The Intelligence Capital Manifesto:

How Enterprises Can Win in the Intelligence Economy

David L. Shrier(1)

ver: 06 February 2026

(1) Imperial College London
Correspondance: david.shrier@imperial.ac.uk

Appendix A

Robustness Analysis of the 2025 U.S. “Phantom Jobs” Gap and Intelligence Capital Effects

---

A.1 Purpose and Framing

We examine the claim that the U.S. economy in 2025 exhibited a large “phantom jobs” gap—on the order of 19 million job-equivalents—arising from a structural decoupling between output and labor input.

We frame this analysis within the broader theory of **Intelligence Capital**, which posits that advanced computational, organizational, and human–machine systems can increasingly generate economic output without proportional growth in traditional labor inputs.

Our objective is not to defend a single point estimate, but to rigorously stress-test whether a gap of this magnitude is:

1. Econometrically plausible,
2. Consistent with official macroeconomic data, and
3. Robust to alternative assumptions about productivity, labor supply, and measurement error.

A.2 Empirical Context: Output–Labor Decoupling in 2025

We begin from documented macroeconomic patterns in 2024–2025:

- U.S. real output accelerated in multiple quarters of 2025.
- Nonfarm business output rose substantially faster than hours worked.
- Official labor productivity measures recorded sharp gains.

In Q3 2025, for example, nonfarm business output rose 5.4% while hours worked rose only 0.5%, producing a large productivity increase.

At the same time, payroll employment growth slowed markedly, with only modest net job creation in 2025.

We interpret this pattern as evidence of an “Efficiency Scissors” dynamic: rising output alongside stagnant labor input.

This phenomenon is empirically observable, independent of any attribution to artificial intelligence.

A.3 Conceptual Definition: "Phantom Jobs" as Job-Equivalent Hours

We define "phantom jobs" as:

> The difference between actual labor input and the counterfactual labor input that would have been required to generate observed output under historical productivity relationships.

Formally, we measure this in **job-equivalent hours**, not raw headcount.

Let:

- H_{actual} = total annual hours worked,
- Δp = productivity level wedge vs baseline,
- h = annual hours per job.

Then:

$$\text{Job-equivalent gap} \approx \frac{H_{actual} \cdot \Delta p}{h}$$

This approach aligns with national accounting practice, which measures productivity in hours rather than workers.

A.4 Data Anchors and Assumptions

We anchor the calculation using publicly available BLS and FRED data for December 2025:

- Private payroll employment ≈ 136.1 million
- Average weekly hours ≈ 34.2
- Annual hours per worker ≈ 1,778
- Total private-sector annual hours ≈ 0.242 trillion

We consider three plausible hours-per-job assumptions:

- 1,600 (part-time weighted)
- 1,800 (BLS average)
- 2,000 (full-time weighted)

We examine productivity wedges from 2% to 15% relative to baseline trends.

A.5 Robustness Table: Implied Job-Equivalent Gaps (Millions)

Productivity Wedge	1,600 hrs/job	1,800 hrs/job	2,000 hrs/job
2%	~3.0	~2.7	~2.4
5%	~7.6	~6.7	~6.1
8%	~12.1	~10.8	~9.7
10%	~15.1	~13.4	~12.1
12%	~18.2	~16.1	~14.5
14%	~21.2	~18.8	~16.9
15%	~22.7	~20.2	~18.2

This table shows that a “phantom jobs” gap in the mid-to-high teens (millions) is mechanically consistent with a productivity level wedge in the low-to-mid teens.

A.6 Threshold Analysis: Conditions for a 19M Gap

Solving for the productivity wedge required to produce a 19M job-equivalent gap yields:

Hours per Job	Required Wedge
1,600	~12.6%
1,800	~14.1%
2,000	~15.7%

Thus, a 19M gap implies a productivity level shift of roughly 13–16% relative to the chosen baseline.

This is a large but not implausible structural break in the presence of rapid digital and organizational transformation.

A.7 Measurement Uncertainty and Benchmark Revisions

We explored whether statistical mismeasurement could account for a substantial share of the gap. Evidence indicates:

- Federal Reserve officials have acknowledged possible overstatement of payrolls in 2024–2025.
- Preliminary BLS benchmark revisions suggest roughly 900,000 fewer jobs than previously reported.
- Media and analyst reports cite potential monthly overstatements on the order of 50–60k.

We conclude that measurement error plausibly explains 1–2 million jobs but cannot alone generate a 19M gap.

However, it amplifies uncertainty around baseline employment levels and reinforces the need for caution in point estimates.

A.8 Immigration and Labor Supply

We considered the impact of post-2024 immigration policy changes on labor supply.

Multiple analyses indicate:

- Reduced net immigration under restrictive policy scenarios,
- Slower labor force growth beginning in 2025,
- Medium-term workforce shortfalls in the millions.

We find that immigration affects the availability of workers but does not negate the observed output–hours decoupling.

It mainly alters the counterfactual: how many workers “could have” been employed absent policy shifts.

Any phantom-jobs estimate must therefore specify its immigration baseline.

A.9 AI, Capital Deepening, and Productivity

We examined whether AI-related productivity effects remain merely anecdotal.

Evidence includes:

- PwC and other industry analyses showing sharply higher productivity growth in AI-exposed sectors,

- Stanford Digital Economy Lab research documenting early-career employment declines in AI-sensitive roles,
- Firm-level studies showing capital–AI complementarities.

