Diffusion Models of Corporate Geometry using Kan Extension Transformers

Sridhar Mahadevan, Adobe Research and U.Mass, Amherst

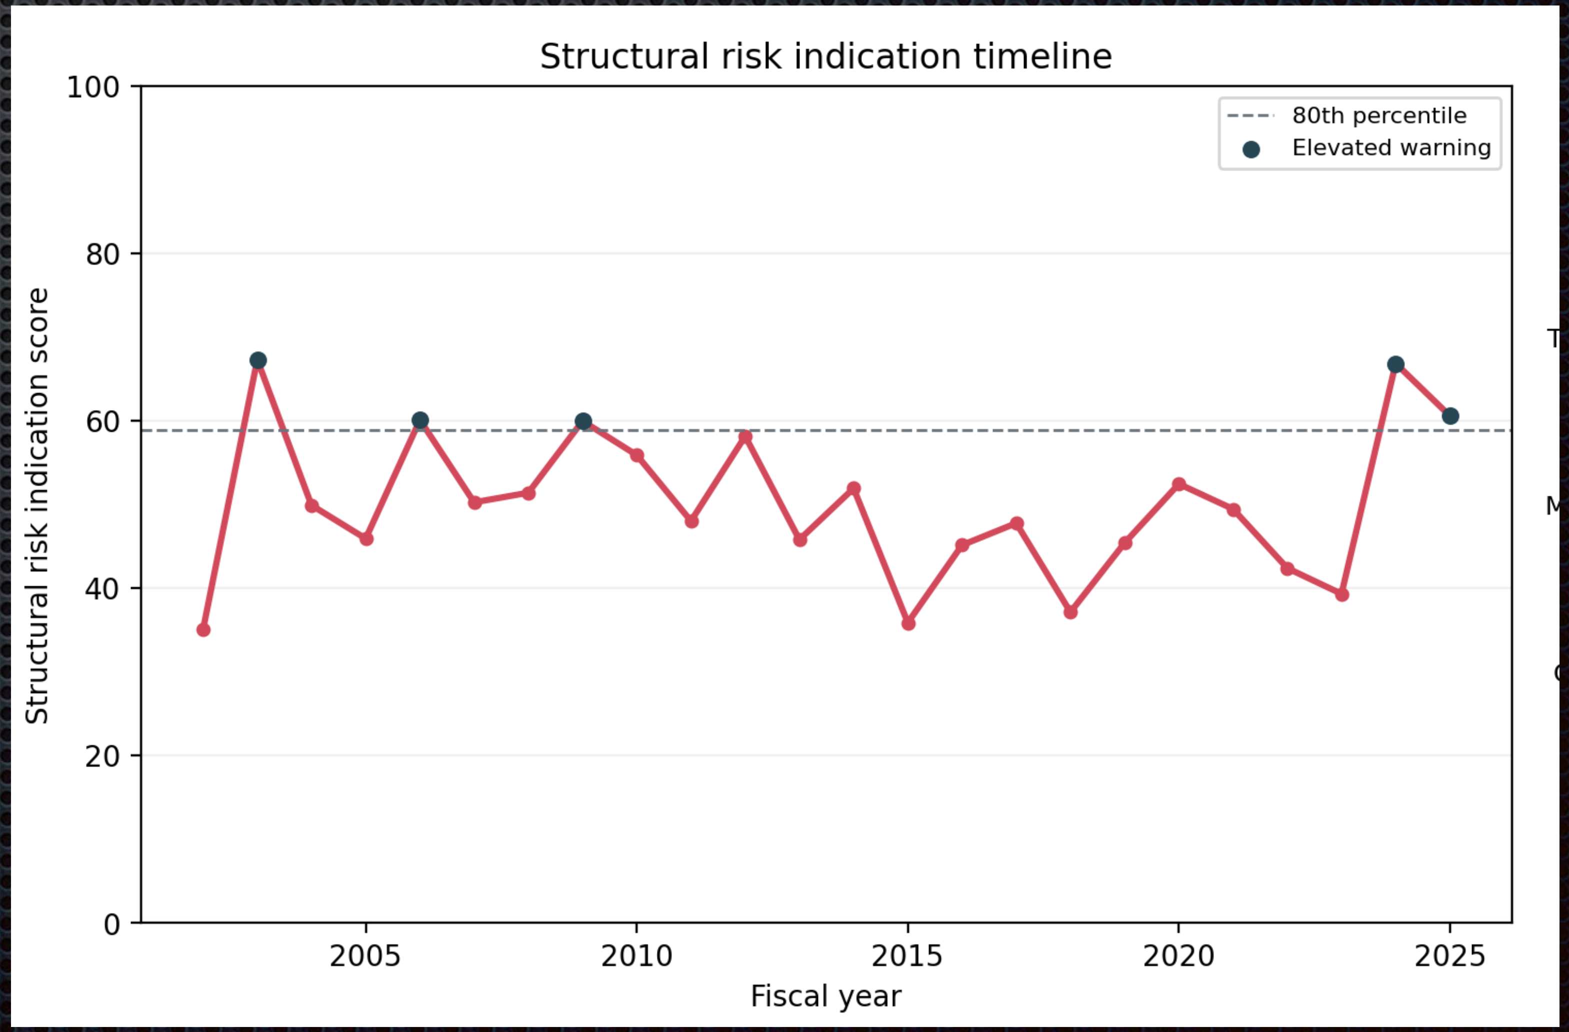

A Grand Challenge Moonshot

- Combine vast amounts of data over several decades
  - Financial reports (yearly 10K filings — each ~100 pages)
  - Social media buzz
  - News articles, press releases, analyst reports,…
- Challenge: how to weave together all this information into a coherent narrative?
  - Can we predict the “health” of a company by giving it an “annual physical”?
  - Can we reveal surprising similarities that lie “hidden” in latent manifolds?

Which companies are more alike?

Adobe

NIKE

NVIDIA

IBM

Brand-Aware Democritus Architecture

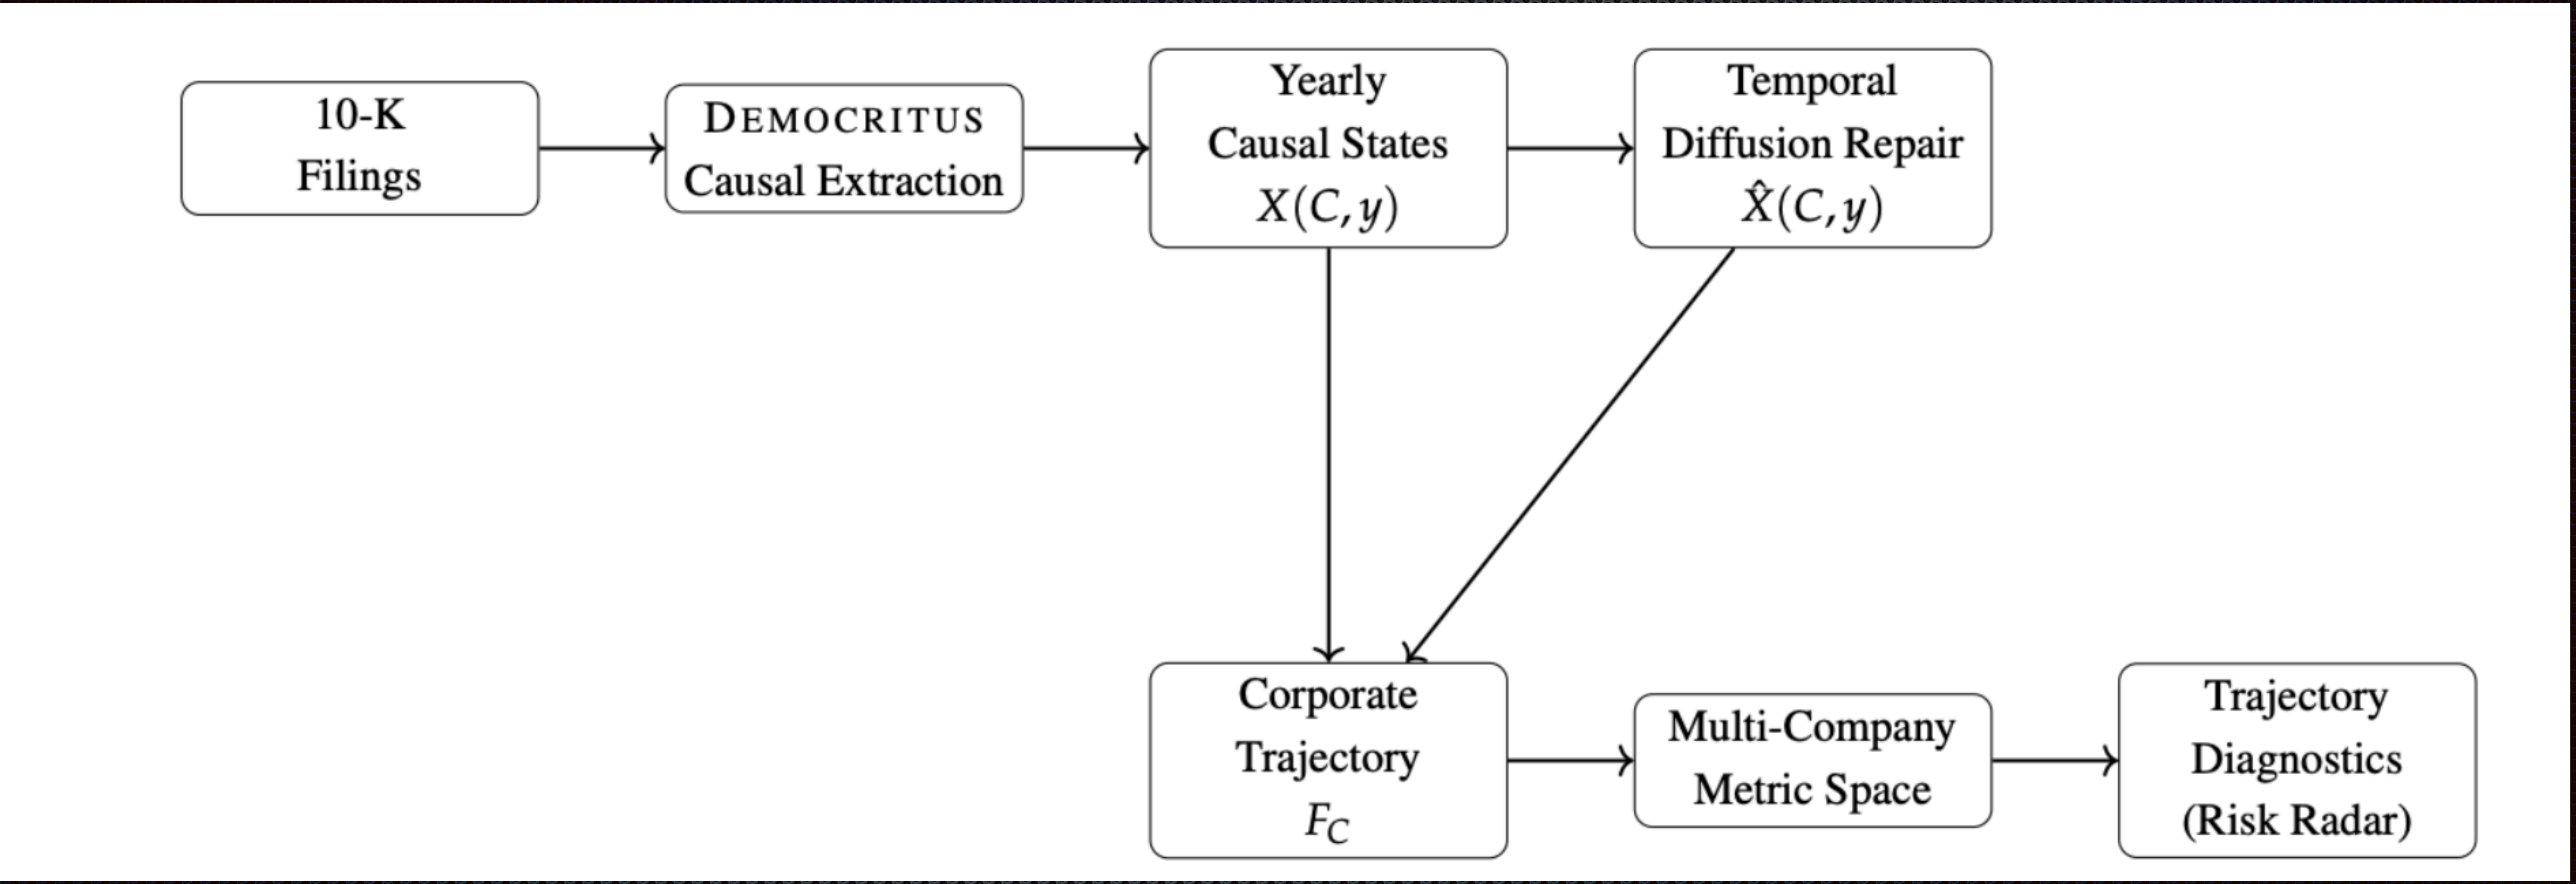

[Mahadevan, Categories for AGI , 2026]

