional codes are X=3, Y=2, Z=2, A=1; M=0.004998, pending direct post-appointment evidence and independent coding. The scores need further independent coding before they are treated as final.

7.2 My uncertainty statement is not a confidence interval

I can state abductive uncertainty ranges for Huang's pre-shock X, Y, and Z reference codes while archival verification continues: X plausibly 4–6, Y 2–4, and Z 2–4. These are judgment ranges, not confidence or credible intervals. A remains numeric; for the current pre-shock profile and the narrowly scoped apology and departure episodes, A=1. I record the uncertain evidence and competing readings in the archive rather than inserting a nonnumeric score.

I report no numerical probability. M is a normalized multiplicative index under the stated equal-spacing assumption. Profile M and event-level M values remain distinct, and financial or political-economic shifts are not substituted for leader behavior.

7.3 A biographical counter-archive and a second abductive cycle

The three-founder history suggests a concrete next archival cycle. I would first assemble contemporary founding accounts, early filings, technical histories, and available oral interviews from all three founders. I would then compare the executive retrospective with engineers' and partners' descriptions of decisions, disagreements, and product pivots. Finally, I would bring those accounts into dialogue with contemporary records concerning supply relationships, power demand, and communities. The goal is not to prove that a childhood experience caused a later business decision. It is to follow how a story becomes an organizational explanation and where the explanation leaves something out.

I would construct a paired timeline: one chronology of technologies and organizational decisions, another of narrated meanings and public accountability claims. The timelines should not be collapsed. A product launch can be dated; the claim that a setback taught humility is a retrospective interpretation; the proposition that humility changed a later decision is a causal hypothesis requiring additional evidence. Keeping these layers distinct is an elementary safeguard against the Ghost Vortex in which every success retrospectively confirms the founder's character.

Here my Concord method has a practical task. One AI collaborator can search for corroborating primary records; another can argue that I have undervalued technical transformation; another can challenge my assumption that ecosystem concentration necessarily produces answerability displacement. I, as the author, must keep their disagreements intact and decide which sources justify revision. Agreement among assistants cannot substitute for the missing founders' voices or for residents who live near infrastructure facilities.

8. Human–AI Concord, disagreement, and the second cycle

I used AI collaborators as provocateurs, not as a jury. Proton, Grok, Claude, Google AI, and ChatGPT contributed distinct interpretations and editorial suggestions during the preliminary inquiry. Some supplied useful leads; others questioned the calculations. Their apparent agreement cannot serve as independent empirical confirmation because they may share training material, web sources, and prompting assumptions. Their disagreement is more valuable: it makes me explain what I mean by probability, which observations are authenticated, and where a critical theory might become an alibi of its own.

This is the practical meaning of Concord in my writing. I do not ask the AI systems to agree with me. I ask them to preserve opposition, point to counterevidence, and expose the blind spots of

my interpretation. A human researcher who later corrects the Huang score will not have spoiled the chapter. That correction will be evidence that the method is working.

I therefore plan the next cycle around a shared source packet. Each human coder receives the same dated records, the same four-axis anchors, and the same rule that A receives a numeric 1–7 code. Coders work independently before comparing decisions. We preserve disagreements in the rationale and confidence fields, revisit the anchors when evidence changes, and record which source moved the bearing. A coefficient of agreement may become appropriate after a real coding exercise; it cannot be computed from empty worksheets or AI-generated consensus.

Table 2.6. Minimum fields for independent coding and correction

Archive field	What I require	Why it matters
Event ID and source	Stable URL, source date, original medium, and event identifier	Reproducibility
Actor and event scope	Leader, organization, board, or network; identify direct action versus exposure context	Prevents false attribution and conflation of firm and leader conduct
Exact language and context	Short verbatim excerpt, page or timestamp, and surrounding paraphrase	Avoid false supplier/operator attribution
Authority and dependencies	Who could decide, warn, intervene, finance, or remedy; record network links	Abductive self-correction
Counter-story / rival reading	Opposing account, affected voice, and evidence that would falsify the initial reading	Auditable navigation
XYZA memo and revision log	All four numeric anchors, confidence, rival score, M calculation, prior code, reason, and revision date	Genuine movement versus reclassification

Reading this table: minimum fields for independent coding and correction provides a working comparison rather than an automatic verdict. The following discussion specifies what evidence could change the interpretation.