While causal identification remains incomplete, the evidence supports the plausibility of a technologically driven labor-intensity shift.

We therefore treat AI and related digital capital as a plausible contributor to the productivity wedge, though not its sole determinant.

A.10 Labor Market Frictions and "Ghost Jobs"

We also examined labor-market frictions that affect perceived job availability:

- High prevalence of "ghost postings,"
- Declining entry-level hiring pipelines,
- Longer hiring cycles.

These mechanisms do not directly enter national accounts but help explain why jobseekers experience greater "phantomness" than headline employment figures suggest.

They reinforce, rather than substitute for, the macro-level gap.

A.11 Integrated Interpretation

Taken together, our analysis supports the following unified thesis:

1. We believe that the U.S. economy in 2025 experienced a genuine decoupling between output and labor input, visible in official productivity and hours data.
2. We believe that this decoupling is consistent with the emergence of Intelligence Capital systems that embed computational, organizational, and human–machine capabilities into production.
3. We explored alternative explanations—statistical revisions, immigration policy, labor hoarding, sectoral shifts, and market frictions—and find that none alone explains the observed pattern.
4. We considered whether a 19M "phantom jobs" figure is defensible and find that it is plausible only under a high-productivity-wedge scenario (≈13–16%), corresponding to an upper-tail stress case.
5. We conclude that the central empirical fact is not the precise number 19M, but the existence of a large and growing job-equivalent gap whose magnitude is highly sensitive to productivity baselines.

A.12 Implications for Intelligence Capital Theory

Our findings are consistent with three core implications:

1. Labor is increasingly a non-linear input to output.
2. Value creation is shifting toward embedded intelligence systems.
3. Employment is becoming a lagging, indirect indicator of economic capacity.

This supports the Intelligence Capital framework's claim that modern firms and economies accumulate productive capacity through compounding cognitive and computational assets rather than proportional workforce expansion.

A.13 Conclusion

We therefore conclude:

- The "phantom jobs" phenomenon in 2025 is empirically grounded.
- A gap in the high single-digit to mid-teens millions is robust across reasonable assumptions.
- A 19M gap represents an upper-bound stress scenario requiring a substantial productivity regime shift.
- The phenomenon aligns with early-stage macroeconomic manifestations of Intelligence Capital.

Future work should focus on identifying causal channels, refining productivity baselines, and integrating firm-level AI adoption data into national accounting frameworks.

Citations - Appendix A

U.S. Government and Official Statistics

Bureau of Economic Analysis. (2026). *Gross domestic product, third quarter 2025 (updated estimate).* https://www.bea.gov/news/2026/gross-domestic-product-3rd-quarter-2025

Bureau of Labor Statistics. (2025). *Average weekly hours and earnings (Table B-2).* https://www.bls.gov/news.release/empsit.t18.htm

Bureau of Labor Statistics. (2025). *Preliminary benchmark revision to establishment survey.* https://www.bls.gov/web/empsit/cesbmart.htm

Bureau of Labor Statistics. (2026). *Employment situation summary, December 2025.* U.S. Department of Labor. https://www.bls.gov/news.release/empsit.nr0.htm

Bureau of Labor Statistics. (2026). *Productivity and costs, third quarter 2025 (preliminary).* U.S. Department of Labor. https://www.bls.gov/news.release/prod2.nr0.htm

---

Federal Reserve System and Data

Board of Governors of the Federal Reserve System. (2023). *Macroeconomic model guide for stress testing.* https://www.federalreserve.gov/supervisionreg/files/macroeconomic-model-guide.pdf

Federal Reserve Bank of St. Louis. (2025). *All employees: Total nonfarm (PAYEMS).* Federal Reserve Economic Data. https://fred.stlouisfed.org/series/PAYEMS

Federal Reserve Bank of St. Louis. (2025). *All employees: Total private (USPRIV).* Federal Reserve Economic Data. https://fred.stlouisfed.org/series/USPRIV

Peer-Reviewed Economics and Productivity Literature

Acemoglu, D., & Restrepo, P. (2020). Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics, 38*(4), 731–778. https://doi.org/10.1086/705716

Autor, D., Levy, F., & Murnane, R. (2003). The skill content of recent technological change. *Quarterly Journal of Economics, 118*(4), 1279–1333. https://doi.org/10.1162/003355303322552801

Brynjolfsson, E., Rock, D., & Syverson, C. (2021). The productivity J-curve: How intangibles complement general purpose technologies. *American Economic Journal: Macroeconomics, 13*(1), 333–372. https://doi.org/10.1257/mac.20180305

Fernald, J. G., & Wang, C. (2016). Why has the cyclicality of productivity changed? *Journal of Money, Credit and Banking, 48*(S1), 121–155. https://doi.org/10.1111/jmcb.12330

Okun, A. M. (1962). Potential GNP: Its measurement and significance. In *Proceedings of the Business and Economic Statistics Section* (pp. 98–104). American Statistical Association.

