DIFFUSION OF AI INNOVATION

Strategy, Adoption, and Market Dynamics

Are Markets Overestimating the Speed of AI Adoption?

Presented by William Kline, Ph.D., CFA, CSCA
Accountant-Lawyer Alliance, LLC

Saturday, August 1, 2026

COURSE PREMISE

The Market Is Pricing a Future That Has Not Fully Arrived.

1. Equity valuations, credit spreads, and capital budgets already embed assumptions about how fast and how broadly AI will be adopted -- across firms, industries, and the workforce.

2. Much of that assumed adoption has not happened yet. Experimentation is widespread; sustained, scaled, workflow-level use remains the exception, not the rule.

3. That gap between priced-in expectations and realized adoption is the analytical thread running through every part of this course -- economics, ecosystem structure, valuation, and diffusion theory.

4. The goal is not to predict an answer, but to give you a disciplined framework for testing, for yourself, whether current prices are ahead of, behind, or roughly in line with the pace of real adoption.

COURSE FRAMEWORK, NOT A PREDICTION

What This Course Examines

1

ECONOMICS

Why the Capital Is Flowing

Capital investment, market expectations, and the investment paradox.

2

ECOSYSTEM

Where Value Is Created

Industry structure, business models, and value capture.

3

VALUATION

What the Numbers Assume

Cash flow, capital allocation, and valuation assumptions.

4

DIFFUSION

How Fast Adoption Really Spreads

Adoption processes, perceived attributes, and adopter categories.

5

CONSTRAINTS

What Could Slow Adoption Down

Energy, infrastructure, regulation, labor, and trust.

6

MEASUREMENT

How Adoption Shows Up in the Numbers

Usage, utilization, margins, cash flow, and returns on capital.

Learning Objectives

1. Explain diffusion theory as applied to the adoption of artificial intelligence.
2. Assess the factors that accelerate or slow AI adoption, including sources of resistance.
3. Evaluate the roles of trust, governance, organizational readiness, and social influence in AI adoption.
4. Apply diffusion theory to specific AI tools, organizational initiatives, and markets.
5. Consider the effects of AI adoption speed on productivity, competition, and valuations.

Course Agenda

PART 1
The Economics of AI
Capital investment, market expectations, productivity, and the investment paradox.

PART 2
The AI Industry Ecosystem
Industry structure, business models, competitive advantage, and value capture.

PART 3
Valuing the AI Investment Cycle
Cash flow, capital allocation, valuation assumptions, risk, and return.

PART 4
The Diffusion of AI Innovation
Adoption processes, perceived attributes, adopter categories, social systems, organizational readiness, and governance.

PART 5
The Constraints and Consequences of AI Diffusion
Energy, infrastructure, regulation, labor, trust, public acceptance, and responsible scaling.

PART 6
Measuring AI Diffusion Through Financial Performance
Usage, sales, capacity utilization, margins, free cash flow, and returns on capital.

YOUR PRESENTER

William Kline, Ph.D., CFA, CSCA

Associate Professor of Management, Penn State Harrisburg

Ph.D. **CFA** **CSCA**

- Ph.D. in Strategic Management, Temple University
- Recipient of four teaching awards
- Author of 14 peer-reviewed journal publications
- CFA charterholder; holds the CSCA (Certified in Strategy and Competitive Analysis) credential
- Valuation consulting experience with PwC, CBIZ Valuation Group, and Curtis Financial Group
- Independent Contractor, Content Developer at HOCK International

PART 1

The Economics of AI

Why Markets Are Investing Before Adoption Has Fully Arrived

~$800B

What six companies plan to spend building AI infrastructure in 2026 alone.

Moody's projects aggregate 2026 hyperscaler capex near $785B, heading toward roughly $1 trillion in 2027. Citi Research estimates AI-related investment now drives over half of real U.S. GDP growth.