Table 2.6 is an instrument rather than a result. It tells future researchers what to collect and how to preserve disagreement. I will not call an empty field a zero, and I will not average incompatible moral interpretations to produce a falsely precise A. I want a visible path from testimony and decisions to coding and then to a corrected Navigation Figure.

The following tables give the common pre-shock reference profiles and the November 3–14 event-window profiles for the ten companies. Table 2.7 preserves the current working baseline; Table 2.8 shows the post-window leadership or firm-event codes, with NVIDIA represented by interim CEO Colette Kress. Huang's departure episode remains a separate event record. The post codes are provisional, and their event descriptions identify where a score reflects direct action and where it reflects company or network context.

Table 2.7. Pre-shock XYZA reference profiles for ten AI leaders through November 2, 2026

Chapter / company / leader	X	Y	Z	A	M	Status
2. NVIDIA / Jensen Huang	5	3	3	1	0.018742	Pre-shock working profile; A1 is role-bound answerability.
3. OpenAI / Sam Altman	5	3	4	1	0.024990	Pre-shock working profile; A1.

Chapter / company / leader	X	Y	Z	A	M	Status
4. Anthropic / Dario Amodei	5	4	4	2	0.066639	Pre-shock working profile; A2 is the highest reviewed anchor in this record.
5. xAI / Elon Musk	4	2	2	1	0.006664	Pre-shock working profile; A1.
6. Oracle / Larry Ellison	5	3	2	1	0.012495	Pre-shock working profile; A1.
7. Meta / Mark Zuckerberg	2	1	2	1	0.001666	Pre-shock working profile; A1.
8. Microsoft / Satya Nadella	5	4	3	1	0.024990	Pre-shock working profile; A1.
9. Amazon / Andy Jassy	5	4	2	1	0.016660	Pre-shock working profile; A1.
10. Google / Alphabet / Sundar Pichai	5	3	2	1	0.012495	Pre-shock working profile; A1.
11. DeepSeek / Liang Wenfeng	5	4	2	1	0.016660	Pre-shock working profile; A1.

Table 2.8. Provisional XYZA leadership and firm-event codes for November 3–14, 2026

Company / focal leader in window	X	Y	Z	A	M	Event basis and interpretation
NVIDIA / Colette Kress, interim CEO	3	2	2	1	0.004998	Interim transition; reported CFO countervoice; no verified post-appointment system repair. Provisional.
OpenAI / Sam Altman	4	2	3	1	0.009996	Compute and financing exposure in the NVIDIA-centered network; no automatic inference of realized loss. Provisional.

Company / focal leader in window	X	Y	Z	A	M	Event basis and interpretation
Anthropic / Dario Amodei	4	3	3	2	0.029988	Compute exposure; existing safety commitments inform A, but are not new November repair. Provisional.
xAI / Elon Musk	3	1	1	1	0.001249	Election and political-economy exposure; leader-specific post-window action needs direct-source review. Low confidence.
Oracle / Larry Ellison	4	2	1	1	0.003332	Data-center capacity and financing exposure; distinguish contractual risk from realized loss. Provisional.
Meta / Mark Zuckerberg	3	2	2	1	0.004998	Data-center and local-opposition context; no verified event-window repair in this record. Provisional.
Microsoft / Satya Nadella	4	4	3	2	0.039983	Reported transparency action on local-government NDAs; scope and effects require verification. Provisional.
Amazon / Andy Jassy	4	4	3	2	0.039983	Reported transparency and community commitments; delivery and local outcomes remain to be verified. Provisional.

Company / focal leader in window	X	Y	Z	A	M	Event basis and interpretation
Google / Alphabet / Sundar Pichai	4	3	3	2	0.029988	Reported NDA and workforce-term changes; action is bounded and scope remains under review. Provisional.
DeepSeek / Liang Wenfeng	4	3	1	1	0.004998	Reported chip-supply adaptation; treat as a company supply-chain episode, not a personal moral verdict. Provisional.

All rows in Table 2.8 carry numeric X, Y, Z, and A values, with M calculated from the same four scores. A1 is the lowest anchored answerability code and applies where the coded episode shows role-bound acknowledgment or no demonstrated intervention for affected others. The context column identifies events at company or network level; those codes do not claim that every CEO personally caused the market or political shift. Financial outcomes and anti-narrative bets remain separate from the XYZA score.