(Static Single-Year) Democritus

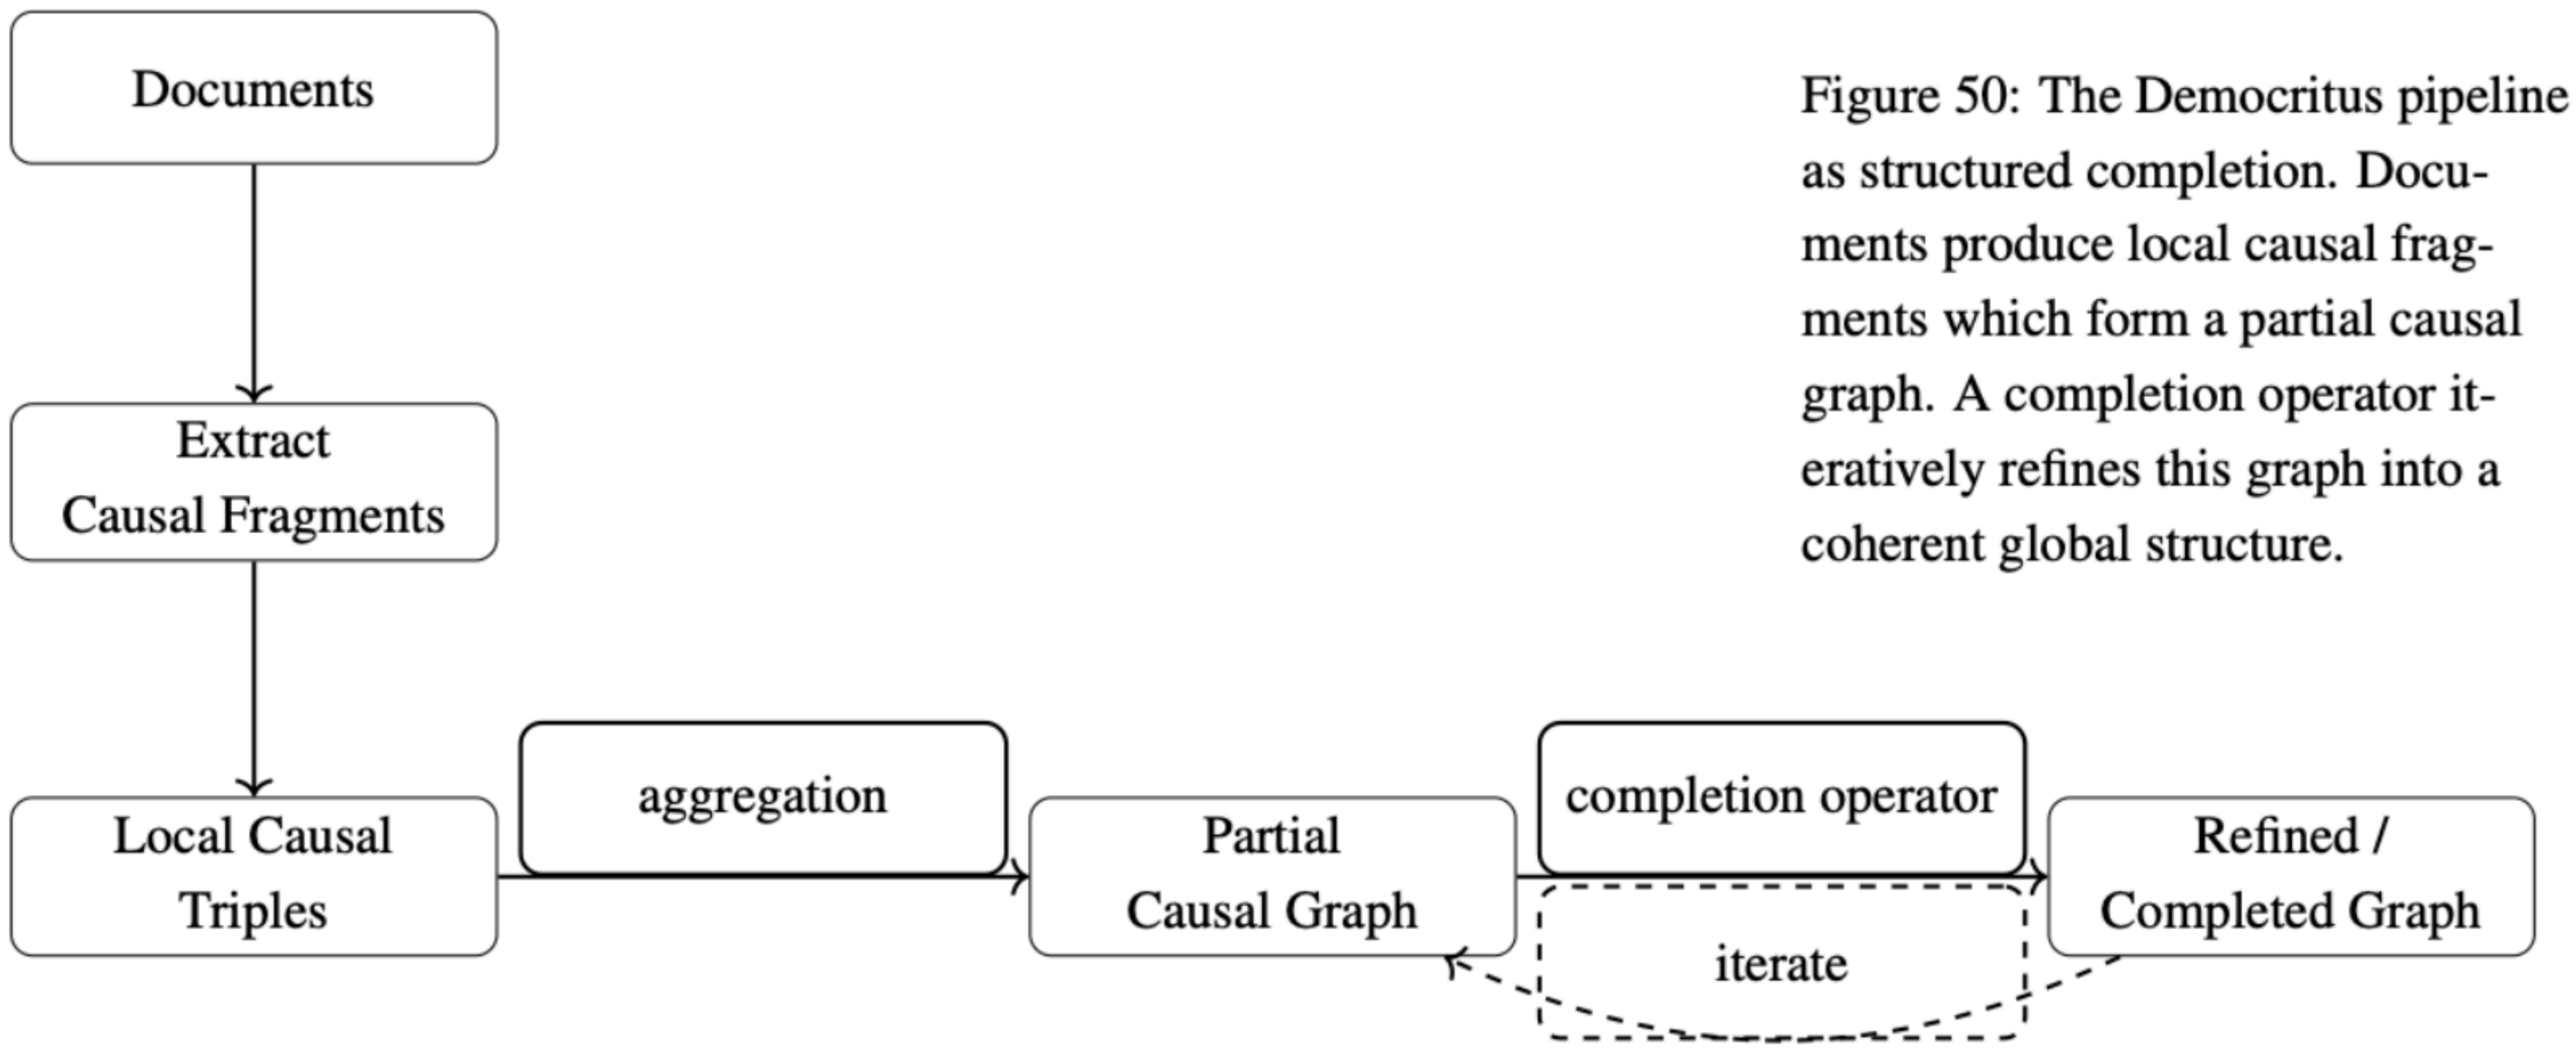

Figure 50: The Democritus pipeline as structured completion. Documents produce local causal fragments which form a partial causal graph. A completion operator iteratively refines this graph into a coherent global structure.

[Mahadevan, Categories for AGI , 2026]

Brand Financial Performance Perception

Brand Revenue Growth Narrative

Brand Profitability and Margin Perception

Brand Demand and Consumer Momentum

Product Innovation and Pipeline Credibility

Pricing Pipeline and Value Capture

Digital Channel and Value Capture

Market Share and Competitive Positioning

Supply Chain and Inventory Management

"question": "To what extent does pricing strategy in key geographies affect demand and operating margin trends disclosed in the 10-K?”

"statement": "Aggressive pricing strategy in key geographies reduces operating margin trends due to increased promotional activity, particularly in high-cost regions.”

"subj": "aggressive pricing strategy in key geographies”
"rel": “reduces"
"obj": "operating margin trends due to increased promotional activity, particularly in high-cost regions”
"domain": "brand financial performance perception"

Brand Profile

Adobe 2002-2025

Brand Profile

Nike 2002-2025

Diffusion of causal relations over time

Discrete Schrödinger Bridges

- Schrödinger bridge problem: given a stochastic process describing random dynamics, what is the most likely process that connects two observed distributions, while remaining as close as possible to the original dynamics?
- Mathematically, solve for $P^\star = \arg \min_{P:B_\mu(\mu_0,\mu_T)} \mathrm{KL}(P\|R)$.
- In our BAD application, each company defines a noisy trajectory through a high-dimensional latent space, and the goal is to find a smooth path

Causal Reasoning in Coend and Kan Extension Transformers (CRICKET)

Model	Attention Geometry
Transformer	Dot Product Similarity
Geometric Transformer	Learned Geometric Convolution Kernel
Kan Extension Transformer	Learned Morphism Kernel
TopoCoend Transformer	Geometric Neighborhood Kernel

[Mahadevan, Categories for AGI, 2026]

PLAYING BAD CRICKET

Companies as Diffusion Functors

- Each company is modeled as a diffusion functor:
  - $F_{\text{Company}} : \mathbf{Year} \to \mathbf{CausalState}$
- Given a noisy temporal block of company causal states, reconstruct a cleaner and more coherent local trajectory
  - For any company C, and years $y_1, \ldots, y_k$ derived from
    - Raw causal relational triples (extracted from company 10K reports)
    - Denoised triples
    - Yearly causal atlas
    - Edge-basis and latent state
  - Denoising reconstructs a cleaner yearly state and local temporal consistency across states