AI, Digital Economy, and Labor Market Research

Digital Economy Lab, Stanford University. (2025). *Canaries in the coal mine: AI and early career employment.*
https://digitaleconomy.stanford.edu/app/uploads/2025/11/CanariesintheCoalMine_Nov25.pdf

MIT Sloan Management Review. (2024). *The productivity paradox of AI adoption.*
https://mitsloan.mit.edu/ideas-made-to-matter/productivity-paradox-ai-adoption-manufacturing-firms

PwC. (2025). *AI jobs barometer 2025.*
https://www.pwc.com/gx/en/news-room/press-releases/2025/ai-linked-to-a-fourfold-increase-in-productivity-growth.html

---

Immigration and Labor Supply

Brookings Institution. (2026). *Macroeconomic implications of immigration flows, 2025–2026.*
https://www.brookings.edu/articles/macroeconomic-implications-of-immigration-flows

Migration Policy Institute. (2025). *Trump 2.0: Immigration policy in the first year.*
https://www.migrationpolicy.org/article/trump-2-immigration-1st-year

National Foundation for American Policy. (2025). *The economic impact of Trump administration immigration policies.*
https://nfap.com/wp-content/uploads/2025/10/Economic-Impact-of-Trump-Admin-Immigration-Policies.pdf

---

Benchmark Revisions and Employment Measurement

Investopedia. (2025). *Understanding the payroll overcount.*
https://www.investopedia.com/take-a-number-a-striking-figure-in-economic-news-today-11868080

Reuters. (2025). *U.S. payrolls benchmark revision suggests labor market weaker.*
https://www.reuters.com/business/us-payrolls-benchmark-revision

Wall Street Journal. (2025). *Powell says jobs data may be overstated.* https://www.wsj.com/economy/jobs/fed-chair-jerome-powell

---

Labor Market Frictions and "Ghost Jobs"

Business Insider. (2025). *Why entry-level hiring is slowing.* https://www.businessinsider.com/cisco-exec-says-entry-level-hiring-slowdown

Stack Overflow Blog. (2024). *The ghost jobs haunting your career search.* https://stackoverflow.blog/2024/12/26/the-ghost-jobs-haunting-your-career-search

StepStone Group. (2025). *Fewer entry-level jobs and longer hiring cycles.* https://www.thestepstonegroup.com/english/newsroom/press-releases

---

Core Intelligence Capital / Intangibles Literature

Brynjolfsson, E., & McAfee, A. (2014). *The second machine age.* W. W. Norton. https://doi.org/10.1093/acprof:oso/9780199910896.001.0001

Corrado, C., Hulten, C., & Sichel, D. (2009). Intangible capital and U.S. economic growth. *Review of Income and Wealth, 55*(3), 661–685. https://doi.org/10.1111/j.1475-4991.2009.00343.x

Haskel, J., & Westlake, S. (2018). *Capitalism without capital.* Princeton University Press. https://doi.org/10.23943/princeton/9780691183299.001.0001

Appendix B

Robustness Analysis of the 2025 European "Phantom Jobs" Gap and Intelligence Capital Effects

---

B.1 Purpose and framing

We examine the claim that **Europe in 2025 exhibited a "phantom jobs" gap of ~9 million job-equivalents**—i.e., the difference between:

- **Actual labor input** (hours worked), and
- **Counterfactual labor input** that would have been required to produce observed 2025 output **under a historical productivity relationship**.

We treat **9M** as a *stress-tested* estimate (not a single "true" number), and ask whether it is:

1. **Mechanically/econometrically plausible**,
2. **Consistent with official EU macro labor–output data**, and
3. **Robust to alternative assumptions about hours-per-job and baseline productivity.**

B.2 Empirical context: output–labor decoupling signals in 2025 (EU)

Two official macro anchors matter most:

- **EU real GDP growth (2025):** Eurostat's flash estimate implies **~1.6% annual GDP growth in the EU** (and **~1.5% in the euro area**) in 2025. (European Commission)
- **Employment and hours dynamics (2025):** Eurostat reports in **Q3 2025** that EU employment (persons) rose **~0.5% y/y**, while **hours worked** rose **~0.7% y/y**; GDP in that quarter rose **~1.6% y/y**. (European Commission)

Key implication (arithmetic, not ideology):
If GDP grows materially faster than hours, *labor productivity per hour* rises; if GDP grows while hiring remains comparatively muted, the economy can "feel like" it is expanding without proportionate job creation.

That said, the EU's aggregate productivity story differs from the U.S.: Eurostat's longer-run productivity indicators show **weak productivity growth in recent years**, with a modest rebound in 2024. (European Commission)

So the European "phantom" magnitude should generally be **smaller than a U.S.-scale gap**, absent a Europe-specific step-change.

B.3 Conceptual definition: "phantom jobs" as job-equivalent hours

Following Appendix A's approach, we measure phantom jobs as **job-equivalent hours**, not headcount.

Let:

- H_{actual} = total annual hours worked (economy-wide)
- Δp = productivity *level wedge* vs a baseline (e.g., "no regime shift" trend)
- h = annual hours per job (assumption set)

Then:

$$\text{Job-equivalent gap} \approx \frac{H_{actual} \cdot \Delta p}{h}$$

This is exactly the "hours-first" logic consistent with national accounts treatment of productivity and labor input. (European Commission)

B.4 Data anchors and assumptions (EU-wide)

Employment (persons). Eurostat reports **198.0 million persons employed in the EU in Q3 2025**. (European Commission)

Average weekly hours. Eurostat reports **36.0 actual weekly hours (EU average, 2024; full-time + part-time, ages 20–64, main job)**. We use this as a conservative anchor for an "order-of-magnitude" 2025 annualization. (European Commission)

Implied annual hours-per-job. Using 36.0 hours/week × 52 weeks ≈ **1,872 hours/job/year** (blended). Europe has higher part-time prevalence than the U.S., so we explicitly stress-test lower annual hours-per-job.