9. Discussion: the dark side and the counter-story I must be willing to hear

The dark side I am investigating is not merely that AI uses electricity. It is the possibility that a highly accurate account of efficiency and contractual separation becomes an alibi when concrete people seek an answer about costs imposed on them. But I must resist the inverse temptation. A critical vocabulary can become its own Ghost Vortex when every counterexample is reinterpreted as further evidence for the theory. The fourth-wave discipline applies to my claims as much as to Huang's.

The strongest counter-story may be that a supplier can acknowledge risks, support verification, and improve outcomes while lacking authority over customer siting decisions. I must not collapse the manufacturer, the developer, the utility, the regulator, and the community into one imaginary actor. Bakhtinian answerability is situated. It asks who could answer to whom in a concrete event. It does not erase differences in legal authority, knowledge, or practical power.

The preliminary sources report that Huang has spoken about AI laboratories' responsibility not to ship uncontrolled products and about engaging communities. I cannot ignore those reports simply because they complicate my starting position. They may reveal a genuine ethics

architecture; they may also reveal a boundary between responsibility demanded of others and responsibility practiced within NVIDIA's own sphere. Only documented action can help decide.

I end the case discussion with a deliberately vulnerable question: under what conditions does a chokepoint supplier's accurate statement about the limits of its control become an inadequate moral answer? The question cannot be resolved by arithmetic alone. It requires a chronology of knowledge, decisions, affected voices, and opportunities to respond. My first score directs that search; it does not finish it.

10. Conclusion: keeping the navigation chart open

I have offered a pre-shock reference profile for Huang and a separate November event-window comparison across ten AI companies. The pre-shock Huang profile (5,3,3,1) yields M=0.018742. The narrow apology event (1,1,1,1) yields M=0.000416; the broader departure episode (2,1,1,1) yields M=0.000833. Kress's provisional interim-leadership profile (3,2,2,1) yields M=0.004998. These M values are arithmetic consequences of the assigned anchors, not causal estimates or measures of guilt. The November evidence changes the story and the archive, while the scores remain open to correction through better records and independent coding.

The substantive contribution is a historically grounded account of how a gaming-graphics company became a central AI infrastructure supplier, and a testable proposal that chokepoint positions can separate technical centrality from responsibility for network consequences. The methodological contribution is a dated event archive that joins qualitative interpretation, event counts, financial indicators, and political-economic context while keeping their units distinct. WAVE codes each event on behaviorally anchored XYZA scales and calculates M as an explicit, revisable multiplicative index.

I invite researchers to bring contrary records and to falsify the proposed mechanisms. Show me a source that changes the chronology, a document that clarifies who controlled an export or financing decision, a hearing or community record that changed a consequential choice, or a coded event that warrants another anchor. The analysis should proceed through abduction, deduction, induction, and self-correction. A future researcher who revises the theory helps build the science of Concord Tesseract™ Leadership.

References

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

Bureau of Industry and Security. (2026, January 13). Department of Commerce revises license review policy for semiconductors exported to China. U.S. Department of Commerce. https://media.bis.gov/sites/default/files/documents/DoC%20Revises%20License%20Review%20Policy%20for%20Semiconductors%20Exports.pdf

Boje, D. M. (1979). Relative centrality in interorganizational networks [Doctoral dissertation, University of Illinois at Urbana-Champaign]. https://www.ideals.illinois.edu/items/67493

Boje, D. M., & Rosile, G. A. (2020). How to use conversational storytelling interviews for your dissertation. Edward Elgar Publishing.

Boje, D. M., & Whetten, D. A. (1981). Effects of organizational strategies and contextual constraints on centrality and attributions of influence in interorganizational networks. Administrative Science Quarterly, 26(3), 378–395.

Charmaz, K. (2006). Constructing grounded theory: A practical guide through qualitative analysis. Sage.

Cherney, M. A., & Dastin, J. (2026, November 13). Ex-Nvidia CEO misled top lawyer on affair with Chinese businesswoman, sources say. Reuters. https://www.reuters.com/world/china/ex-nvidia-ceo-misled-top-lawyer-on-affair-with-chinese-businesswoman-sources-say-2026-11-14/

Clarke, A. E. (2005). Situational analysis: Grounded theory after the postmodern turn. Sage.