nVidia Diffusion Model

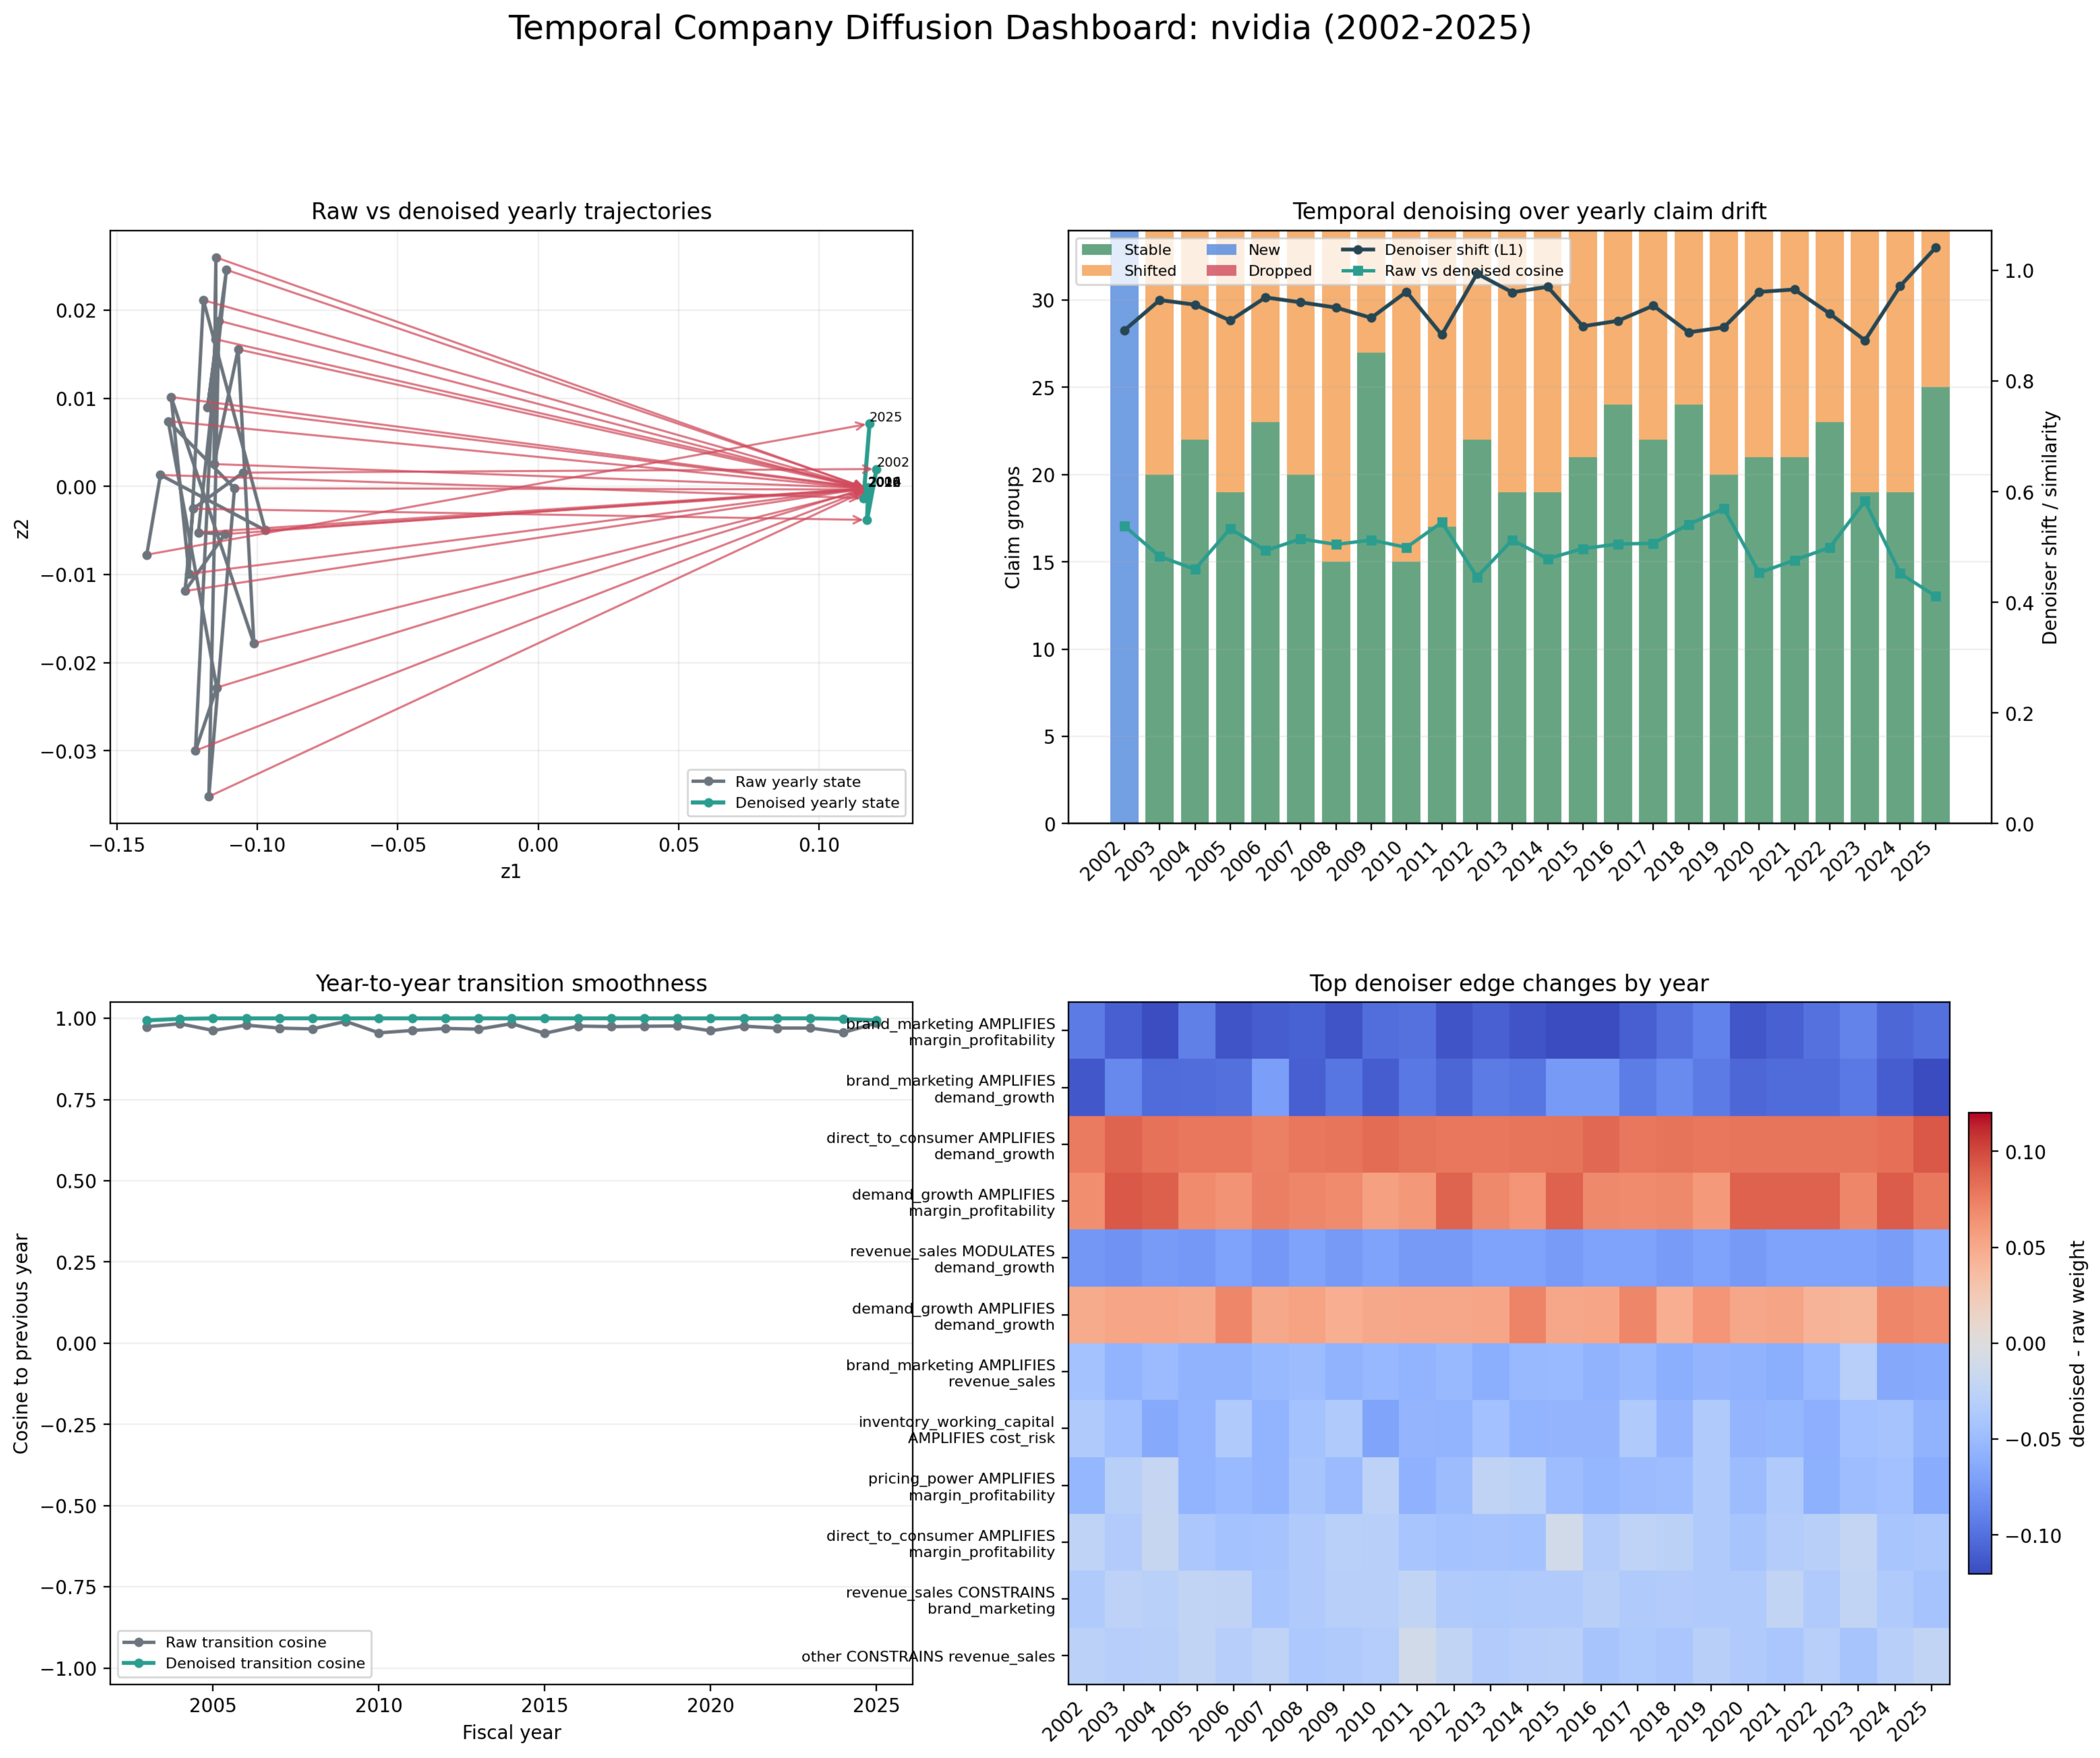

Adobe Diffusion Model

nVidia vs IBM

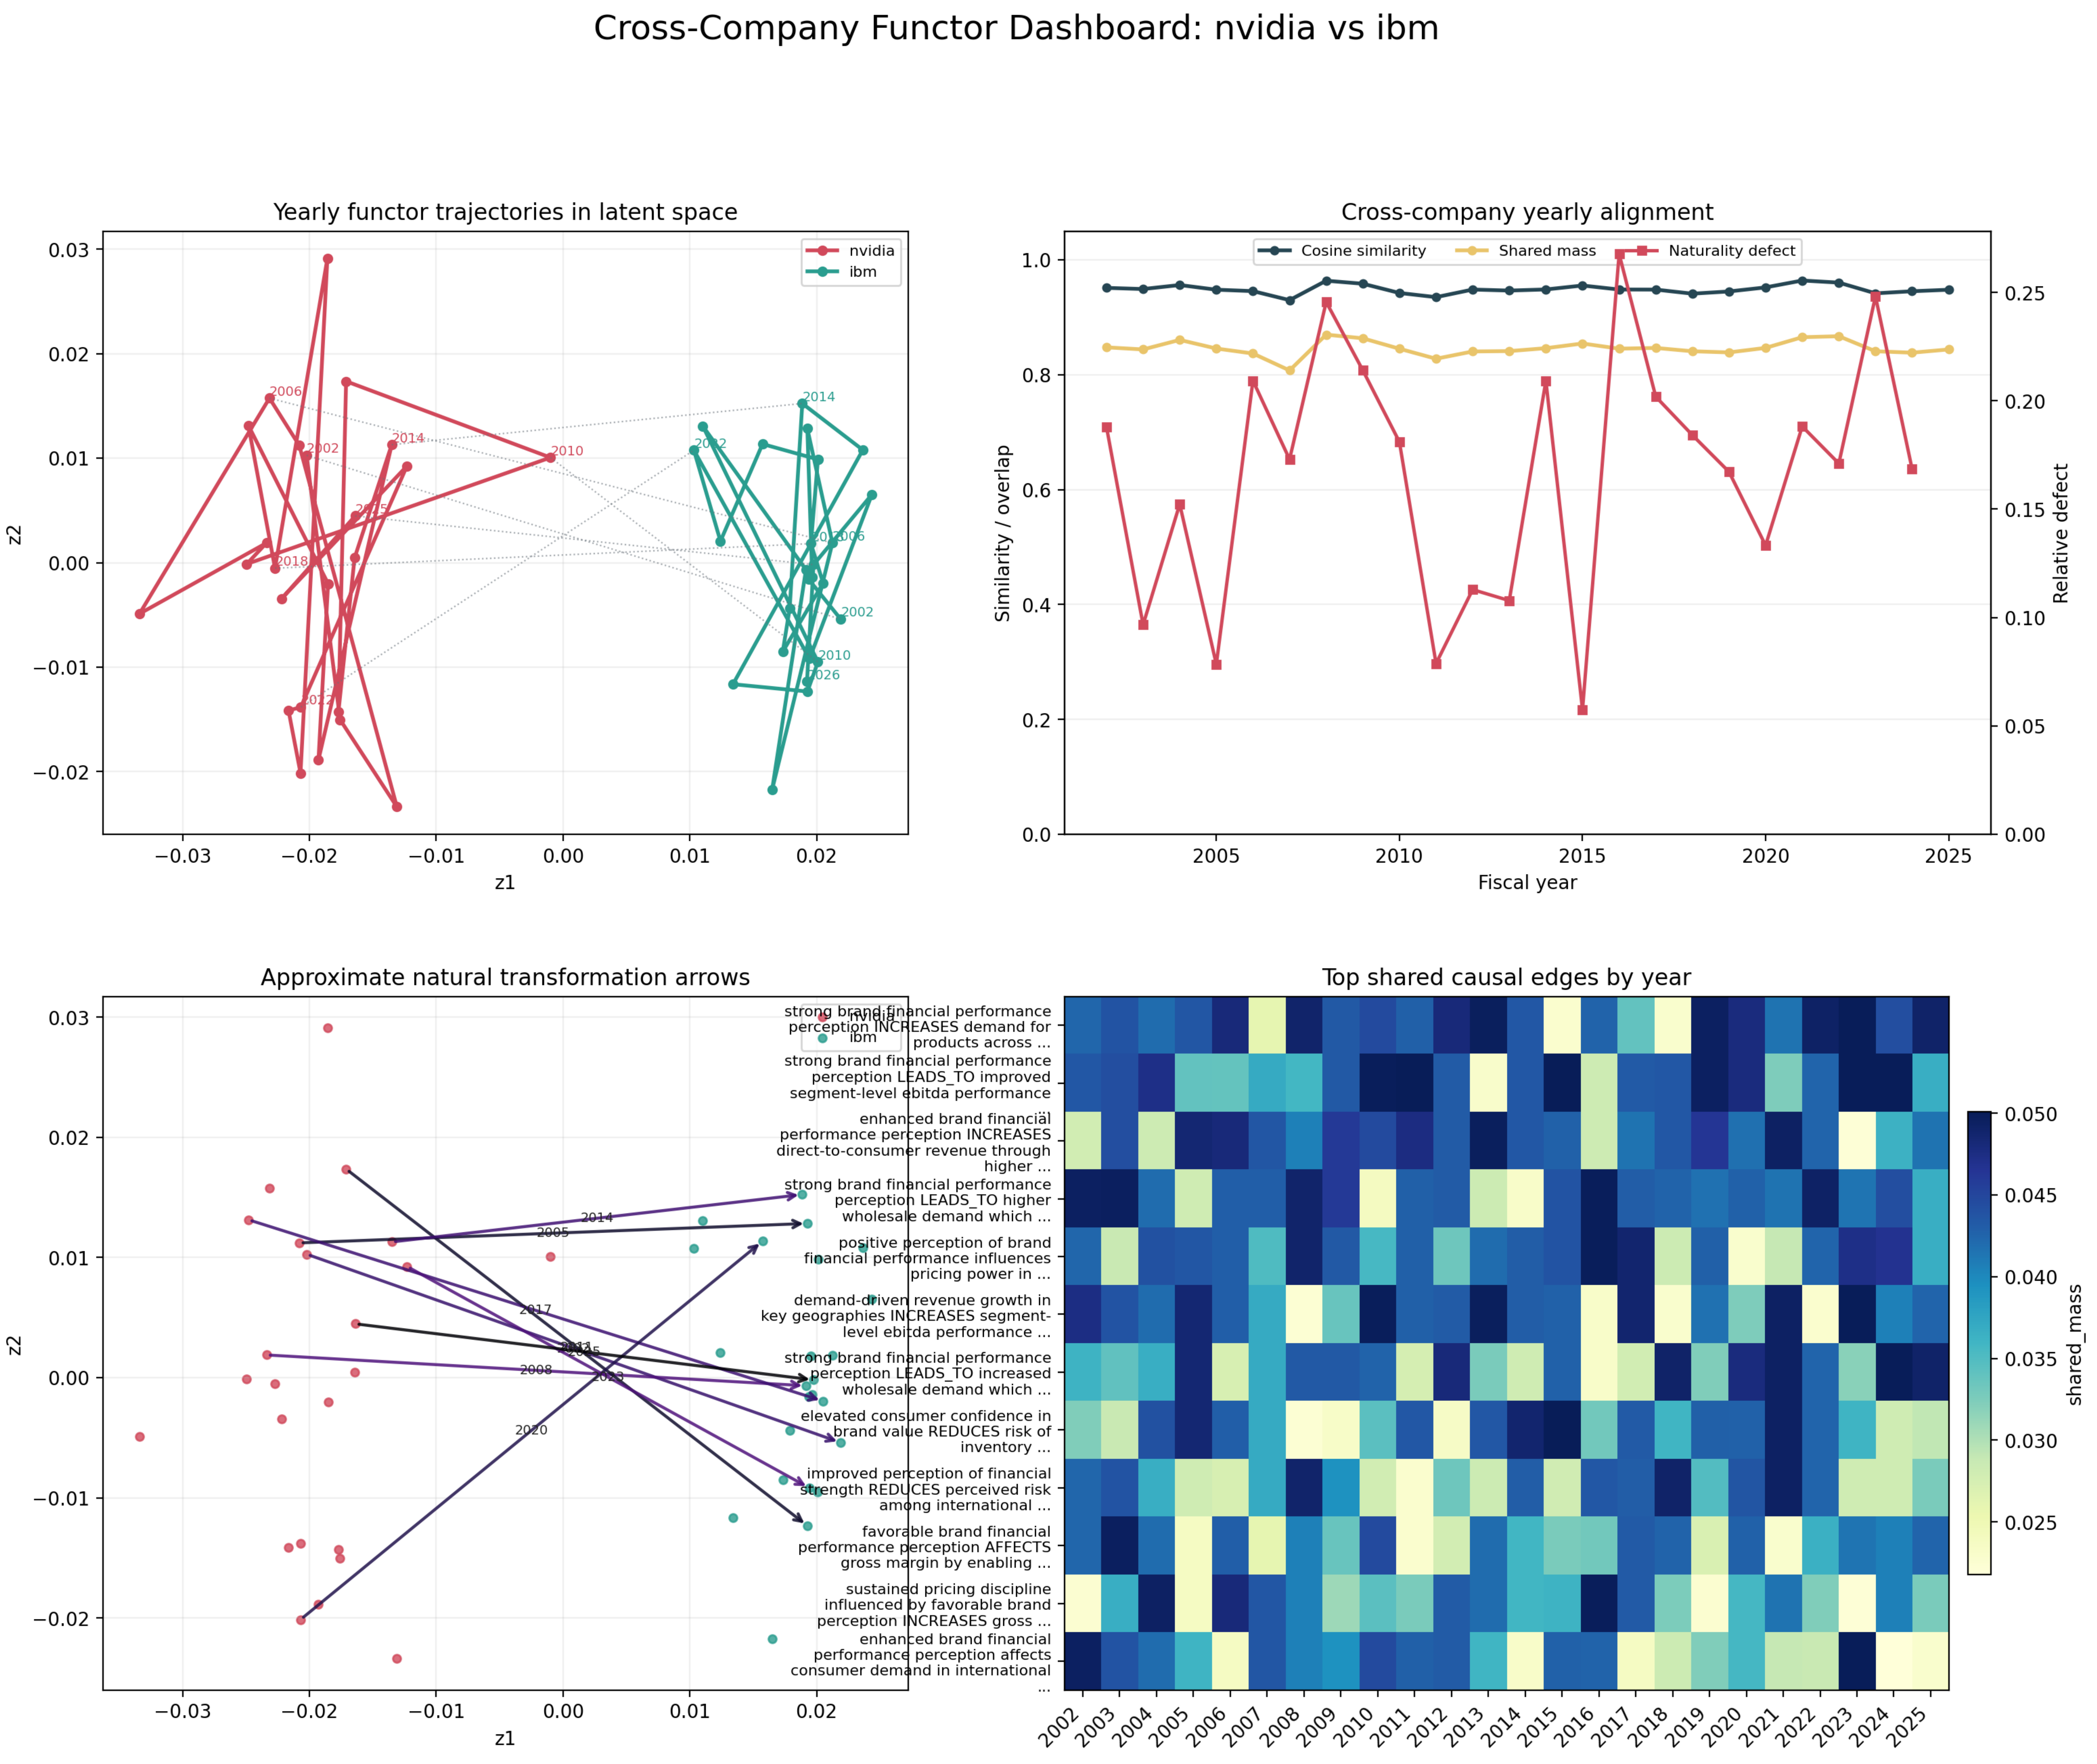

Apple vs. nVidia

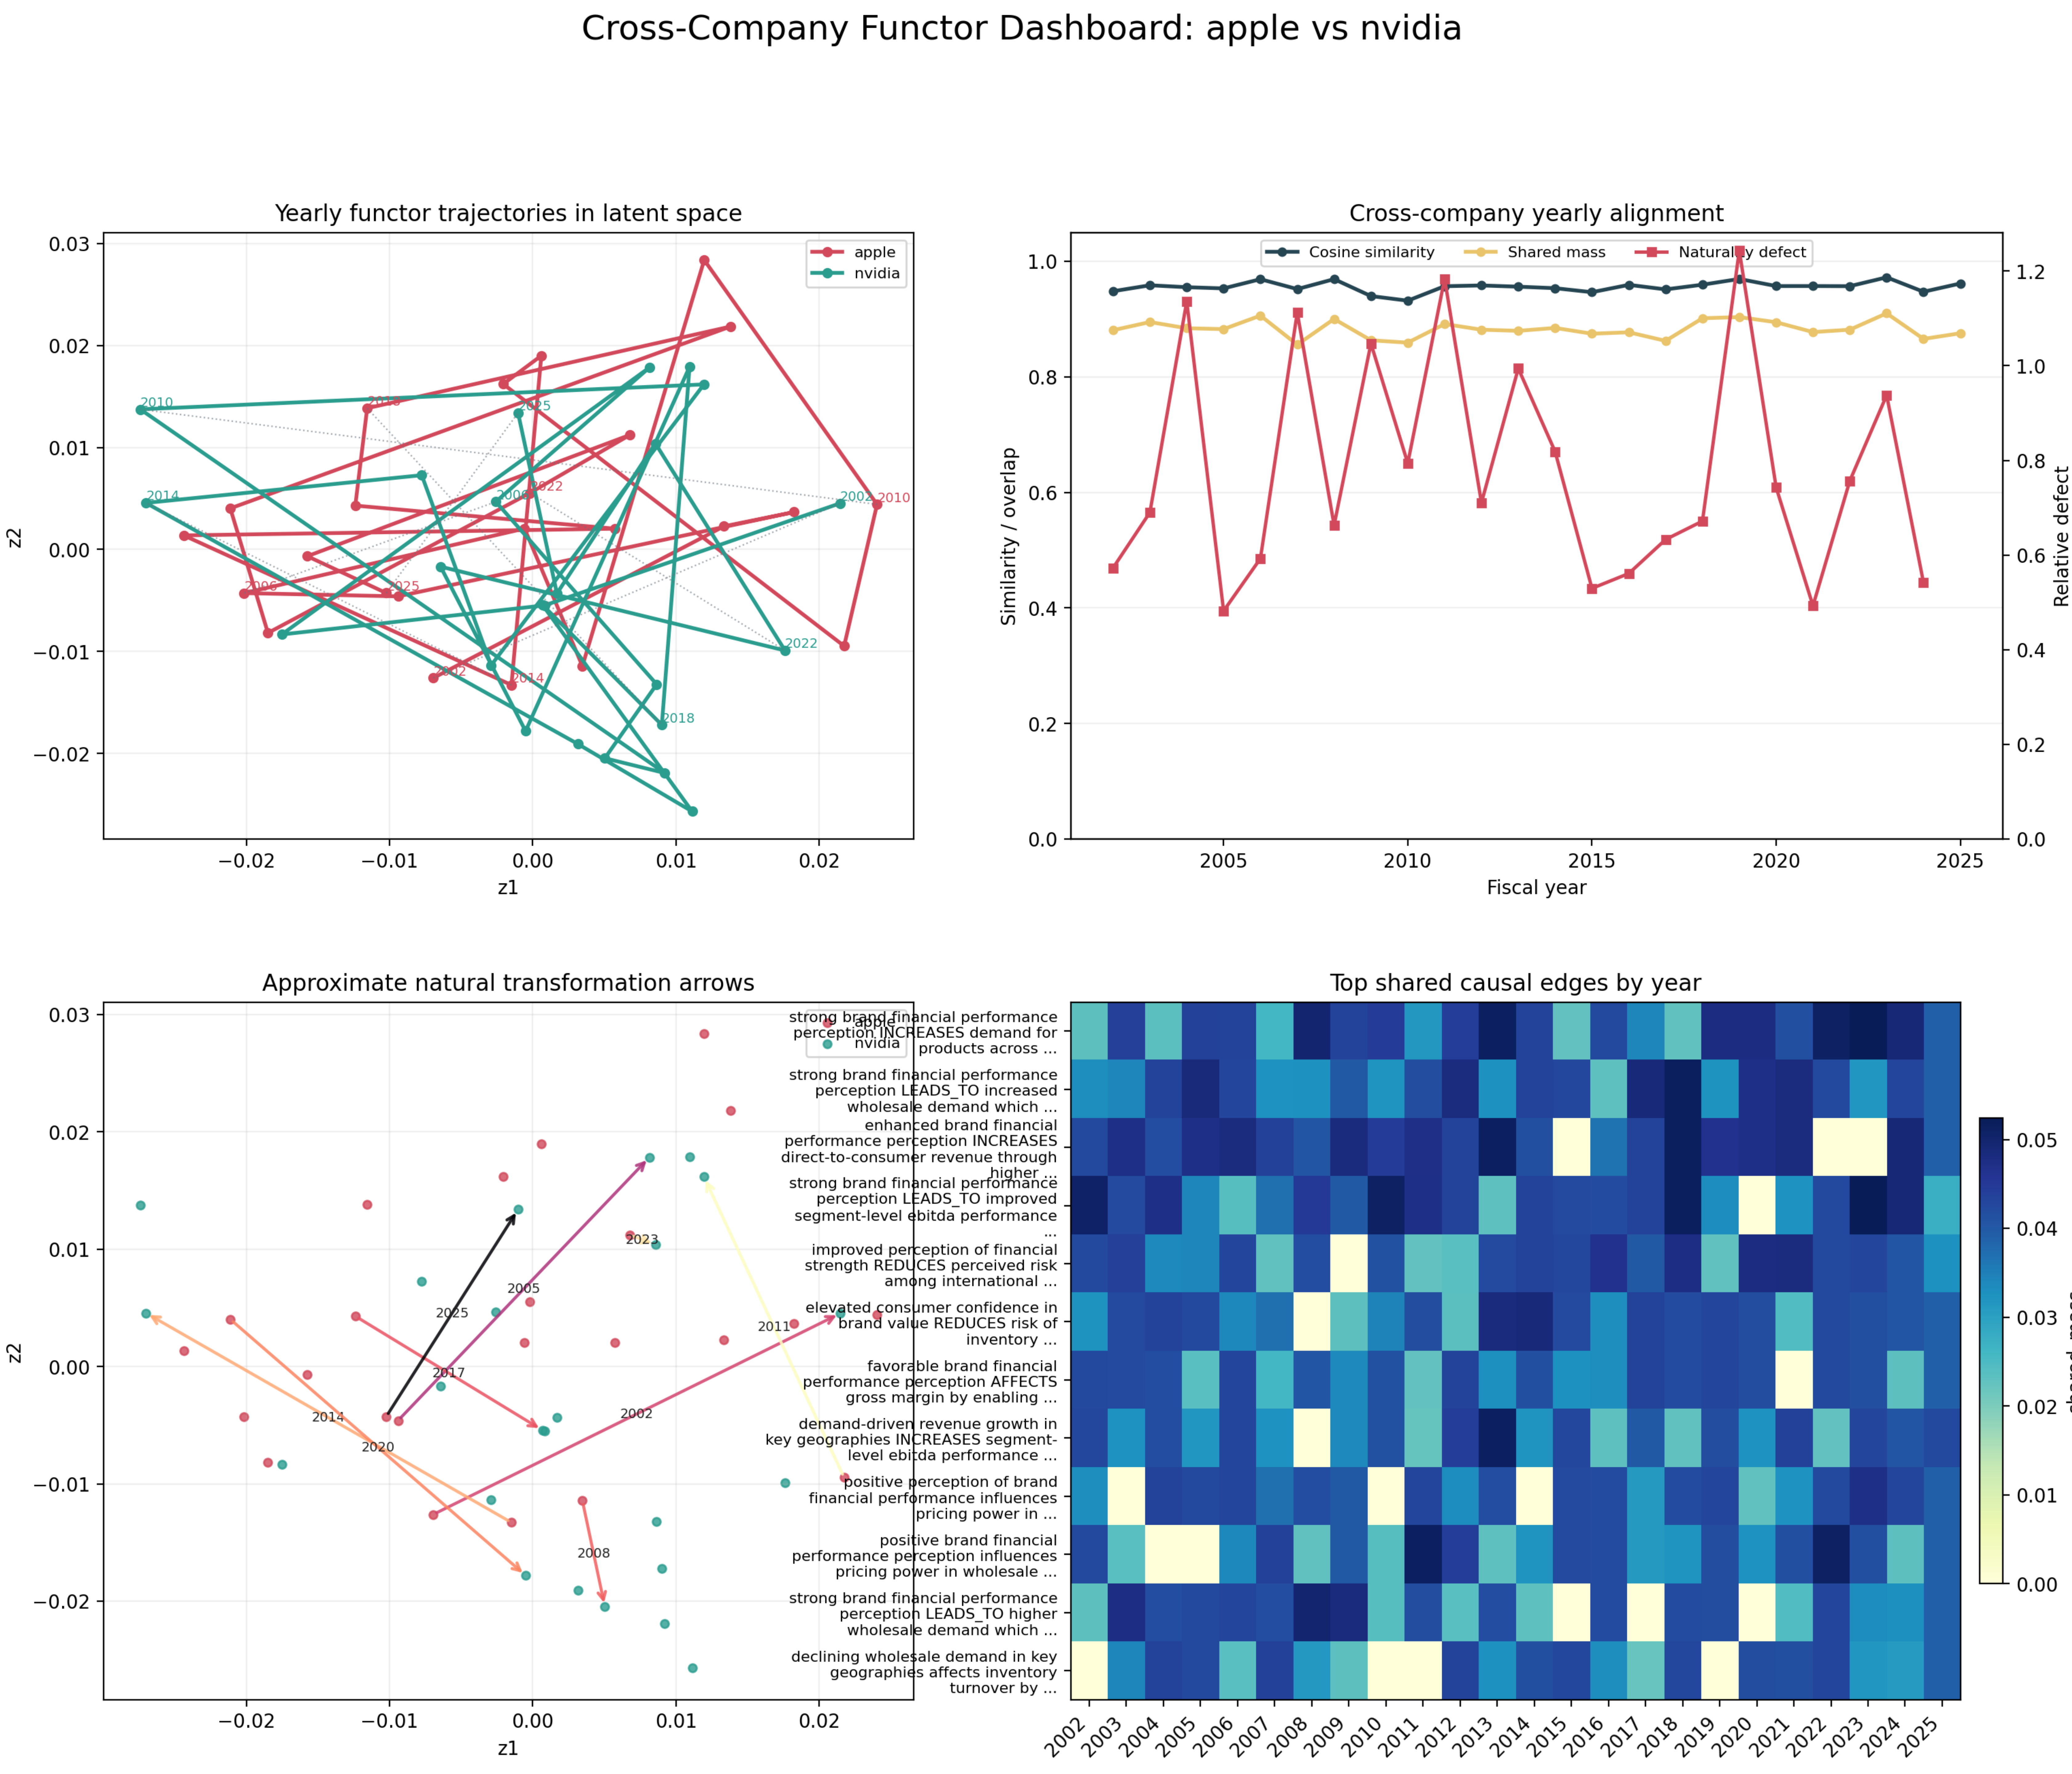

Nike vs. Adobe

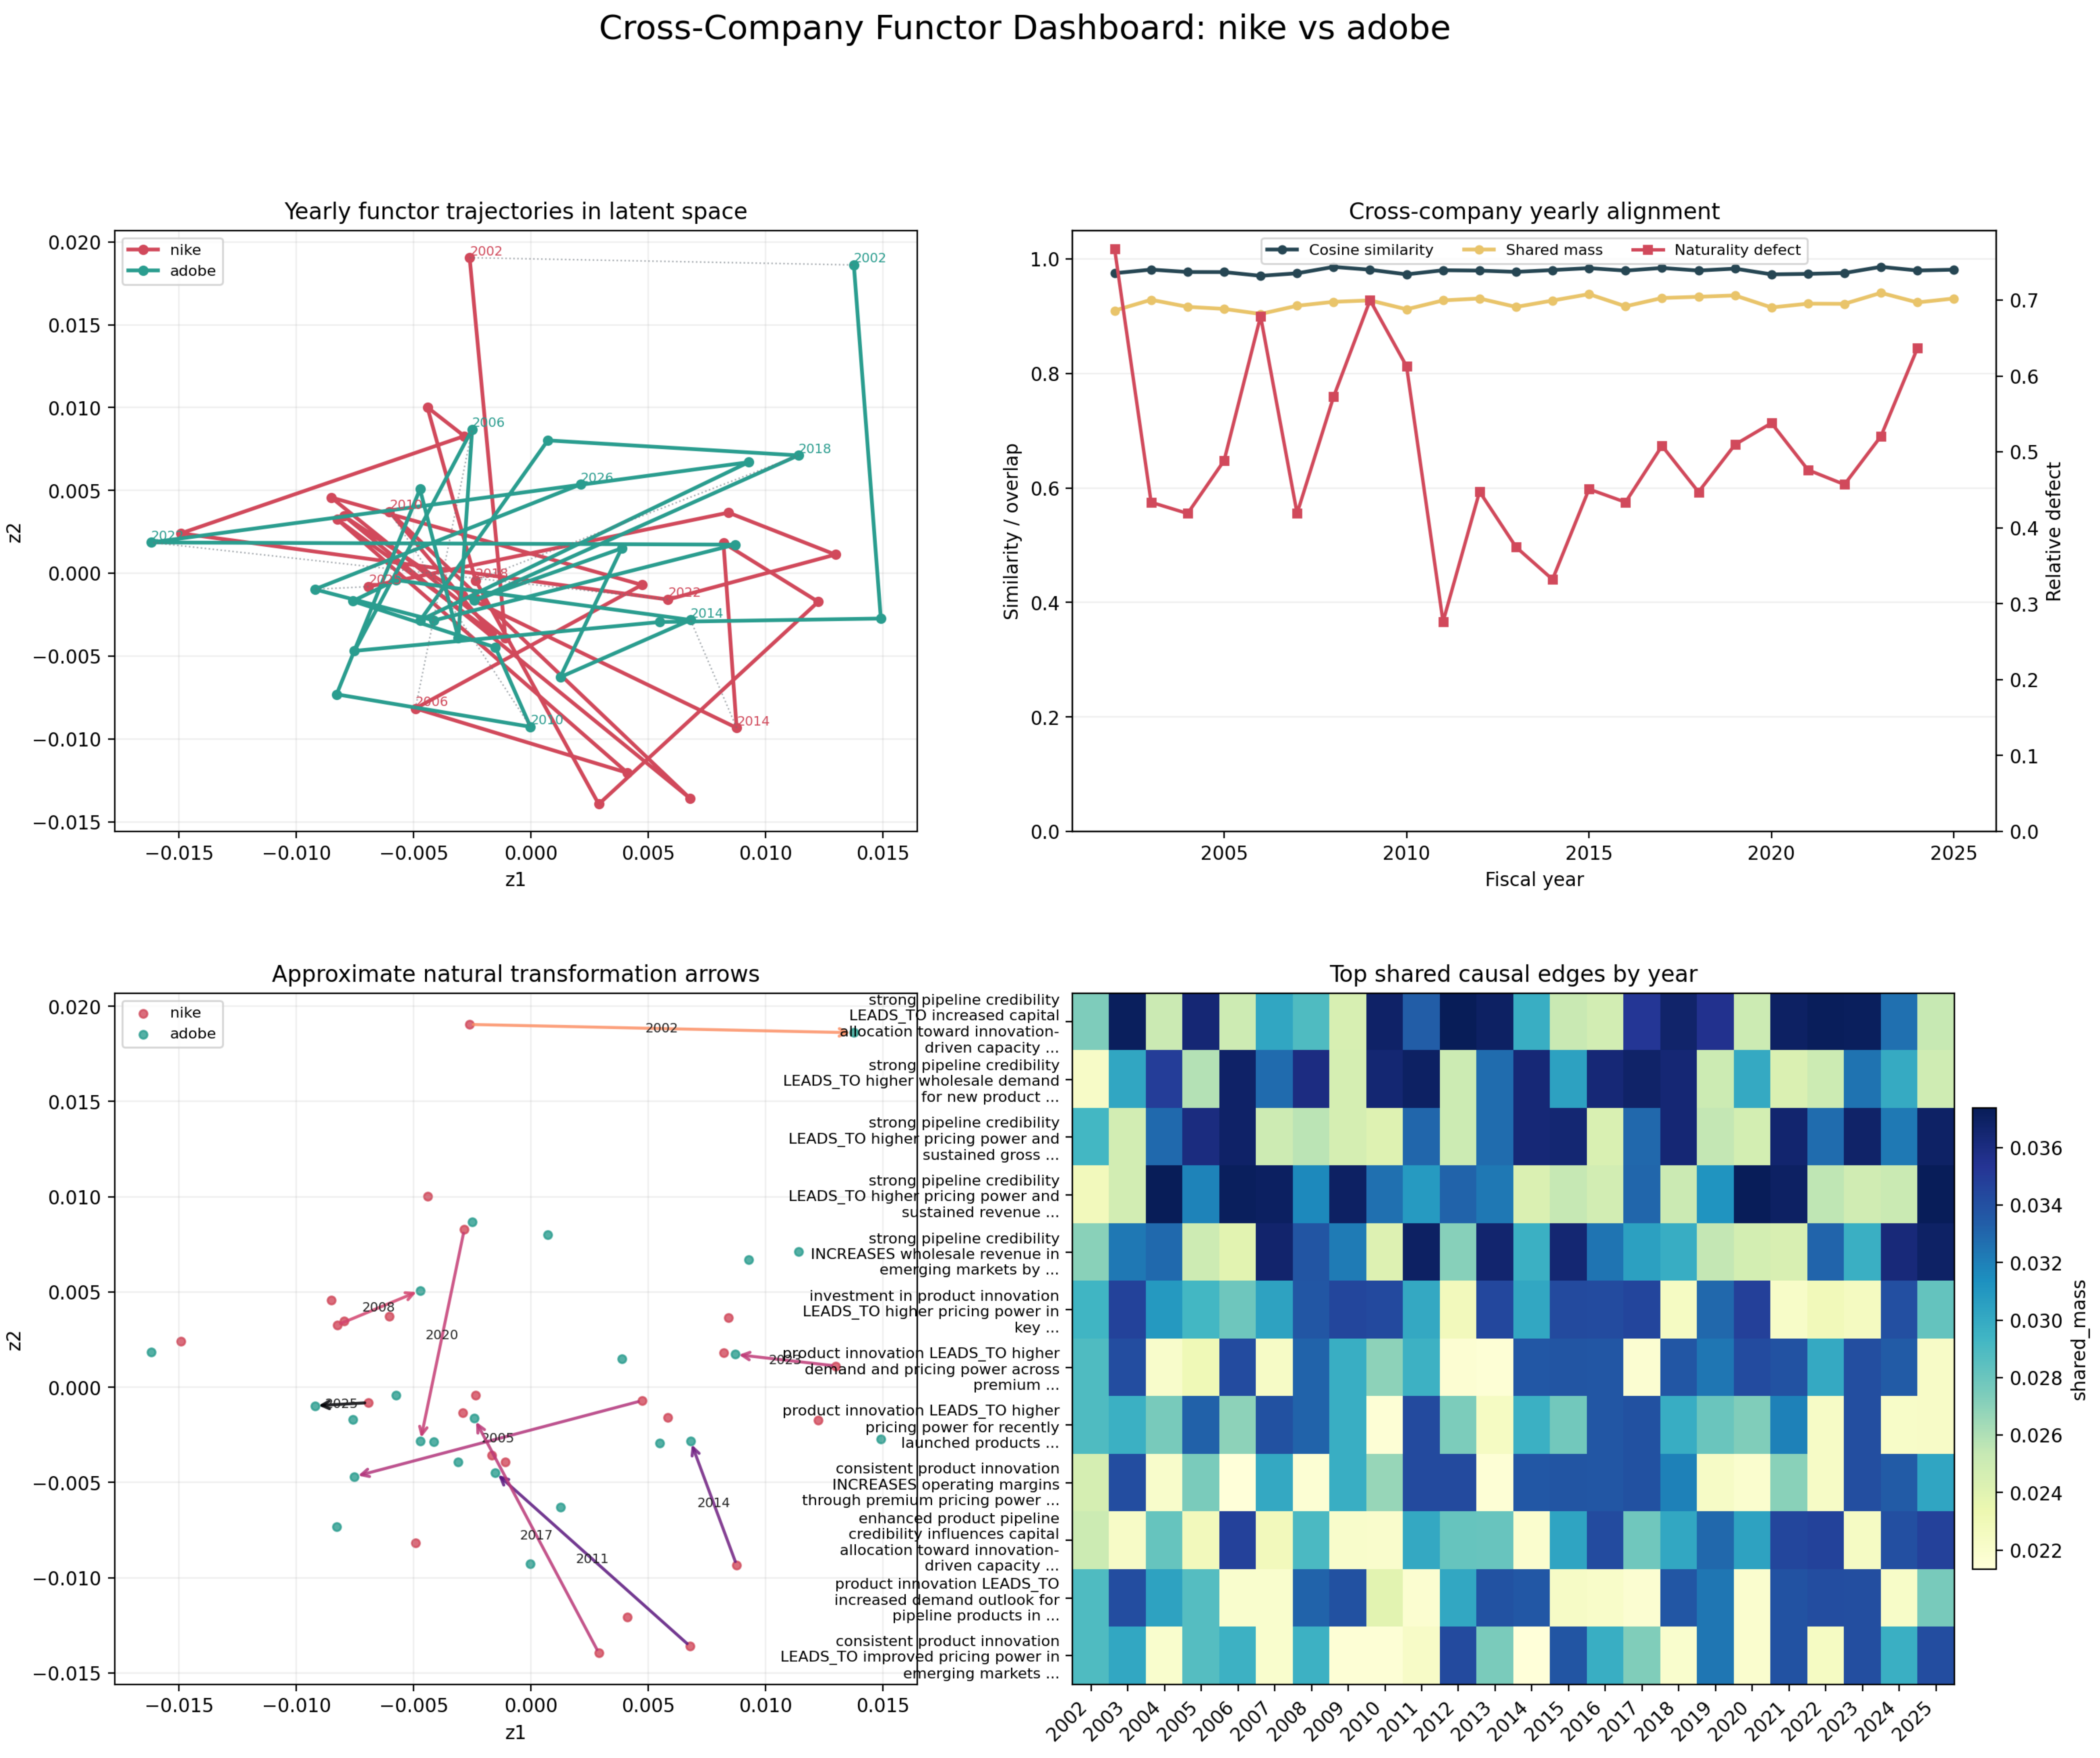

Structural Risk Assessment

- The score is constructed from several components derived from the repaired causal states:
  - Instability: magnitude of the denoising correction applied to the yearly causal state.
  - Transition break: deviation between consecutive yearly transitions.
  - Margin pressure: accumulation of constraining relations associated with operational stress roles.
  - Growth fragility: erosion of motifs amplifying core growth targets such as demand growth and pricing power.
  - Opacity: concentration of causal mass or reliance on unspecified roles.
- These components are aggregated into a single radar score
  - $R(C, y) = w_I I(C, y) + w_T T(C, y) + w_M M(C, y) + w_G G(C, y) + w_O O(C, y)$