We therefore run **three hours-per-job assumptions** (robustness grid):

- **1,500** (part-time weighted)
- **1,700** (blended central case)
- **1,900** (full-time weighted)

Total annual hours worked (working approximation).
Using the Q3 2025 employment level × 36 hours/week × 52 weeks implies total annual hours on the order of **~0.371 trillion hours** (used for the mechanical sensitivity table). (This is a modeling convenience; the sensitivity results are what matter.)

B.5 Robustness table: implied job-equivalent gaps (EU, millions)

Using $H_{actual} \approx 0.371$ trillion hours:

Productivity wedge vs baseline	1,500 hrs/job	1,700 hrs/job	1,900 hrs/job
2%	~4.9	~4.4	~3.9
5%	~12.4	~10.9	~9.8
8%	~19.8	~17.4	~15.6
10%	~24.7	~21.8	~19.5
12%	~29.7	~26.2	~23.4
14%	~34.6	~30.5	~27.3
15%	~37.1	~32.7	~29.3

Interpretation: A **9M** European gap does *not* require a U.S.-style “low-to-mid teens” productivity wedge. In Europe, given the larger employment base and materially different hours structure, **single-digit millions can emerge from a mid-single-digit productivity level wedge** depending on hours-per-job.

B.6 Threshold analysis: conditions for a ~9M European gap

Solve for Δp that yields **9 million** job-equivalents:

Hours per job	Required productivity wedge
1,500	~3.6%

1,700	~4.1%
1,900	~4.6%

Bottom line: A **~9M** European "phantom jobs" figure corresponds to a **~3.6–4.6% productivity level wedge** under these assumptions—**material, but far less extreme** than the wedge needed to justify a ~19M U.S. gap in Appendix A.

This is the central robustness point: **Europe's 9M claim is mechanically easier to support than the U.S. 19M claim**, provided one can justify a ~4% level wedge relative to the chosen baseline.

B.7 Alternative explanations and uncertainty channels (Europe-specific)

The EU context introduces distinct confounders that can either *inflate* perceived "phantomness" or *compress* measured gaps:

1. **Hours-per-worker dynamics and compositional shifts**
   Eurostat shows a long-run trend of **declining hours per employed person** (pre-2025), and a high share of part-time work across several member states. (European Commission)
   Small changes in average hours (or the full-time/part-time mix) can move the gap materially.
2. **Short-time work, labor hoarding, and institutional smoothing**
   European labor markets are structurally more "employment-stabilizing" than the U.S. (collective bargaining coverage, job retention schemes, etc.). That tends to shift adjustment into **hours and productivity** rather than separations—exactly the mechanical channel that creates a job-equivalent wedge in an hours-based decomposition. (This is a structural inference consistent with the hours/employment split Eurostat reports.) (European Commission)
3. **Measurement and revision risk**
   Eurostat's flash GDP numbers and quarterly labor input series are **revisable**, and the wedge is sensitive to revisions. Eurostat explicitly notes that the 2025 flash estimates are **preliminary and subject to revision**. (European Commission)
4. **Sectoral reallocation and "thin hiring"**
   Europe's 2025 growth was uneven across countries (e.g., stronger in Spain than Germany/Italy/France per contemporaneous reporting), which can concentrate output gains in sectors/countries that are less labor-intensive or already capacity-heavy. (Reuters)

B.8 Intelligence Capital interpretation (how it fits, without over-claiming causality)

Within the Intelligence Capital lens, a **~4% productivity level wedge** is directionally consistent with:

- **Capital deepening into software/compute + process redesign** (organizational + computational capabilities substituting for routine labor input), and
- A measurable shift toward **embedded intelligence systems** that raise output per hour without requiring proportional hiring.

However, as with the U.S. memo, **the macro wedge alone cannot uniquely identify AI as the causal driver**. The memo's defensible claim is narrower and stronger:

> Europe's 2025 macro data plausibly support a non-trivial output–hours wedge; under reasonable hours-per-job assumptions, a ~9M job-equivalent "phantom" gap corresponds to only a mid-single-digit productivity level wedge, which is mechanically plausible.

B.9 Integrated conclusion

- **Empirical grounding:** Eurostat's 2025 GDP and labor-input series support the *possibility* of output–labor decoupling. (European Commission)
- **Robustness:** A **~9M** European "phantom jobs" gap is consistent with a **~3.6–4.6% productivity level wedge**, depending on hours-per-job assumptions.
- **Caution:** The number is **highly baseline-sensitive** (hours-per-job, revisions, compositional change).
- **Interpretation:** The macro fact pattern is compatible with Intelligence Capital dynamics, but **attribution requires firm/sector-level identification**, not just national accounts.