SECTION 1 · THE AI INVESTMENT PARADOX

From ChatGPT to a Trillion-Dollar Capital Cycle

November 2022 -- ChatGPT launches and triggers an enterprise race to deploy generative AI.

Hyperscaler capital expenditure has more than doubled since ChatGPT's release (Stanford HAI, 2026 AI Index).

Global private AI investment grew 127.5% year-over-year in 2025 and is now roughly 60% of total AI investment.

Generative AI alone captured about half of all private AI funding in 2025 -- up more than 200% year-over-year.

THE CENTRAL TENSION

The AI Investment Paradox

CAPITAL IN

— 88% of organizations report regular AI use in at least one business function

— Hyperscaler capex has more than doubled since late 2022

— Global private AI investment +127.5% YoY in 2025

VALUE OUT

— Only 39% of firms report any EBIT impact -- most under 5%

— Only ~6% of firms qualify as 'AI high performers'

— PwC: technology is only ~20% of an AI initiative's value; the other 80% requires redesigning the workflow

Sound familiar? Economist Robert Solow, 1987: "You can see the computer age everywhere but in the productivity statistics."

MARKET EXPECTATIONS

Where the Money Is Expected to Go by 2030

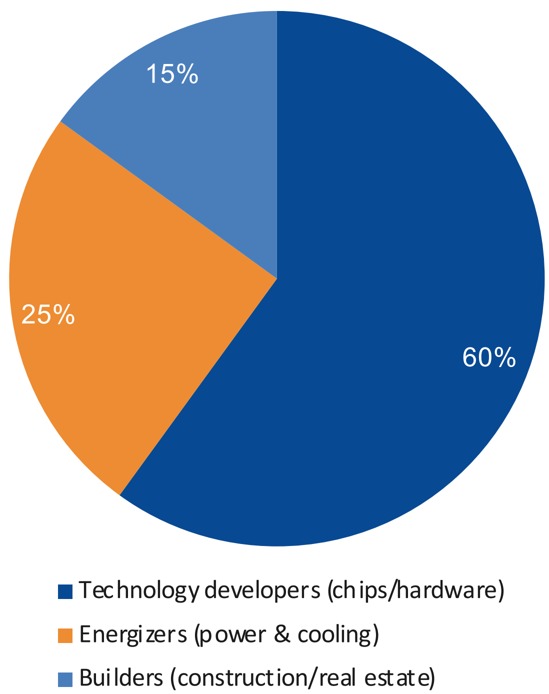

McKinsey estimates a $5.2 trillion AI-only data center buildout by 2030, split across three investor archetypes shown above.

?

AUDIENCE POLL

In your organization, has AI investment shown up yet in your P&L?

A
Cost only
(spend, no measured return)

B
Benefit only
(efficiency gains, minimal new spend)

C
Both
(spending and measurable return)

D
Neither yet
(too early to tell)

THE FINANCE LENS

Why We Should Care

1. CapEx classification and depreciation policy is now a headline accounting judgment -- Microsoft extended data center/office useful life from 15 to 25 years effective FY2027, lowering reported depreciation without reducing actual cash spend.
2. Free cash flow across hyperscalers is compressing or turning negative -- Alphabet posted its first negative quarterly FCF since its 2004 IPO in Q2 2026.
3. Debt issuance is surging to fund the buildout -- JPMorgan projects roughly $375B in hyperscaler debt issuance from 2026-2030.
4. Valuation increasingly hinges on backlog/RPO (remaining performance obligation) disclosures, not just trailing revenue.

02

PART 1 · SECTION 2

Financial Markets

How much of today's stock market is really one concentrated bet on AI?