nVidia Movie

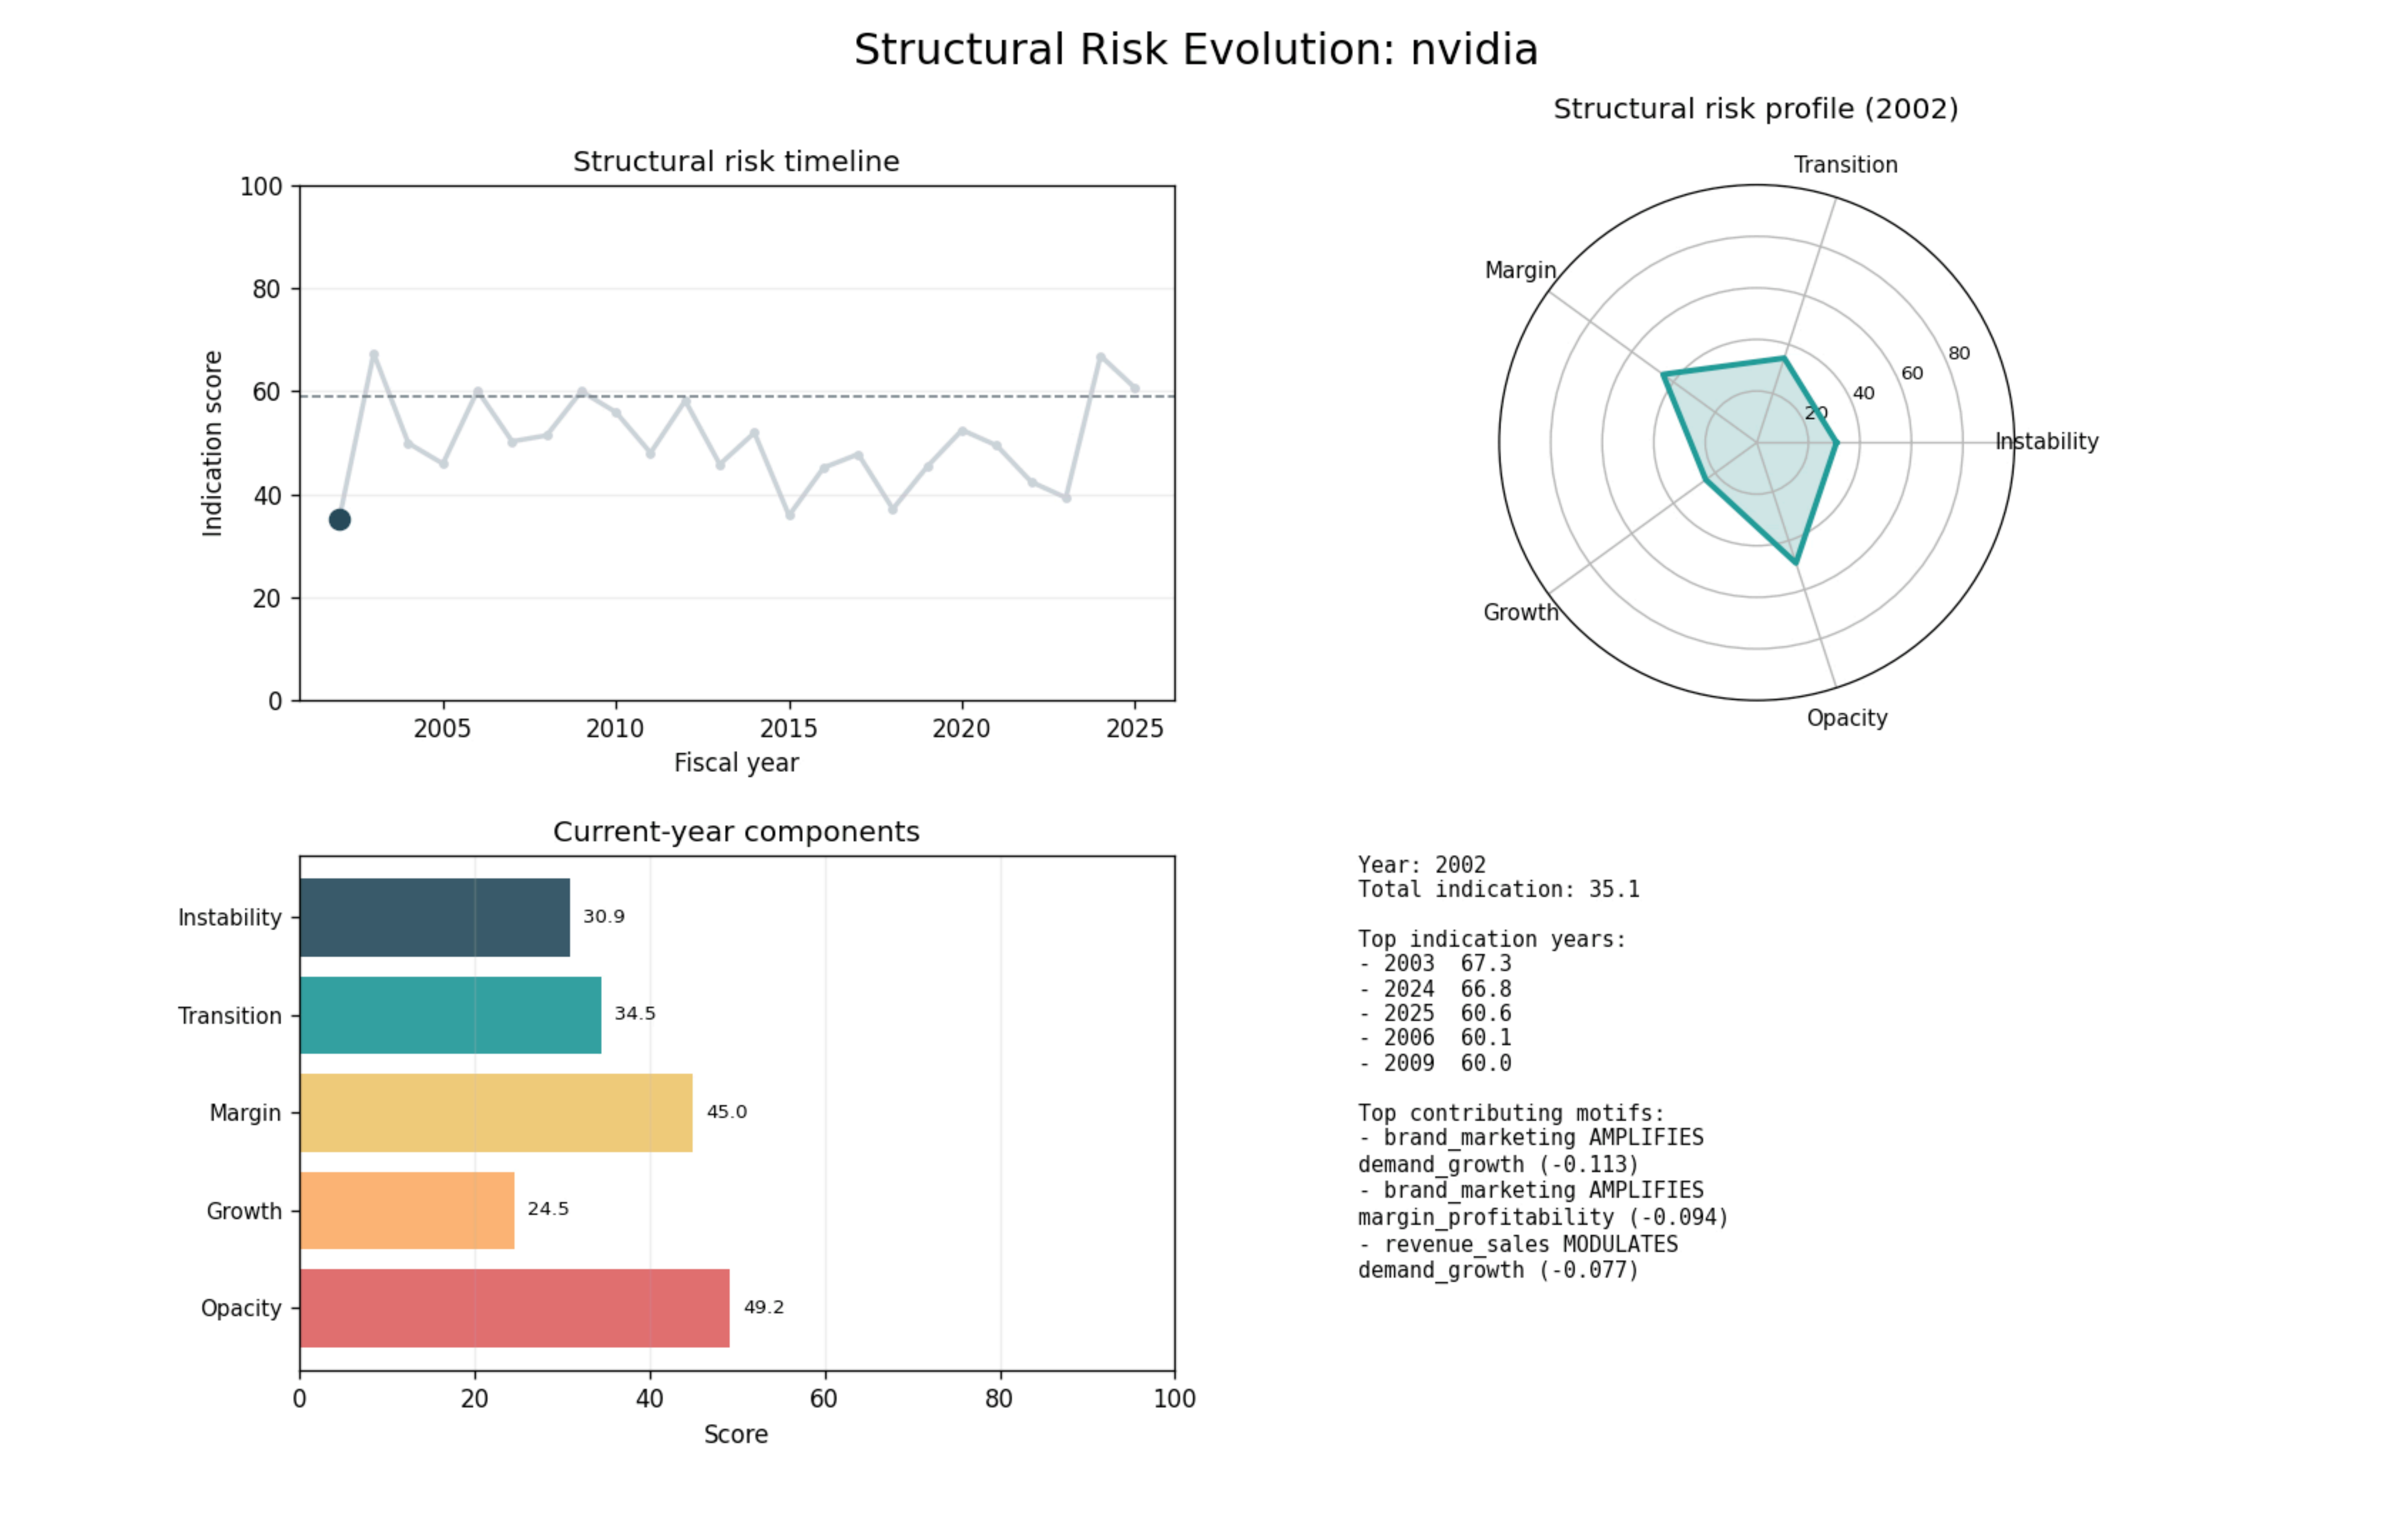

IBM Movie

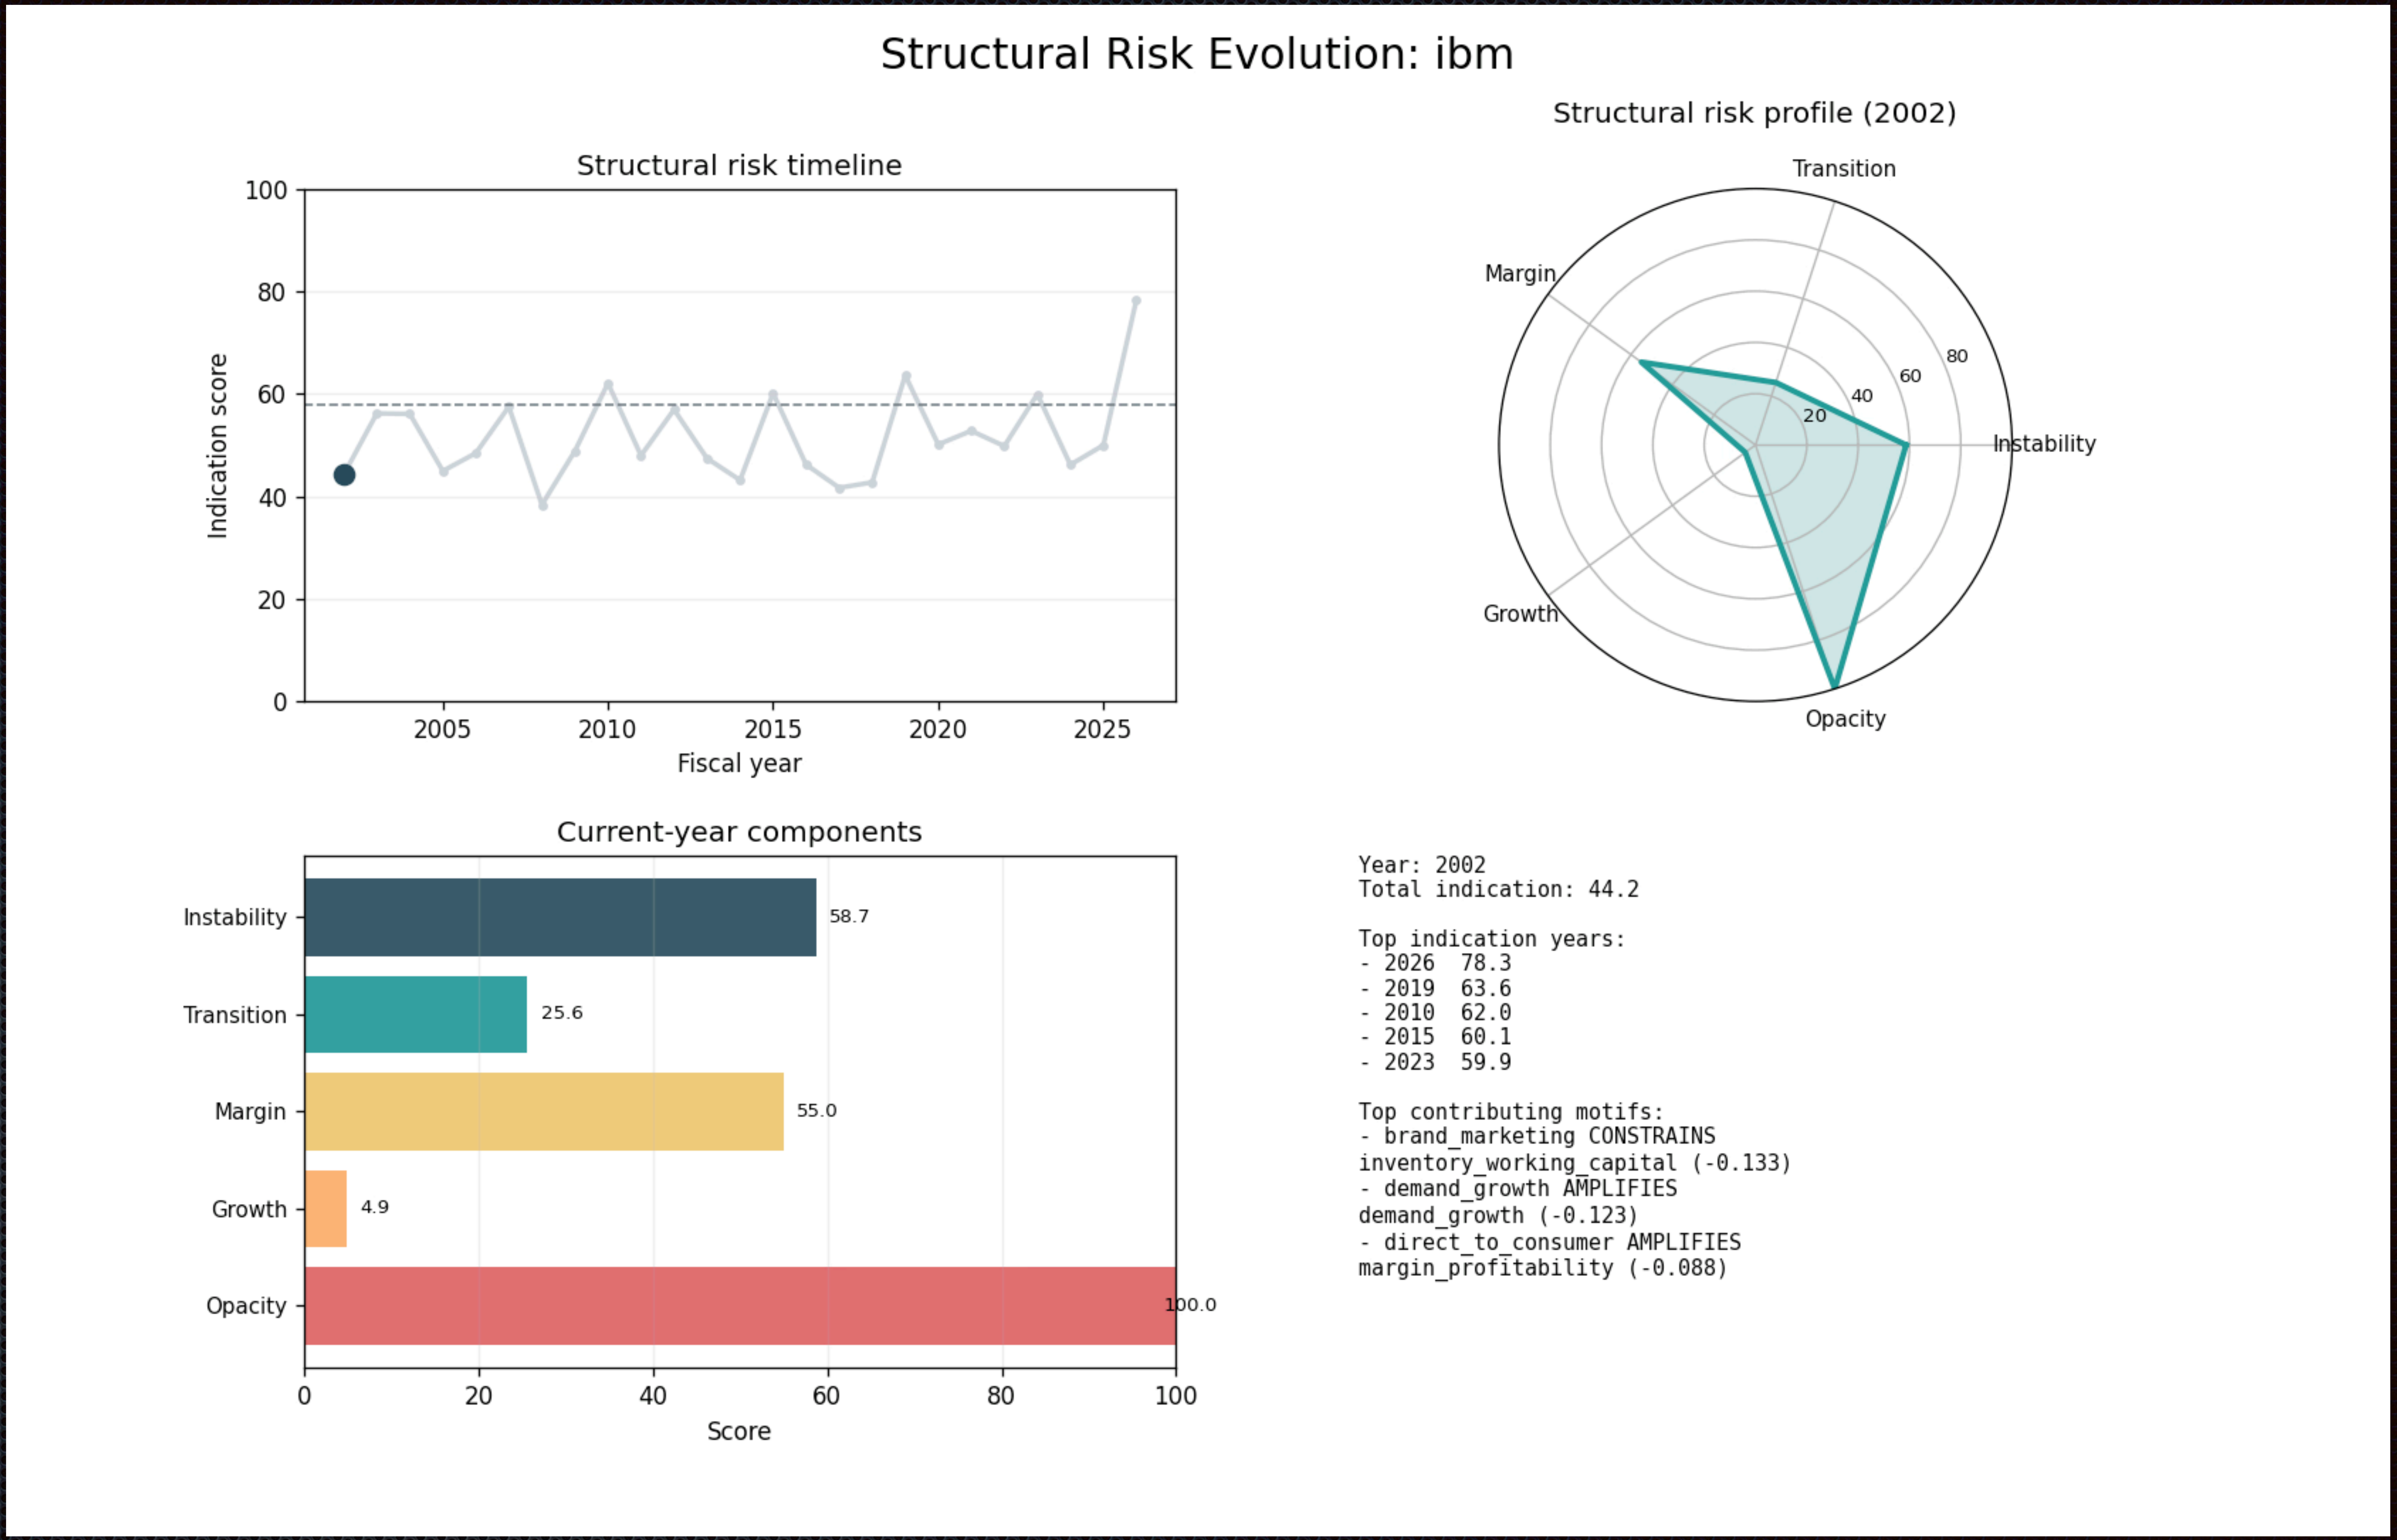
Structural Risk Evolution: ibm
Structural risk timeline
Indication score
Fiscal year
Structural risk profile (2002)
Transition
Margin
Instability
Growth
Opacity
Current-year components
Instability 58.7
Transition 25.6
Margin 55.0
Growth 4.9
Opacity 100.0
Score
Year: 2002
Total indication: 44.2
Top indication years:
- 2026 78.3
- 2019 63.6
- 2010 62.0
- 2015 60.1
- 2023 59.9
Top contributing motifs:
- brand_marketing CONSTRAINS inventory_working_capital (-0.133)
- demand_growth AMPLIFIES demand_growth (-0.123)
- direct_to_consumer AMPLIFIES margin_profitability (-0.088)