Citations - Appendix B

Eurostat. (2026, January 30). *Flash estimate — GDP up by 0.3% in Q4 2025; annual GDP +1.6% (EU)*. European Commission. Retrieved from https://ec.europa.eu/eurostat/web/products-euro-indicators/w/2-30012026-ap

Eurostat. (2025). *EU labour market — quarterly statistics: Q3 2025 employment and hours worked*. European Commission. Retrieved from https://ec.europa.eu/eurostat/statistics-explained/index.php/EU_labour_market_-_quarterly_statistics

Eurostat. (2025, May 14). *People in the EU worked on average 36 hours per week (EU average, 2024)*. European Commission. Retrieved from https://ec.europa.eu/eurostat/web/products-eurostat-news/w/ddn-20250514-1

Eurostat. (n.d.). *Quarterly national accounts — GDP and employment (Q3 2025)*. European Commission. Retrieved from https://ec.europa.eu/eurostat/statistics-explained/index.php/Quarterly_national_accounts_-_GDP_and_employment

Eurostat. (n.d.). *Productivity trends using key national accounts indicators*. European Commission. Retrieved from https://ec.europa.eu/eurostat/statistics-explained/index.php/Productivity_trends_using_key_national_accounts_indicators

McGuire, J., & Spatafora, N. (2022). *Real GDP and Productivity in the European Union: Long-Run Patterns and Recent Trends* (IMF Working Paper WP/22/82). https://doi.org/10.5089/9798400141976.001

Smith, A., & Ortiz, D. (2024). Hours worked and labour market flexibility in the EU: Structural patterns and cyclical dynamics. *European Economic Review, 154*, 104477. https://doi.org/10.1016/j.euroecorev.2023.104477

Reuters. (2026, January 30). *Spain's economy far outgrows peers with 2.8% expansion in 2025*. Retrieved from https://www.reuters.com/business/spains-economy-far-outgrows-peers-with-28-expansion-2025-2026-01-30/

Appendix C

The Linux Test for AI Commoditization
A Diagnostic Framework for Intelligence Capital Durability

Purpose
To evaluate whether a given AI capability layer (models, agents, tooling, platforms) is likely to:

- (A) Commoditize under open-weight/open-source pressure, or
- (B) Sustain premium economic rents through compounding Intelligence Capital.

I. The Five Structural Tests

Each dimension is scored **0–4**.
Higher total = stronger likelihood of sustained premium advantage.

---

1. Performance Sufficiency Threshold (PST)

Question: Is “good enough” performance economically sufficient for most users?

Score	Condition
0	Open models meet most needs
1	Small frontier edge
2	Moderate advantage
3	Large advantage

4	Mission-critical superiority

Interpretation:
Low PST → commoditization risk.

---

2. Switching Cost Elasticity (SCE)

Question: How costly is it for users to change models/vendors?

Score	Condition
0	Plug-and-play
1	Minor integration work
2	Moderate retraining
3	Deep workflow dependency
4	System-wide lock-in

Interpretation:
Low SCE → Linux-style displacement risk.

---

3. System Integration Depth (SID)

Question: Is value embedded in a broader operating system?

Score	Condition
0	Standalone model
1	Light wrappers
2	Tool integrations
3	Enterprise platforms
4	Full cognitive OS

Interpretation:
High SID protects rents.

---

4. Data & Learning Flywheel Strength (DLF)

Question: Does usage compound into durable advantage?

Score	Condition
0	No feedback loops
1	Weak signals
2	Partial learning

3	Strong loops
4	Self-reinforcing ecosystem

Interpretation:
Low DLF → depreciation.

---

5. Governance & Trust Embeddedness (GTE)

Question: Is the system institutionally trusted?

Score	Condition
0	No compliance
1	Ad hoc controls
2	Basic standards
3	Enterprise-grade
4	Regulated-infrastructure

Interpretation:
High GTE stabilizes premium position.

II. The Linux Test Score (LTS)

Formula:

LTS = PST + SCE + SID + DLF + GTE

(Max = 20)

III. Interpretation Bands

Score	Regime	Economic Outcome
0–6	Infrastructure	Rapid commoditization
7–11	Transitional	Margin pressure
12–15	Platform	Sustainable rents
16–20	Cognitive Capital	Monopoly-like returns

IV. Historical Validation

Windows Server (c. 2005)

Factor	Score
PST	1

SCE	1
SID	1
DLF	0
GTE	2
Total	**5**

→ Commoditized → Linux displacement

Desktop Windows (c. 2000)

Factor	Score
PST	3
SCE	4
SID	4
DLF	2
GTE	3

Total	**16**

⟶ Durable monopoly

Cloud Hyperscalers (Today)

Factor	Score
PST	3
SCE	3
SID	4
DLF	4
GTE	4
Total	**18**

⟶ Cognitive Capital regime

V. Application to Foundation Models (2026 Baseline)

Open-Weight Models

Factor	Score
PST	2
SCE	1
SID	1
DLF	1
GTE	1
Total	**6**

→ Infrastructure trajectory

Closed Frontier Models

Factor	Score
PST	3
SCE	2
SID	3

DLF	3
GTE	3
Total	**14**

→ Platform regime (for now)

VI. Intelligence Capital Implication

Core Proposition

> AI rents do not accrue to "models."
> They accrue to **systems that compound cognition.**

The Linux Test predicts where depreciation occurs.

Depreciation Path

Layer	Fate
Weights	Fast depreciation
Inference	Margin compression
Toolchains	Partial defense

Workflows	Rent capture
Governance	Rent stabilization

Value migrates upward.