One Trade, Two Indices

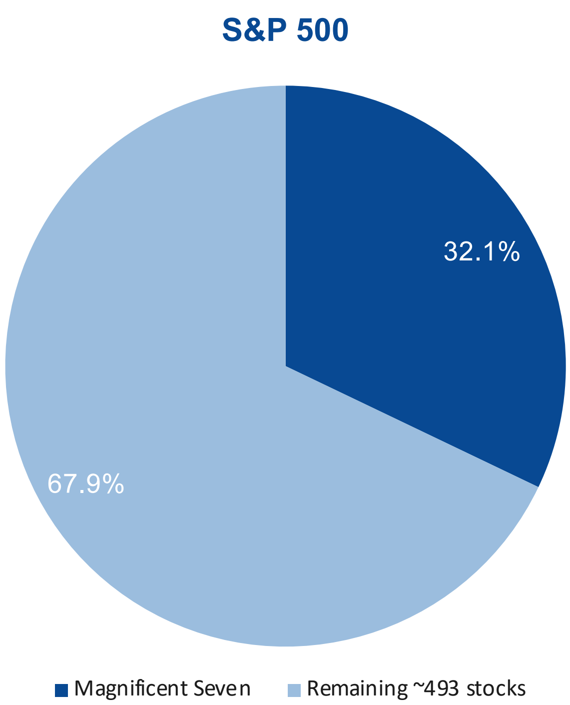

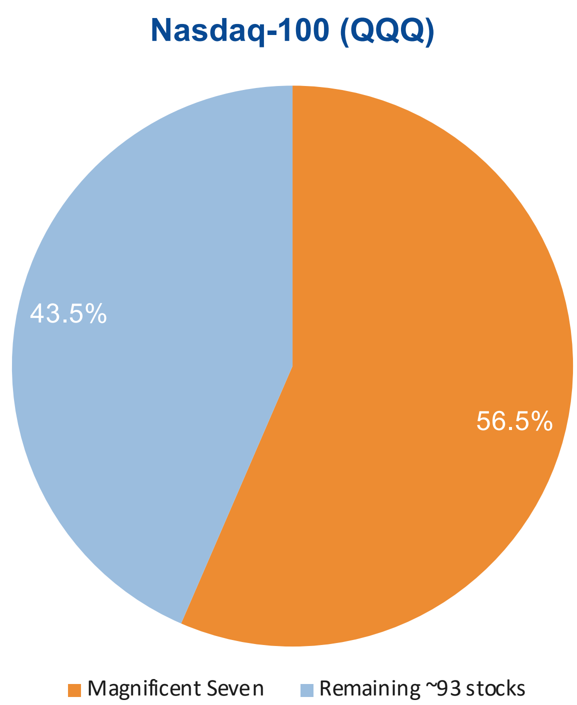

S&P 500 weights: stockanalysis.com SPY holdings data, 7/27/26. Nasdaq-100 weights: slickcharts.com, 7/30/26.