Glaser, B. G., & Strauss, A. L. (1967). The discovery of grounded theory: Strategies for qualitative research. Aldine.

Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet classification with deep convolutional neural networks. In F. Pereira, C. J. C. Burges, L. Bottou, & K. Q. Weinberger (Eds.), Advances in neural information processing systems (Vol. 25, pp. 1097–1105). Curran Associates.

Li, J., & Li, S. (2026, November 4). Analysis: Trump's midterm drubbing demonstrates the power of data center opponents. Reuters. https://www.reuters.com/world/us/trumps-midterm-drubbing-demonstrates-power-data-centre-opponents-2026-11-04/

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

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

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

NVIDIA. (2026, August 10). NVIDIA partners with Apollo, BlackRock, Blackstone, Brookfield, Goldman Sachs, and KKR to establish AI compute infrastructure financing platforms to mobilize over $500 billion of third-party capital. NVIDIA Investor Relations. https://investor.nvidia.com/news/press-release-details/2026/NVIDIA-Partners-With-Apollo-BlackRock-Blackstone-Brookfield-Goldman-Sachs-and-KKR-to-Establish-AI-Compute-Infrastructure-Financing-Platforms-to-Mobilize-Over-500-Billion-of-Third-Party-Capital/default.aspx

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

Office of Government Ethics. (2026, September 8). President Donald J. Trump: Periodic transaction report. The White House. https://www.whitehouse.gov/wp-

content/uploads/2026/09/President-Donald-J.-Trump-Periodic-Transaction-Report-09.08.26.pdf

NVIDIA Corporation. (2026a). Annual report (Form 10-K) for the fiscal year ended January 25, 2026. U.S. Securities and Exchange Commission. https://www.sec.gov/Archives/edgar/data/1045810/000104581026000038/a2026-annualxreportxwebxfi.pdf

NVIDIA Corporation. (2026b, November 12). Current report (Form 8-K). U.S. Securities and Exchange Commission. https://www.sec.gov/Archives/edgar/data/1045810/000104581026000082/nvda-20261111.htm

NVIDIA Corporation. (2026c, August 27). Quarterly report (Form 10-Q) for the quarter ended July 26, 2026. U.S. Securities and Exchange Commission. https://www.sec.gov/Archives/edgar/data/1045810/000104581026000075/nvda-20260726.htm

Peirce, C. S. (1931–1958). Collected papers of Charles Sanders Peirce (C. Hartshorne, P. Weiss, & A. W. Burks, Eds., Vols. 1–8). Harvard University Press.

Reuters. (2026f, October 8). Trump discloses Nvidia stock trades as he prepares to honor CEO Huang. Reuters. https://www.reuters.com/business/trump-discloses-nvidia-stock-trades-he-prepares-honor-ceo-huang-2026-10-08/

Reuters. (2026a, November 12). Nvidia CEO Jensen Huang resigns in blow to AI boom he helped foster. Reuters. https://www.reuters.com/business/retail-consumer/nvidia-ceo-jensen-huang-step-down-big-blow-ai-boom-2026-11-12/

Reuters. (2026b, November 13). Former Nvidia CEO Huang had a relationship with a second employee, sources say. Reuters. https://www.reuters.com/technology/former-nvidia-ceo-huang-had-relationship-with-second-employee-sources-say-2026-11-12/

Reuters. (2026c, November 14). What we know about Chinese businesswoman Alice Huang. Reuters. https://www.reuters.com/world/asia-pacific/what-we-know-about-chinese-businesswoman-alice-huang-2026-11-14/

Reuters. (2026d, November 13). Neoclouds, AI stocks slide again as Nvidia turmoil rattles sector. Reuters. https://www.reuters.com/business/neoclouds-ai-stocks-slide-again-nvidia-turmoil-rattles-sector-2026-11-13/

Reuters. (2026e, November 4). Analysis: Trump's midterm drubbing demonstrates the power of data center opponents. Reuters. https://www.reuters.com/world/us/trumps-midterm-drubbing-demonstrates-power-data-centre-opponents-2026-11-04/

Strauss, A., & Corbin, J. (1990). Basics of qualitative research: Grounded theory procedures and techniques. Sage.

Trafimow, D. (2023). Methodological issues in psychology: Concept, method, and measurement. Routledge. https://doi.org/10.4324/9781003365167