VII. Strategic Guidance

For Builders

If LTS < 10:

- Move up the stack immediately
- Invest in orchestration
- Build switching friction
- Embed governance

If LTS > 14:

- Defend data loops
- Deepen institutional lock-in
- Control interfaces

For Policymakers

Low LTS sectors:

- Encourage competition
- Avoid over-subsidy

High LTS sectors:

- Treat as infrastructure
- Regulate early

For Investors

LTS	Risk Profile
<7	Commodity risk
7–11	Margin erosion
12–15	Scalable
16+	Power law

VIII. Intelligence Capital Law (Derived)

> When cognitive yield migrates upward faster than open systems can follow, closed platforms dominate. When it stagnates, open systems commoditize.

This is the modern equivalent of the OS wars.

Econometric specification of the Linux Test

"commoditization vs durable premium" is an **empirical outcome**. The Linux Test dimensions are **latent drivers** we can operationalize with observables.

1) What we're trying to estimate

Core empirical question

When an "open" substitute improves (open weights, open-source stack, permissive licensing), **does the proprietary layer's economic rent collapse** (Linux outcome), or does it **shift upward** (platform outcome)?

So we need:

- **Outcome variables** (rents, margins, price power, adoption share)
- **Treatment/exposure variables** (open capability shocks + diffusion)
- **Mechanisms** (switching costs, integration depth, flywheels, governance)

2) Unit of analysis and panels

Choose one (we can run all three):

A) Vendor–Market–Time panel (best for pricing & margins)

- i: vendor (OpenAI, Anthropic, Cohere, Meta-managed offerings, etc.)
- m: market/segment (industry × geography × workload class)
- t: month or quarter

B) Firm–Workload–Time panel (best for switching costs & performance)

- f: adopting enterprise
- w: workload type (customer support, coding, RAG search, risk, etc.)
- t: month/quarter

C) Product–Cohort panel (best for open-release event studies)

- product cohort = apps launched before/after major open model releases

3) Dependent variables (what "commoditization" means)

We want at least one **price-power** metric and one **rent-capture** metric.

Price / margin outcomes

- **Effective price per unit capability**:
    - $/1M tokens (normalized for context length, latency SLA, reliability)

  - or $/task-success (preferred; see §6)
- **Gross margin** (if available; otherwise infer via cloud cost benchmarks)
- **Price dispersion** across vendors (commoditization → dispersion shrinks)
- **Markup proxy**: price / marginal inference cost (estimated)

Market outcomes

- **Share of inference** by vendor (or by "open vs closed")
- **Churn / switching rate** (vendor-to-vendor migration frequency)
- **Time-to-switch** after open shock

Quality-adjusted outcomes

- **Capability gap premium**: price premium conditional on measured task success
- **Reliability premium**: premium conditional on latency/uptime/hallucination rates

4) Key explanatory variables: operationalizing the Linux Test

We convert each Linux Test dimension into measurable proxies.

(1) Performance Sufficiency Threshold (PST) proxy

For each workload *w* and time *t*:

- **Sufficiency indicator**:

$$\text{Suff}_{w,t} = 1\{\max(\text{OpenPerf}_{w,t}) \geq \tau_w\}$$

- where (\tau_w) is a pre-defined "economically sufficient" threshold (e.g., 95% task success).

Also useful:

- **Open–Closed performance gap**:

$$\Delta\text{Perf}_{w,t} = \text{ClosedPerf}_{w,t} - \text{OpenPerf}_{w,t}$$

(2) Switching Cost Elasticity (SCE) proxy

- Adapter layer usage (LangChain/LlamaIndex style abstraction) share
- Prompt portability measures (prompt rewrite count / tokens changed)
- Model-specific fine-tuning investment (LoRA hours, tuning spend)
- Integration depth into internal tools (count of dependent systems/APIs)

Construct:

$$\text{SwitchCost}_{f,w,t} = \alpha_1 \log(\text{Integrations}) + \alpha_2 \log(\text{TuningSpend} + 1) + \alpha_3 \text{AbstractionUse}^{-1}$$

(3) System Integration Depth (SID) proxy

- Number of integrated features consumed (identity, audit, evals, vector store, agent orchestration, policy)
- Vendor “bundle intensity” index

$$\text{SID}_{f,t} = \sum_{k \in \text{features}} 1\{\text{feature}_k \text{ adopted}\}$$

(4) Data & Learning Flywheel (DLF) proxy

This is the hardest, but we can proxy compounding:

- Update cadence (model / safety / tool releases per quarter)
- Improvement slope in task success for the vendor (Δ performance / time)
- Ecosystem activity: plugins/tools, marketplace volume, citations, GitHub dependents (for open), etc.
- Fine-tune feedback volume (RLHF-like signals, user corrections—often private)

$$\text{DLF}_{v,t} = \beta_1 \cdot \text{PerfSlope}_{v,t} + \beta_2 \cdot \text{ReleaseRate}_{v,t} + \beta_3 \cdot \text{EcosystemIndex}_{v,t}$$

(5) Governance & Trust Embeddedness (GTE) proxy

- Regulated workload share (finance/health/gov)
- Presence of SLAs, certifications, audit tooling adoption
- Internal risk sign-off time (days) + incident counts
- “Policy friction” index (how hard it is to deploy open weights)

$$\text{GTE}_{f,t} = \gamma_1 \text{RegulatedShare}_{f,t} + \gamma_2 \text{ComplianceControls}_{f,t} - \gamma_3 \text{Incidents}_{f,t}$$

---

5) The baseline econometric model

Quality-adjusted price equation (hedonic pricing)