CONTRIBUTION TO RETURNS

Seven Stocks, Most of the Gains

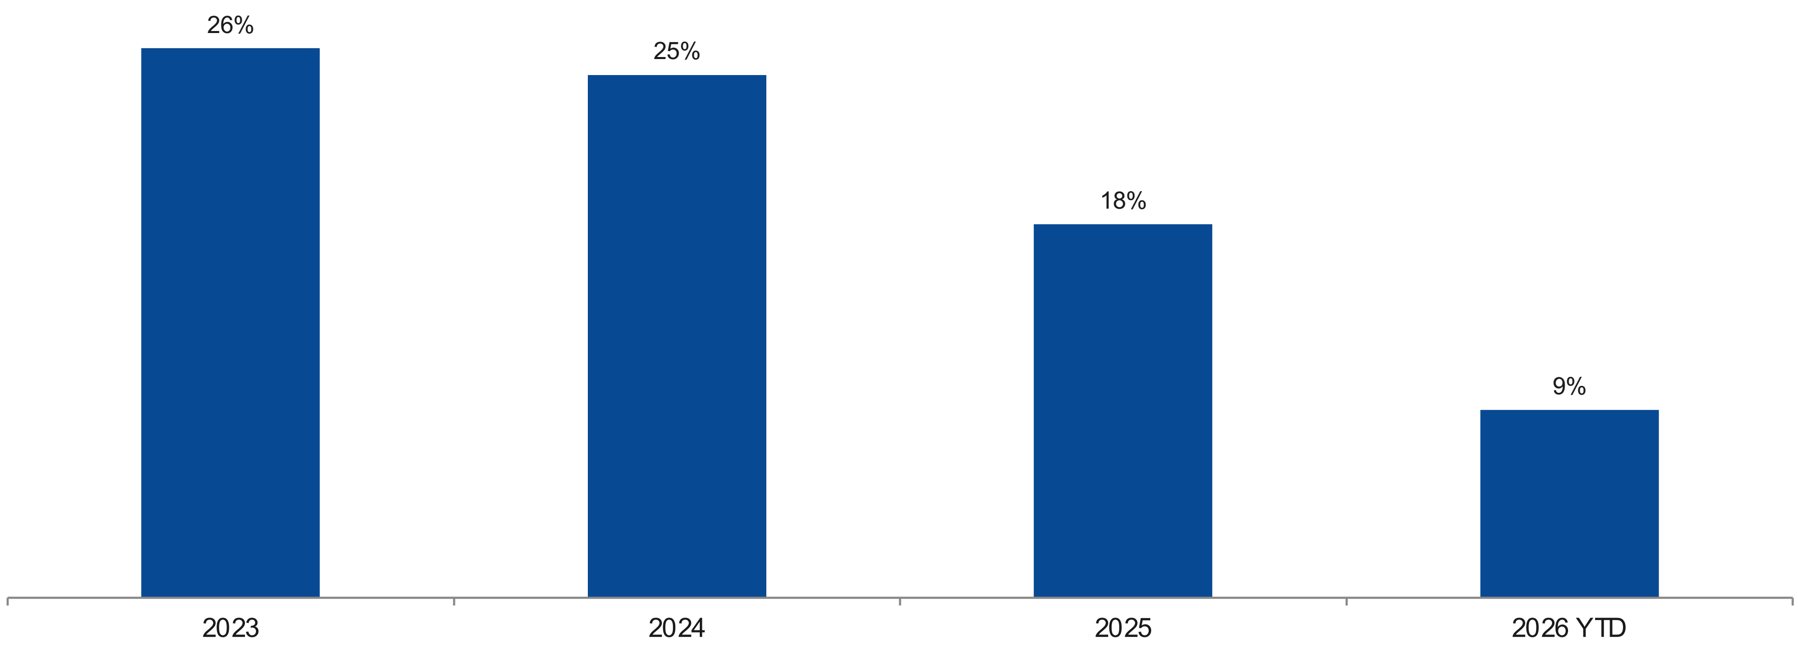

S&P 500 total annual returns (slickcharts.com, as of 7/24/26). In 2023, the Magnificent Seven delivered a 107% return and were credited with roughly two-thirds of the S&P 500's total gain (Reuters/Deutsche Bank).

CONCENTRATION RISK

$797B

Lost by the Magnificent Seven in a single stretch as AI-skeptic selling hit tech stocks (Bloomberg, via Fortune).

This followed an earlier ~$1 trillion Big Tech wipeout in February 2026 driven by the same concerns (CNBC). Concentration cuts both ways -- the same seven stocks that drive the gains also drive the drawdowns.

THE BULL CASE

Why Investors Pay High Multiples: It's About Backlog

AWS remaining performance obligations (RPO/backlog), including a new $100B+ Anthropic commitment and $138B from OpenAI.

Microsoft's RPO, up 99% year-over-year -- read by analysts as evidence of contracted future demand, not speculation.

Amazon CEO Andy Jassy: "We're not investing approximately $200 billion in capex in 2026 on a hunch... we already have customer commitments."

D.A. Davidson's Gil Luria estimates Alphabet's compute buildout is already generating $15-20B in direct profit this year.

03

PART 1 · SECTION 3

The AI Capital Spending Boom

Six companies. Roughly $800 billion in 2026. Where is it actually going on the balance sheet?