Nike
Movie

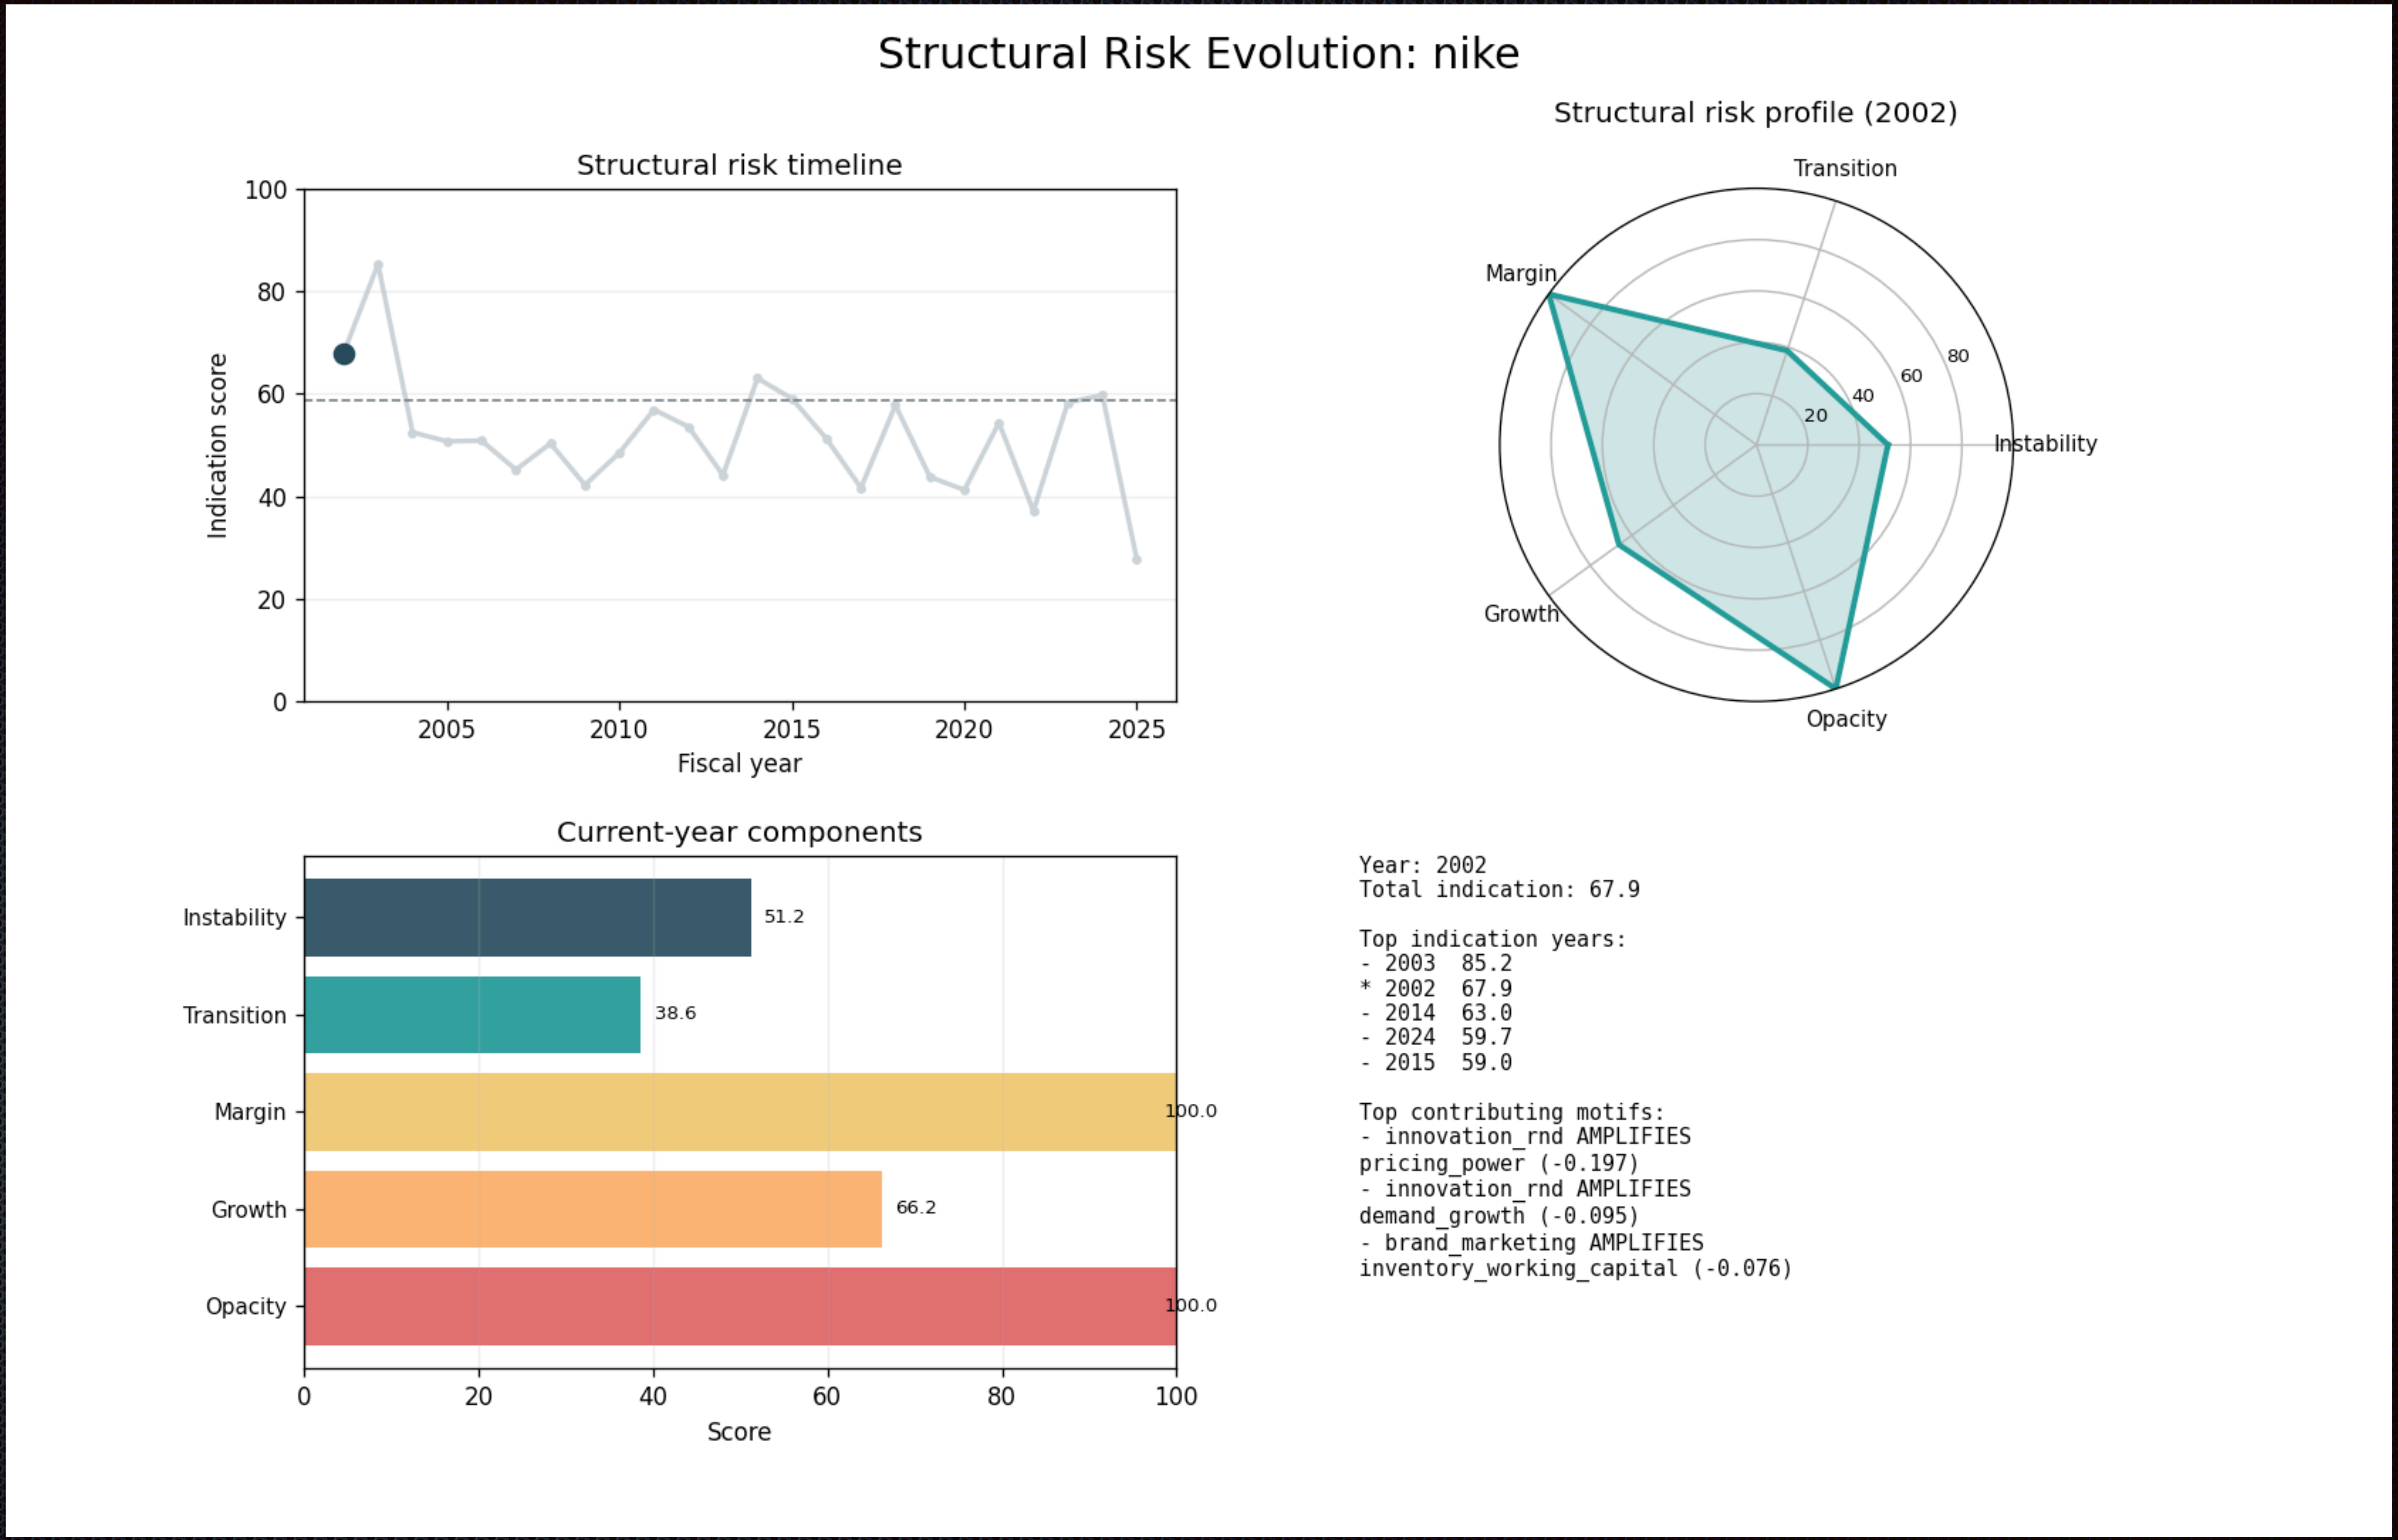
Structural Risk Evolution: nike
Structural risk timeline
Indication score
Fiscal year
Structural risk profile (2002)
Transition
Margin
Instability
Growth
Opacity
Current-year components
Instability
Transition
Margin
Growth
Opacity
51.2
38.6
100.0
66.2
100.0
Score
Year: 2002
Total indication: 67.9
Top indication years:
- 2003 85.2
* 2002 67.9
- 2014 63.0
- 2024 59.7
- 2015 59.0
Top contributing motifs:
- innovation_rnd AMPLIFIES pricing_power (-0.197)
- innovation_rnd AMPLIFIES demand_growth (-0.095)
- brand_marketing AMPLIFIES inventory_working_capital (-0.076)

Apple Movie

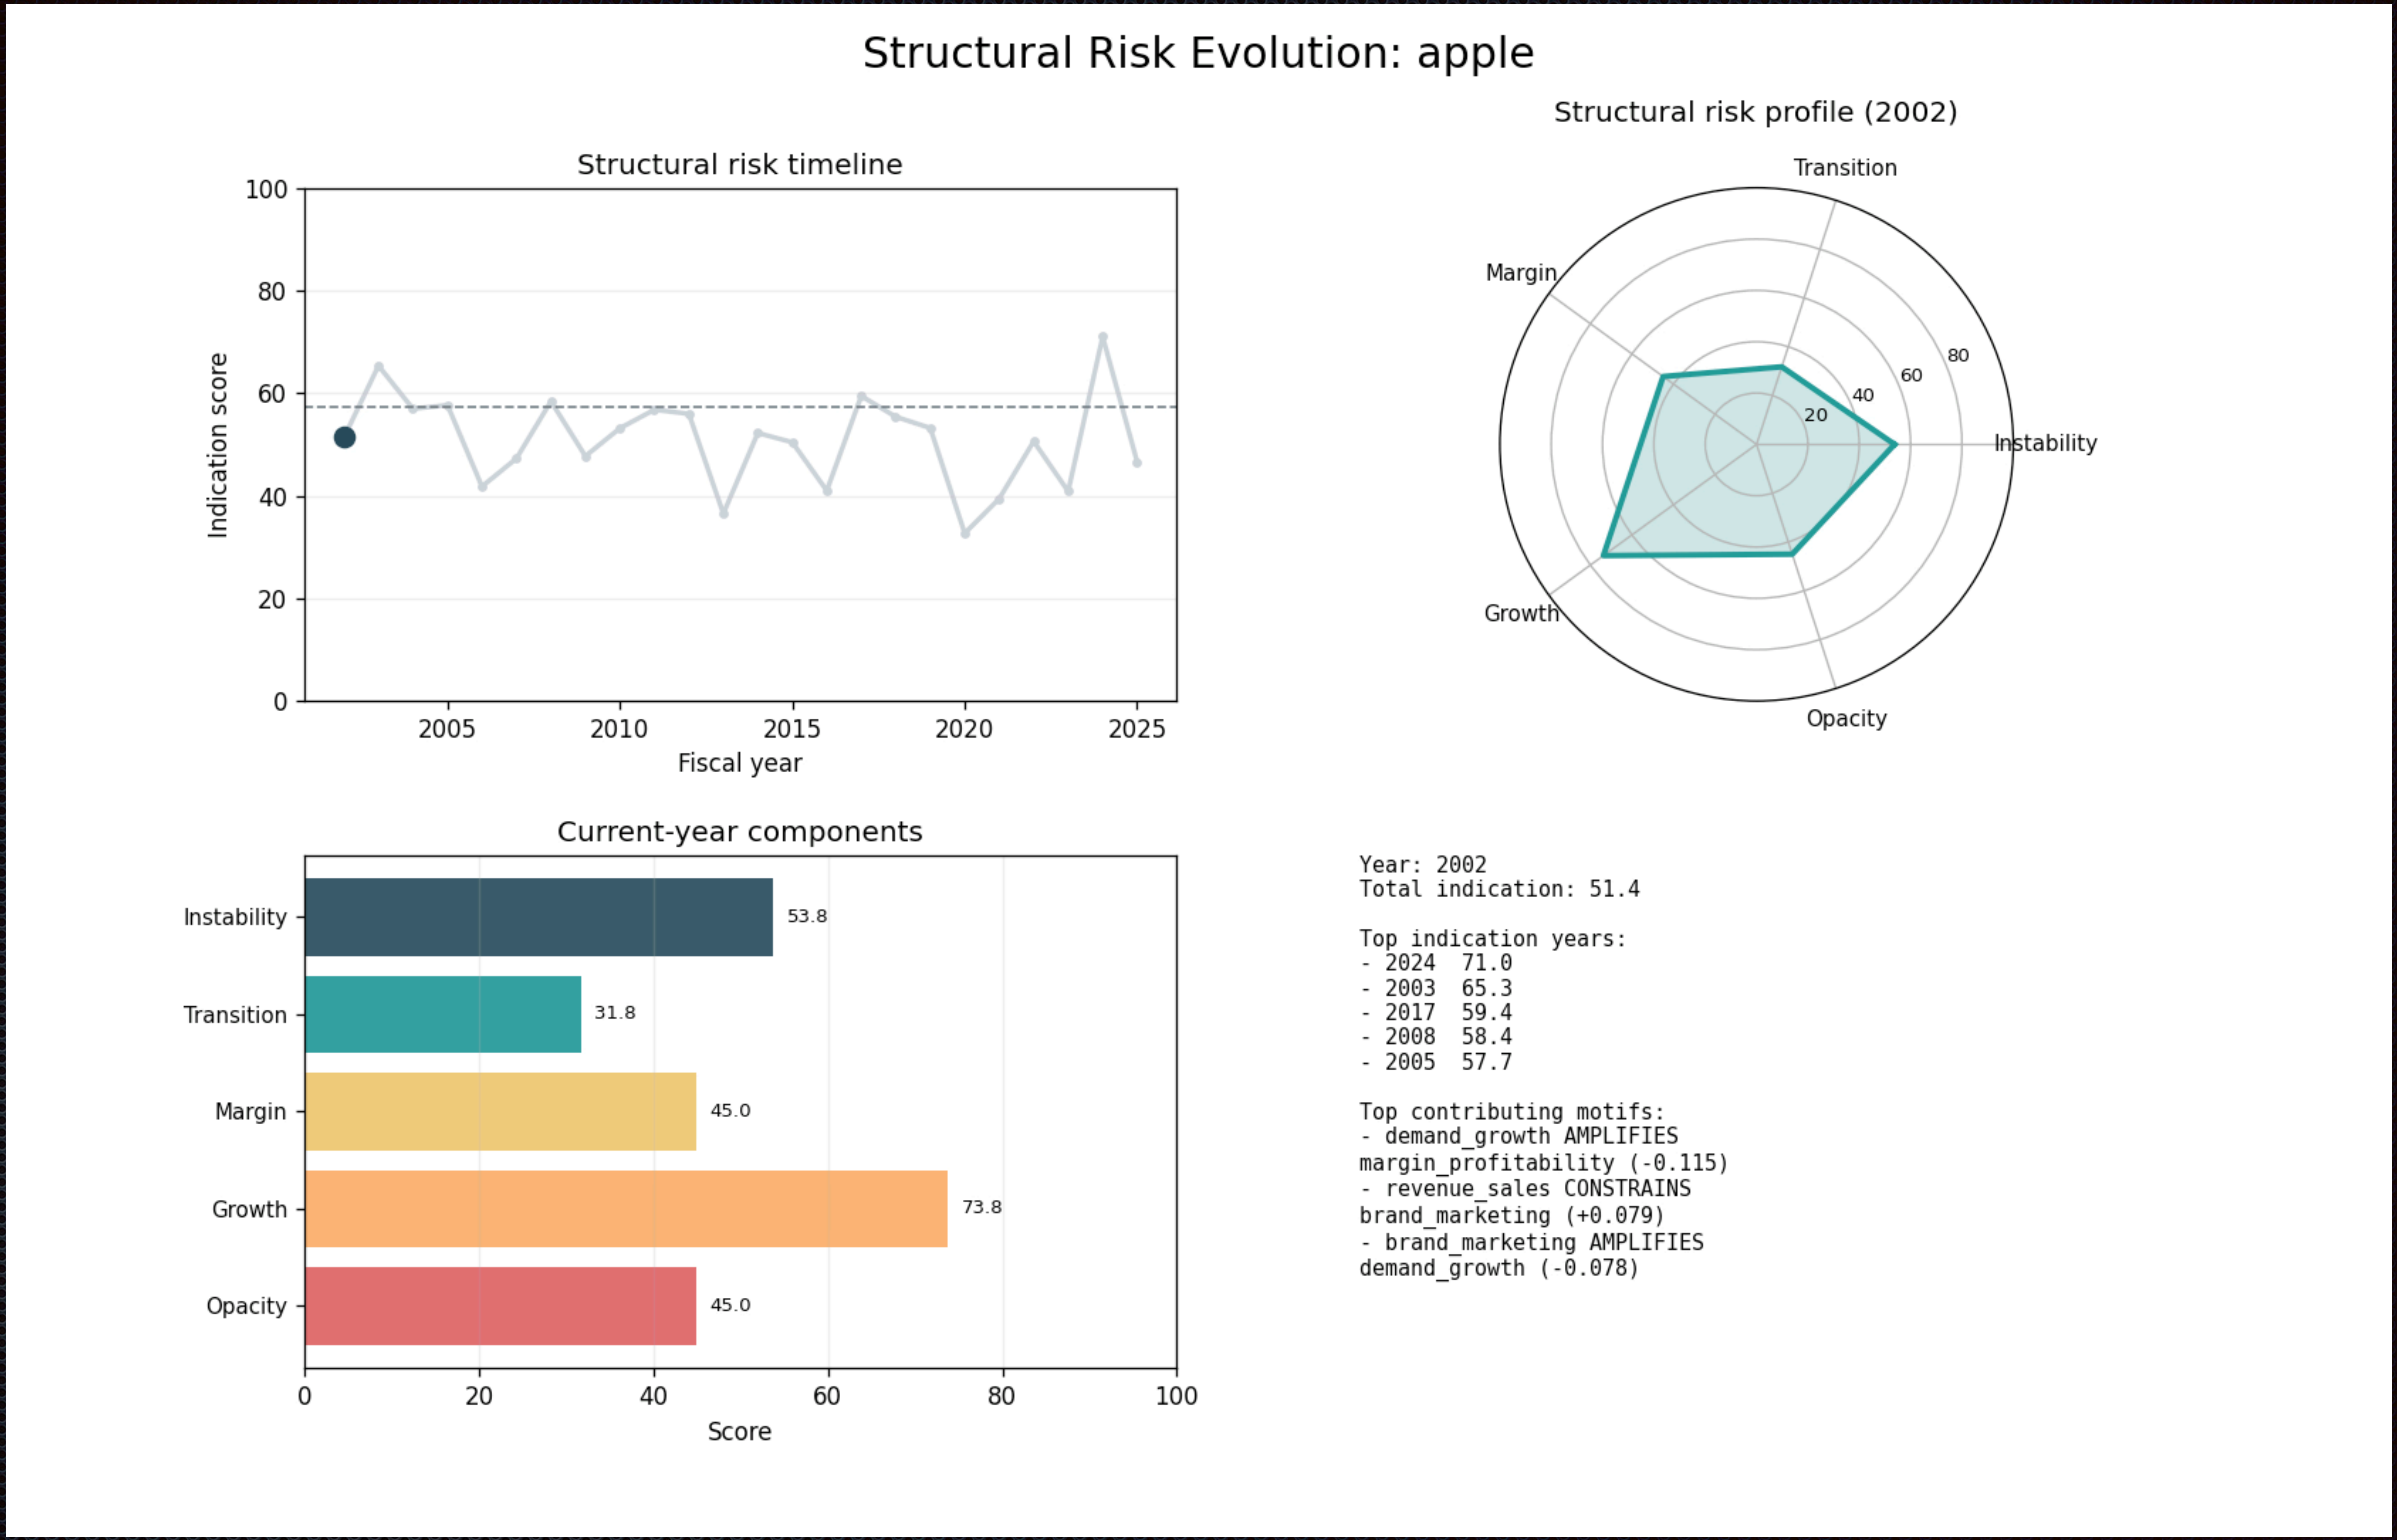

Structural Risk Evolution: apple
Structural risk timeline
Indication score
Fiscal year
Structural risk profile (2002)
Transition
Margin
Instability
Growth
Opacity
Current-year components
Instability
Transition
Margin
Growth
Opacity
53.8
31.8
45.0
73.8
45.0
Score
Year: 2002
Total indication: 51.4
Top indication years:
- 2024 71.0
- 2003 65.3
- 2017 59.4
- 2008 58.4
- 2005 57.7
Top contributing motifs:
- demand_growth AMPLIFIES margin_profitability (-0.115)
- revenue_sales CONSTRAINS brand_marketing (+0.079)
- brand_marketing AMPLIFIES demand_growth (-0.078)