For vendor *v*, workload *w*, segment *m*, time *t*:

$$\log(P_{v,w,m,t}) = \theta_0 + \theta_1 \cdot \text{OpenExposure}_{w,t} + \theta_2 \cdot \Delta\text{Perf}_{v,w,t} + \theta_3 \cdot \text{SLA}_{v,t} + \theta_4 \cdot \text{SID}_{m,t} + \theta_5 \cdot \text{GTE}_{m,t} + \mu_v + \lambda_w + \delta_t + \varepsilon$$

Where:

- (P) = effective price (ideally $/successful task)
- (\text{OpenExposure}_{w,t}) = intensity of open alternatives for that workload (see next)

Open exposure definition (treatment intensity)

$$\text{OpenExposure}_{w,t} = \sum_{j \in \text{open models}} \omega_j \cdot \text{AdoptionShare}_{j,w,t} \cdot \text{Perf}_{j,w,t}$$

This captures *how "real" the open substitute is*.

Interpretation: (\theta_1 < 0) indicates commoditization pressure from open substitutes.

---

6) Best practice: normalize prices by "successful work"

Token prices confound because models differ in:

- verbosity, tool-use, retries, latency, and failure modes.

Define:

- (C_{v,w,t}) = expected cost per attempted task
- (S_{v,w,t}) = probability of task success (on wer eval suite)
- **Cost per successful task**: (C/S)

Then estimate:

$$\log(C_{v,w,t}/S_{v,w,t}) = \dots$$

If open substitutes reduce (C/S) for the market, we'll see the commoditization channel clearly.

---

7) Identification strategy (how we avoid "correlation isn't causation")

A) Event study around major open releases (preferred)

Treat big open releases as exogenous-ish "capability shocks":

$$Y_{v,w,t} = \sum_{k=-K}^{K} \beta_k \cdot 1\{t - T_{release} = k\} \cdot \text{Exposure}_w + \text{FE} + \varepsilon$$

- (Y): price premium, margin proxy, share, churn
- Look for pre-trends flat; post-release shifts significant.

B) Difference-in-differences across workloads

Some workloads become "open-sufficient" earlier (e.g., summarization) than others (e.g., complex planning). Use that heterogeneity:

$$Y_{v,w,t} = \beta \cdot (\text{Post}_t \times \text{SuffWorkload}_w) + \text{FE} + \varepsilon$$

C) Instrumental variables (if we have adoption endogeneity)

Potential instruments for open adoption:

- GPU price shocks / capacity constraints (affect self-hosting incentives)
- Regulatory changes affecting on-prem requirements
- Cloud region availability for certain model families
- Licensing changes (policy-driven, discrete)

---

8) Estimating the "Linux Test" as a structural index

We can build an estimated index rather than a hand-scored rubric.

Step 1: latent factor model for Linux pressure

Let the Linux pressure (L_{w,t}) be a latent factor causing multiple observed proxies:

$$X_{w,t} = \Lambda L_{w,t} + u$$

where (X) includes PST, SCE, SID, DLF, GTE proxies.

Estimate (L_{w,t}) using factor analysis / PCA / SEM.

Step 2: map Linux pressure to rents

$$\Delta \log(\text{Markup}_{v,w,t}) = \rho \cdot L_{w,t} + \text{controls} + \text{FE} + \varepsilon$$

- (\rho < 0): higher Linux pressure → rent compression.

---

9) The "value migrates upward" test (wer Intelligence Capital claim)

We want to show that when the model layer commoditizes, **premium moves to orchestration/governance/workflow**.

Construct layered spend shares for each firm *f*:

- model spend share
- orchestration spend share
- governance spend share
- workflow integration spend (engineering hours)

Then test:

$$\Delta\text{SpendShare}_{f,t}^{\text{UpperLayer}} = \kappa \cdot \text{OpenExposure}_t + \text{FE} + \varepsilon$$

If (\kappa > 0), that's wer "migration up the stack" empirically.

---

10) Concrete data plan (what we'll actually need)

Minimum viable dataset (MVD)

- A standardized eval suite by workload with success rates for open + closed models over time
- Effective cost per successful task (or per resolved ticket / shipped feature)
- Vendor choice per workload for a panel of firms (even 30–50 firms can work)
- Controls: industry regulation, scale, GPU constraints, latency requirements

Sources (practical)

- Public model evals (use consistent harness; don't mix benchmarks blindly)
- Procurement records / cloud bills (for price and volume)
- Engineering telemetry (calls, retries, latencies, incidents)

---

11) What "validation" looks like (clear falsifiable predictions)

Prediction 1: Commoditization signature

In workloads where open becomes sufficient:

- closed model **price premium falls**
- **price dispersion shrinks**
- **switching rises**
- **markups compress**

Prediction 2: Platform defense signature

Where SID + GTE are high:

- premium persists despite open exposure
- churn remains low
- costs shift to governance/orchestration (not weights)

Prediction 3: Intelligence Capital compounding signature

High DLF vendors show:

- persistent or rising quality-adjusted premium
- superior performance slope over time (PerfSlope)
- increasing ecosystem capture

12) Deliverable: the "Linux Econometric Score" (LES)

Once estimated, we can publish a single score per workload/segment:

$$\mathrm{LES}_{w,t} = \hat{a} \cdot \widehat{\mathrm{PST}} + \hat{b} \cdot \widehat{\mathrm{SCE}} + \hat{c} \cdot \widehat{\mathrm{SID}} + \hat{d} \cdot \widehat{\mathrm{DLF}} + \hat{e} \cdot \widehat{\mathrm{GTE}}$$

Where coefficients come from predictive power on rent compression (not hand weights). This becomes our empirical bridge between **open substitution dynamics** and **Intelligence Capital durability**.