FOUNDATIONAL CONCEPT

What Is CapEx -- and Why Does It Matter Here?

OPERATING EXPENSE (OpEx)

— Expensed immediately on the income statement

— Reduces net income in the period incurred

— Example: cloud subscription fees, salaries

CAPITAL EXPENDITURE (CapEx)

— Capitalized on the balance sheet as an asset

— Depreciated/amortized over its useful life

— Example: GPUs, servers, data center buildings

The judgment call that matters most right now: how long is a GPU's "useful life"? Get that assumption wrong and reported earnings shift materially -- with no change in actual cash spent.

THE AGGREGATE PICTURE

Six Companies, Roughly $800 Billion in 2026

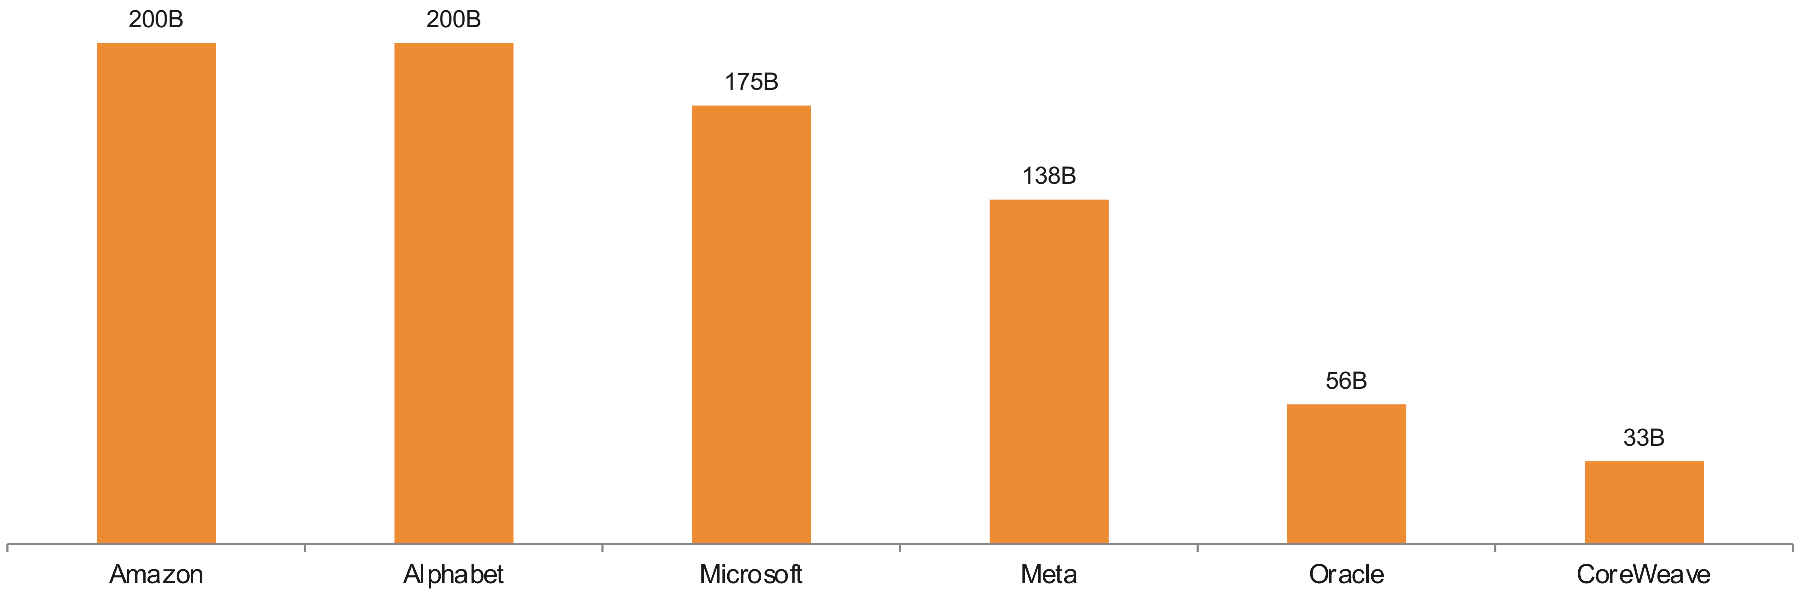

Figures are 2026 capex guidance (midpoints where a range was given) or actual, per company disclosures; fiscal years differ (Oracle FY ends May, Microsoft FY ends June). See individual company slides for exact figures and sourcing.

COMPANY DEEP DIVE

Oracle & CoreWeave: Financing the Boom With Debt

Oracle FY2026 capex: $55.7B vs. FY2025's $21.2B (+162% YoY). Long-term liabilities rose to $176.9B from $114.7B; plans to raise another $40B this fiscal year. CLSA estimates Oracle may need up to $500B by 2030, with internal cash covering only ~20%.

CoreWeave FY2025 capex: $10.3B; 2026 guidance of $31-35B against trailing-twelve-month revenue of just $6.2B. Backlog of $99.4B (including a new $21B Meta commitment), financed via more than $20B in debt and equity raised year-to-date through May 2026. Q1 2026 free cash flow: -$4.7B.

THE SKEPTIC'S CASE, PART 1

The Circular Financing Web

Microsoft ↔ OpenAI: cloud commitments and equity ties running in both directions.

Amazon ↔ Anthropic: AWS backlog includes over $100B tied to Anthropic; Amazon is also an Anthropic investor.

Google ↔ Anthropic: Alphabet discloses roughly $51B in off-balance-sheet "backstops" of other firms' data centers, per D.A. Davidson's Gil Luria.

Nvidia ↔ CoreWeave: chip supply and customer financing relationships that blur vendor and investor roles.