Adobe Movie

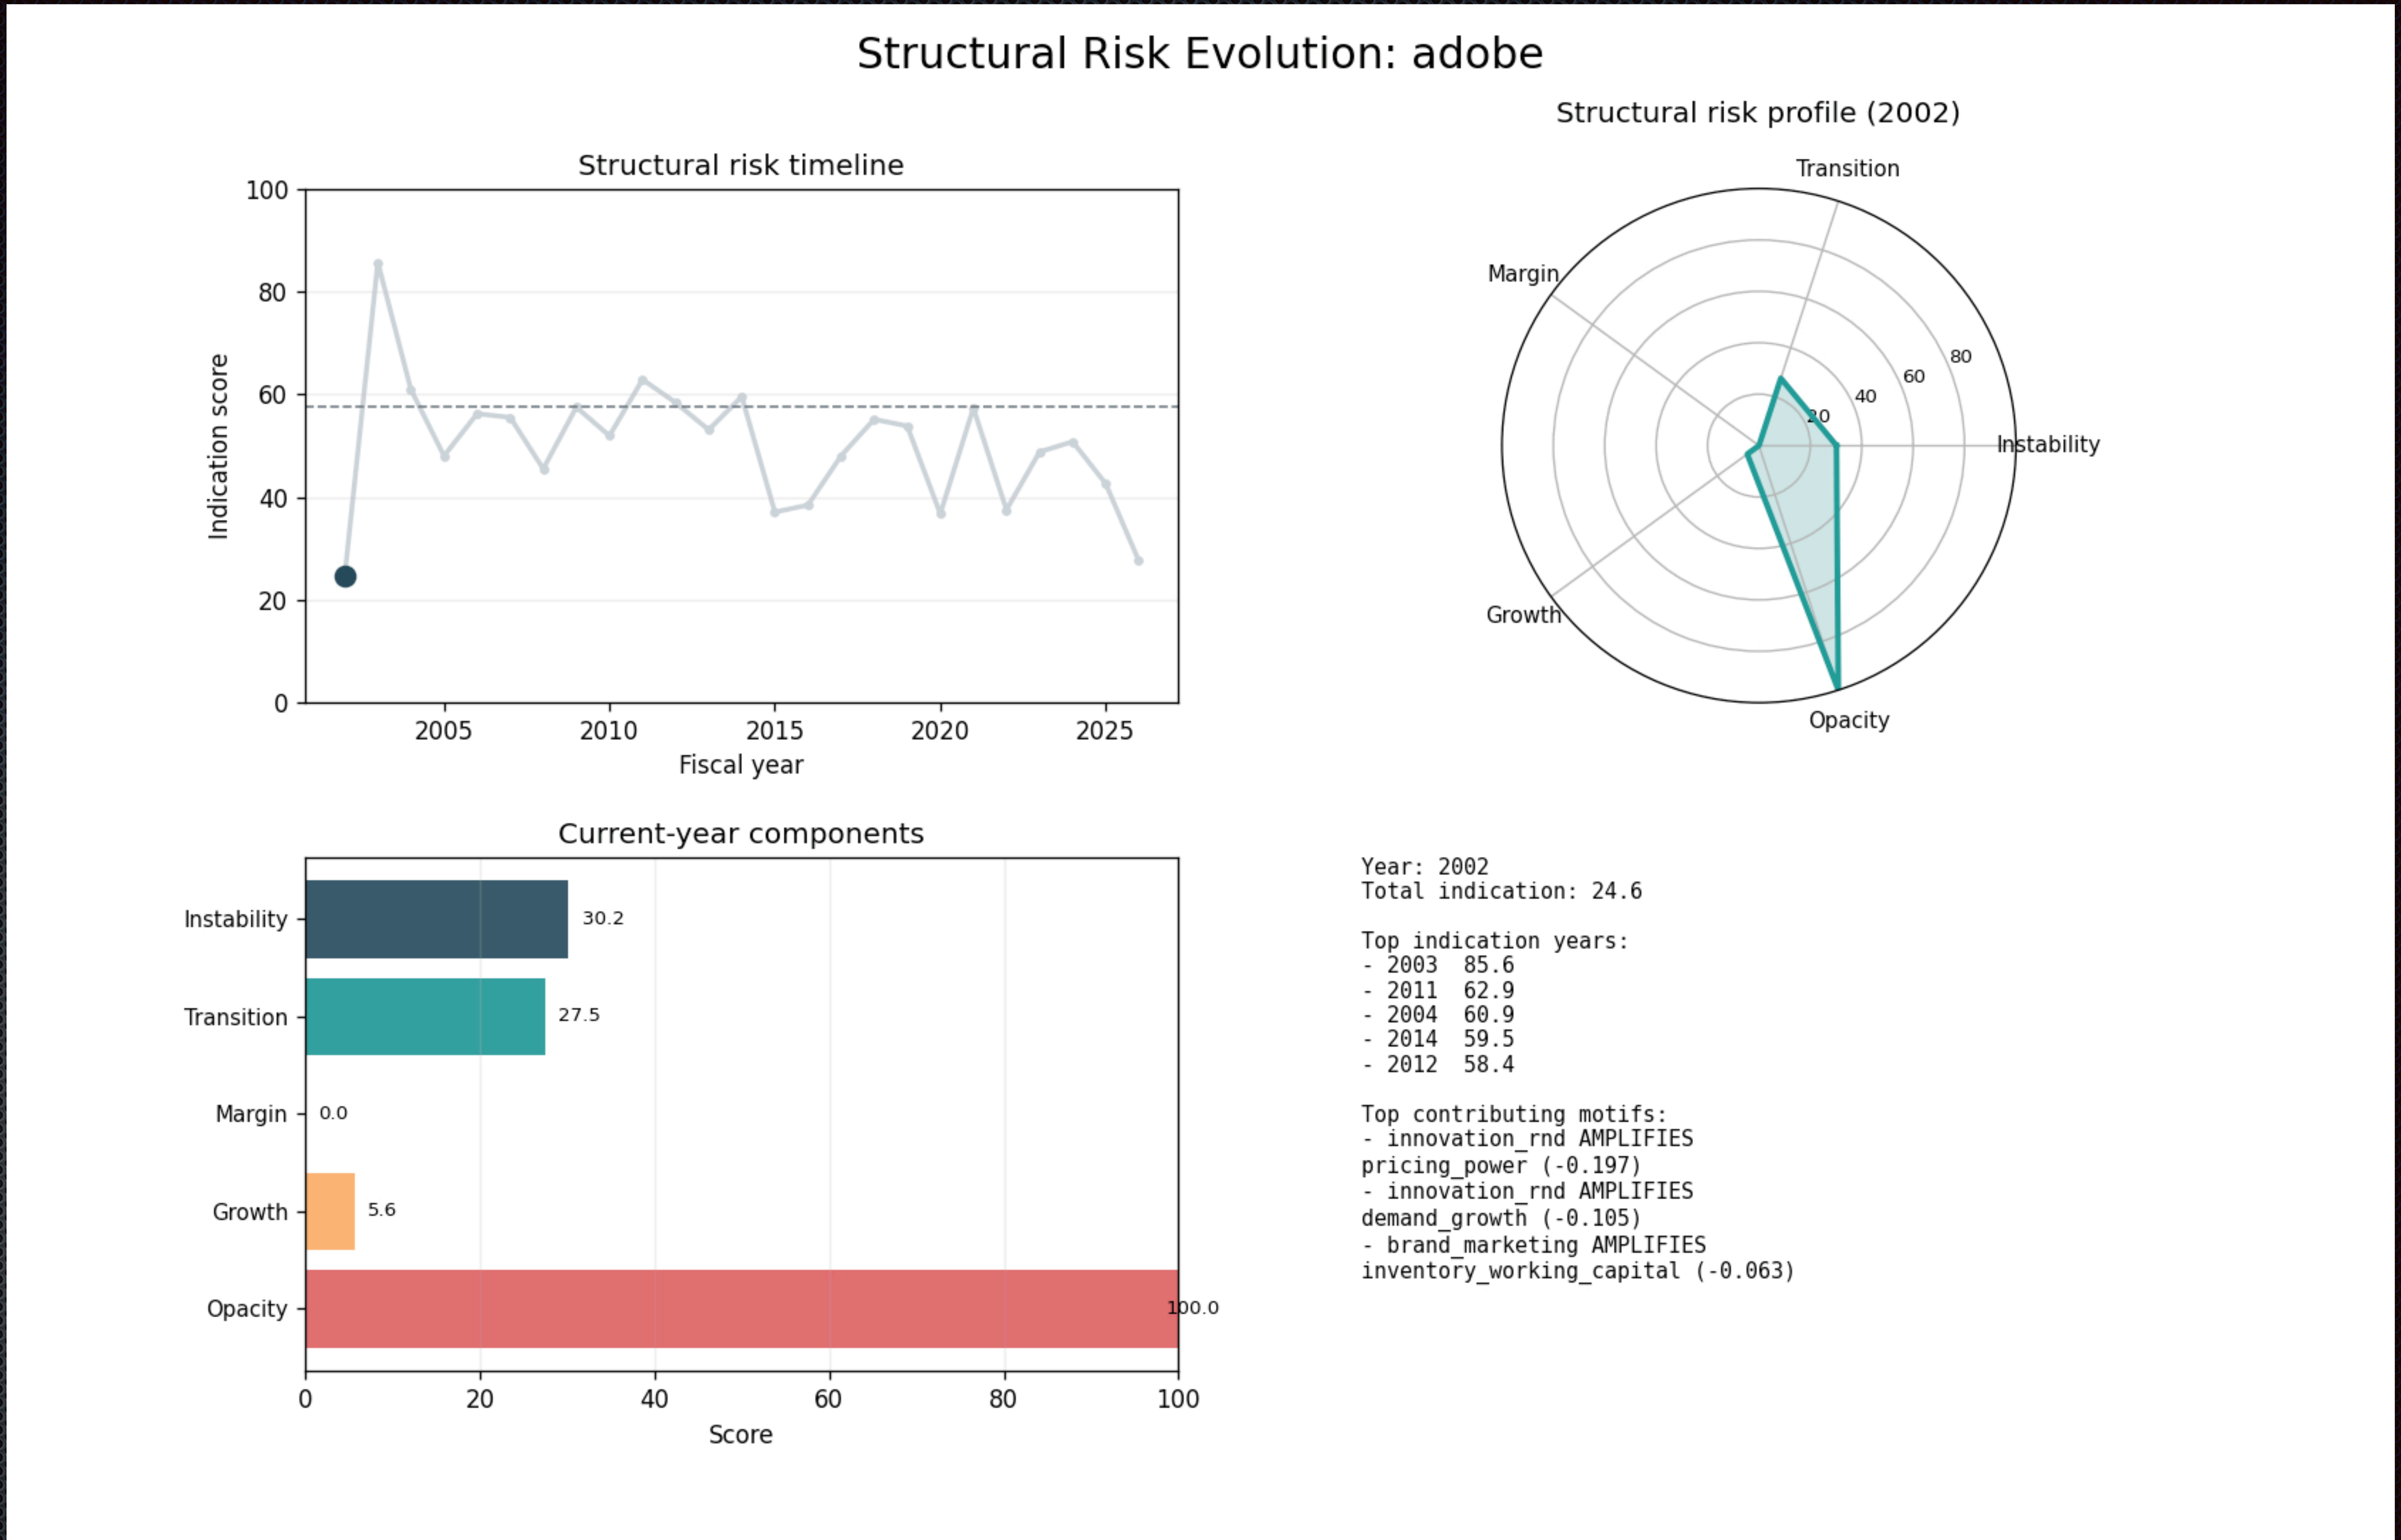

Structural Risk Evolution: adobe
Structural risk timeline
Indication score
Fiscal year
0
20
40
60
80
100
2005
2010
2015
2020
2025
Structural risk profile (2002)
Transition
Margin
Instability
Growth
Opacity
20
40
60
80
Current-year components
Instability
Transition
Margin
Growth
Opacity
30.2
27.5
0.0
5.6
100.0
Score
Year: 2002
Total indication: 24.6
Top indication years:
- 2003 85.6
- 2011 62.9
- 2004 60.9
- 2014 59.5
- 2012 58.4
Top contributing motifs:
- innovation_rnd AMPLIFIES pricing_power (-0.197)
- innovation_rnd AMPLIFIES demand_growth (-0.105)
- brand_marketing AMPLIFIES inventory_working_capital (-0.063)

Company-Level Takeaways

Company	Mean risk	Peak year	Read
Nike	51.76	2003 (85.24)	Early-years stress dominates; latest year is low overall with transition-break signal stronger than broad margin stress.
Adobe	50.13	2003 (85.64)	Similar early spike; latest profile is low-risk and looks more like controlled reconfiguration than distress.
IBM	51.72	2026 (78.33)	Latest year is broad multi-axis stress, with instability, transition break, and growth fragility all elevated.

Pairwise Relationship Matrix

Pair	Radar corr	Mean abs gap	Shared high-risk years	Read
Nike–Adobe	0.412	8.84	2003, 2014	Closest pair on radar layer.
Nike–IBM	0.033	10.25	2015	Very weak synchrony.
Adobe–IBM	0.227	10.15	none	Mild overlap, still clearly separated.

- Tightest three-way convergence years: 2021, 2012.
- Highest collective stress year: 2003, but with large cross-company dispersion.

Company Functor Similarity

Pair	Cosine	JS	Naturality defect	Read
Nike–Adobe	0.9791	0.0085	0.5000	Strong shared state geometry, but materially different dynamics.
Nike–IBM	0.0000	1.0000	0.0034	No shared-state overlap in current basis; low defect is degenerate.
Adobe–IBM	0.0000	1.0000	0.0064	Same degeneracy warning.

- Key distinction: state similarity and dynamical similarity are not the same object.
- Nike and Adobe are similar in state space, but their transitions do not commute well.

Corporate Geometry Classes

- Smooth manifolds
  - Nike
- Piecewise-smooth
  - Adobe
- High-curvature
  - IBM

Summary

- We introduced causal diffusion models to analyze corporate geometry
  - Brand-Aware Democritus uses block-denoising categorical attention models operating on higher-order simplicial complexes of causal geometry
  - It constructs a deep latent-space model of each company as a dynamical system, based on an analysis of decades of 10K filings
- We used a locally hosted Qwen 235B model running on 2 Mac Studios with 512GB of RAM
  - The yearly snapshots and summaries are available for further analysis
- For technical details on CRICKET: Mahadevan, Categories for AGI, 2026
  - https://people.cs.umass.edu/~mahadeva/papers/catagi.pdf