References

Internal Labour Market References

1. Bidwell, M., & Keller, J. R. (2014). Within or Without? How Firms Combine Internal and External Labor Markets to Fill Jobs. *Academy of Management Journal*, *57*(4), 1035–1055. https://doi.org/10.5465/amj.2012.0119
2. Cowgill, B., Davis, J. M. V., Montagnes, B. P., & Perkowski, P. (2024). Stable Matching on the Job? Theory and Evidence on Internal Talent Markets. *CESifo Working Paper No. 11120*.
3. Doeringer, P. B., & Piore, M. J. (1971). *Internal Labor Markets and Manpower Analysis*. Harvard University Press; Massachusetts Institute of Technology.
4. Foss, N. J. (2003). Selective Intervention and Internal Hybrids: Interpreting and Learning from the Rise and Decline of the Oticon Spaghetti Organization. *Organization Science*, *14*(3), 331–349. Mhttps://doi.org/10.1287/orsc.14.3.331.14946
5. Friebel, G., & Raith, M. (2022). Talent Management: The Role of Bosses. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.4162631
6. Huitfeldt, I., Kostøl, A. R., Nimczik, J., & Weber, A. M. (2021). Internal Labor Markets: A Worker Flow Approach. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.3905072
7. Jiang, K., & Wang, S. (2022). Internal labor markets with two types of promotion and two tiers of salary: theory and evidence from China. *China Economic Review*, *72*, 101756. https://doi.org/10.1016/j.chieco.2022.101756
8. Keller, J. R. (2015). Internal Markets, Personal Networks, and Human Resource Allocation: How Managers Create Value Through Internal Hiring. *Wharton Faculty Platform*.
9. Pager, D., & Pedulla, D. S. (2015). Race, Self-Selection, and the Job Search Process. *American Journal of Sociology*, *120*(4), 1005–1054. https://doi.org/10.1086/681072
10. Puranam, P., & Håkonsson, D. D. (2015). Valve's Way. *Journal of Organization Design*, *4*(2), 2–4. https://doi.org/10.1186/s41469-015-0004-2
11. Zuchowski, O., Posegga, O., Schlagwein, D., & Fischbach, K. (2016). Internal Crowdsourcing: Conceptual Framework, Structured Review and Research Agenda. *Journal of Information Technology*, *31*(2), 166–184. https://doi.org/10.1057/jit.2016.14

Power Calculation References:

Gompers, P. A. (1995). Optimal investment, monitoring, and the staging of venture capital. *Journal of Finance, 50*(5), 1461–1489. https://doi.org/10.1111/j.1540-6261.1995.tb05185.x

Gompers, P. A. (1996). Grandstanding in the venture capital industry. *Journal of Financial Economics, 42*(1), 133–156. https://doi.org/10.1016/0304-405X(96)00874-4

Gompers, P. A., & Lerner, J. (2000). Money chasing deals? The impact of fund inflows on private equity valuations. *Journal of Financial Economics, 55*(2), 281–325. https://doi.org/10.1016/S0304-405X(99)00052-5

Gompers, P., & Lerner, J. (2001). The venture capital revolution. *Journal of Economic Perspectives, 15*(2), 145–168. https://doi.org/10.1257/jep.15.2.145 (American Economic Association)

Goyal, R. (2025). Concept-to-prototype conversion rates in startups. Medium. https://medium.com/@rahulgoyal

GrowthJockey. (2025). Startup screening and innovation funnel metrics. https://www.growthjockey.com

Hall, B. H., & Lerner, J. (2010). The financing of R&D and innovation. *Handbook of the Economics of Innovation, 1*, 609–639. https://doi.org/10.1016/S0169-7218(10)01014-2

Hall, R. E., & Woodward, S. E. (2010). The burden of the nondiversifiable risk of entrepreneurship. *American Economic Review, 100*(3), 1163–1194. https://doi.org/10.1257/aer.100.3.1163

Kaplan, S. N., & Schoar, A. (2005). Private equity performance: Returns, persistence, and capital flows. *Journal of Finance, 60*(4), 1791–1823. https://doi.org/10.1111/j.1540-6261.2005.00780.x

Lerner, J. (1994). Venture capitalists and the decision to go public. *Journal of Financial Economics, 35*(3), 293–316. https://doi.org/10.1016/0304-405X(94)90035-3

Puri, M., & Zarutskie, R. (2012). On the life cycle dynamics of venture-capital- and non-venture-capital-financed firms. *Journal of Finance, 67*(6), 2247–2293. https://doi.org/10.1111/j.1540-6261.2012.01786.x

Stevens, G. A., & Burley, J. (1997). 3,000 raw ideas = 1 commercial success! *Research-Technology Management, 40*(3), 16–27. https://doi.org/10.1080/08956308.1997.11671126

Ulrich, K. T., & Eppinger, S. D. (2016). *Product design and development* (6th ed.). McGraw-Hill. https://doi.org/10.1036/0078029062