THE SKEPTIC'S CASE, PART 2

Credit and Macroeconomic Risk

Moody's: return on investment for roughly $785B of 2026 hyperscaler capex (heading toward ~$1T in 2027) "is unclear."

Citi Research: AI investment now drives over half of real U.S. GDP growth -- a pullback could trigger a "mild recession."

Apollo's Torsten Sløk: hyperscaler bond coverage ratios fell from roughly 5x (Feb. 2026) to below 2x (July 2026).

JPMorgan projects roughly $375B in hyperscaler debt issuance from 2026-2030; Amazon reportedly had to sweeten a recent bond sale to attract buyers.

04

PART 1 · SECTION 4

Following the Money

A complete map of the AI value chain -- and who actually captures the economic value.

THE AI VALUE CHAIN -- PART 1 OF 2

Upstream: Power, Chips & Infrastructure

Hyperscalers
Microsoft Azure, Amazon AWS, Google Cloud

GPU Designers
Nvidia, AMD, Broadcom (custom AI ASICs)

CPU Designers
Intel, AMD, Nvidia (Grace)

Semiconductor Manufacturing
TSMC, Samsung Foundry, Intel Foundry

Semiconductor Equipment
ASML, Applied Materials, Lam Research

Networking
Nvidia (Mellanox/Spectrum-X), Arista Networks, Broadcom

Memory (HBM)
SK Hynix, Samsung, Micron

Power Generation
Constellation Energy, Talen Energy, NextEra/Vistra

Utilities
Duke Energy, Entergy, Dominion Energy

THE AI VALUE CHAIN -- PART 2 OF 2

Downstream: Models, Software & End Customers

Cooling Vertiv, Schneider Electric, nVent	**Data Centers** Equinix, Digital Realty, CoreWeave	**Cloud Infrastructure** Oracle Cloud Infrastructure, CoreWeave, Lambda
Foundation Models OpenAI, Anthropic, Google DeepMind, Meta AI	**Enterprise Software** Microsoft (Copilot), Salesforce (Agentforce), SAP, ServiceNow	**Cybersecurity** CrowdStrike, Palo Alto Networks, Microsoft Security
Consulting Firms McKinsey (QuantumBlack), Deloitte, BCG	**Systems Integrators** Accenture, IBM Consulting, Infosys, TCS	**Enterprise & Consumer Customers** JPMorgan Chase, Walmart, health systems | ChatGPT, Gemini, Copilot, Midjourney

VALUE CAPTURE

Where Does the Value Actually Go?

WHERE VALUE CONCENTRATES

— Infrastructure (chips + cloud) captures the majority of dollars flowing through the stack

— Nvidia singled out as the single biggest winner -- durable moat via CUDA software lock-in and hardware lead

— McKinsey: a small number of semiconductor firms control market supply -- a structural chokepoint

WHERE MARGIN IS UNDER PRESSURE

— Application-layer companies: 50-90% gross margins, but weak differentiation and retention

— Most foundation-model providers have not yet reached large-scale commercial profitability relative to usage

— Apps and models increasingly rely on similar underlying architectures and training data -- commoditization risk rising

KEY INSIGHT: In this cycle, the picks-and-shovels layer -- chips and infrastructure -- is capturing more durable economic value than most of the applications built on top of it.

05

PART 1 · SECTION 5

Technology Primers

Ten terms, explained for a professional with zero engineering background -- no jargon left unexplained.

1

EXECUTIVE TECH PRIMER

What Is a Semiconductor?

PLAIN-ENGLISH DEFINITION

A chip made from a material -- usually silicon -- that can precisely control the flow of electrical current, forming the microscopic "switches" (transistors) that perform all computing. It matters economically because it is the physical chokepoint for all AI capability: whoever controls the supply of advanced chips controls the pace at which AI can be deployed, and captures outsized margin as a result.

THINK OF IT LIKE...

The semiconductor is the engine block of the entire digital economy -- everything else in the value chain is built on top of it.

SOURCES: Stanford AI Index 2026; industry-standard definition, verified via live search.

2

EXECUTIVE TECH PRIMER

GPU vs. CPU: Why AI Needs Parallel Processing

PLAIN-ENGLISH DEFINITION

A CPU (central processing unit) is the general-purpose "brain" of a computer -- optimized to handle a wide variety of tasks, one after another, very quickly and in the correct order. A GPU (graphics processing unit) is built to perform thousands of simple calculations simultaneously. AI training and inference are essentially massive parallel matrix math -- exactly what GPUs excel at -- which is why Nvidia became the AI industry's central chip supplier.

THINK OF IT LIKE...

A CPU is a skilled all-around manager handling tasks one at a time; a GPU is an army of specialized assembly-line workers, each doing one small calculation at once, all in parallel.

SOURCES: verified via live search, mid-2026.

3

EXECUTIVE TECH PRIMER

What Is HBM (High Bandwidth Memory)?

PLAIN-ENGLISH DEFINITION

Ultra-fast memory stacked directly next to or on top of the GPU, so data can feed the chip without creating a bottleneck. It has become a hot commodity because GPU compute power has outpaced memory supply -- AI chips are frequently "starved," waiting for data -- and only a small number of firms (SK Hynix, Samsung, Micron) can manufacture it at the required specification.

THINK OF IT LIKE...

HBM is the high-speed conveyor belt that keeps the assembly line (the GPU) constantly fed; without it, the workers stand idle no matter how fast they could otherwise work.

SOURCES: verified via live search, mid-2026. Exact market-share splits across SK Hynix/Samsung/Micron should be sourced from TrendForce before citing a specific number.

4

EXECUTIVE TECH PRIMER

What Is a Foundry?

PLAIN-ENGLISH DEFINITION

A factory that manufactures chips designed by other companies. Nvidia, Apple, and AMD design their own chips but do not build them -- they rely on foundries like TSMC to physically manufacture the silicon. This "fabless plus foundry" split is why a handful of foundries, heavily concentrated in Taiwan, are geopolitically critical to the entire global AI supply chain.

THINK OF IT LIKE...

The foundry is a specialized printing press: design firms bring the blueprint, and the foundry prints the physical chip.

SOURCES: TSMC investor relations (verified July 2026 filings); industry-standard definition.

5

EXECUTIVE TECH PRIMER

What Is Inference (vs. Training)?

PLAIN-ENGLISH DEFINITION

Training is the extremely expensive, largely one-time process of teaching a model by feeding it vast datasets so it learns underlying patterns. Inference is using that already-trained model to answer a real question or perform a task -- far cheaper per use, but it happens billions of times, so its aggregate industry-wide cost and revenue now rivals or exceeds training cost.

THINK OF IT LIKE...

Training is like years of medical school; inference is the doctor seeing patients afterward, one visit at a time.

SOURCES: verified via live search, mid-2026.

6

EXECUTIVE TECH PRIMER

What Are Tokens?

PLAIN-ENGLISH DEFINITION

The basic unit that AI language models read and generate text in -- typically a word fragment, not a whole word and not a single letter. AI vendors bill customers by the number of tokens processed, making tokens effectively the AI industry's unit of production and revenue.

THINK OF IT LIKE...

Tokens are to AI language processing what kilowatt-hours are to electricity billing, or what word-count units are to a translator's invoice.

SOURCES: verified via live search, mid-2026.

7

EXECUTIVE TECH PRIMER

What Is a Context Window?

PLAIN-ENGLISH DEFINITION

The maximum amount of text -- measured in tokens -- that a model can "see" and reason over at one time; effectively its short-term working memory. A larger context window lets an AI read an entire contract or codebase in one pass, rather than in disconnected fragments, which directly determines what business tasks are practically feasible.

THINK OF IT LIKE...

The context window is the size of the desk the AI can spread its documents across while working -- a bigger desk means it can consider more at once.

SOURCES: verified via live search, mid-2026.

8

EXECUTIVE TECH PRIMER

What Is Agentic AI?

PLAIN-ENGLISH DEFINITION

AI systems built on foundation models that can autonomously plan and execute multi-step tasks in the real world -- not just answer a single question -- often calling other tools, software, or systems along the way. Per McKinsey's November 2025 survey, 62% of firms are at least experimenting with AI agents, but scaled enterprise use remains under 10% in any given business function -- this is still an early-stage capability.

THINK OF IT LIKE...

Agentic AI is a junior employee who can be handed an entire project and complete multiple steps unsupervised, versus a chatbot that only answers one question at a time.

SOURCES: McKinsey, 'The State of AI in 2025' (Nov. 5, 2025).

9

EXECUTIVE TECH PRIMER

What Is Retrieval-Augmented Generation (RAG)?

PLAIN-ENGLISH DEFINITION

A technique where an AI model, before answering, first retrieves relevant facts from an external database or document set -- rather than relying solely on what it memorized during training -- and then generates its answer grounded in that retrieved material. This reduces "hallucination" and lets AI work accurately with a firm's private, current data, such as internal financial records.

THINK OF IT LIKE...

RAG is an open-book exam: the AI looks up the reference material before answering, instead of answering purely from memory.

SOURCES: verified via live search, mid-2026.

10

EXECUTIVE TECH PRIMER

What Is Model Context Protocol (MCP)?

PLAIN-ENGLISH DEFINITION

An open-source standard, originally introduced by Anthropic in November 2024 and now broadly adopted industry-wide (across Claude, ChatGPT, VS Code, Cursor, and others), that standardizes how AI applications connect to external data sources, tools, and workflows. Rather than every AI application needing a custom, one-off integration for every database or tool, MCP provides one common connector.

THINK OF IT LIKE...

MCP is the USB-C port of AI: one standard connector instead of a different cable for every device.

SOURCES: modelcontextprotocol.io, "Introduction" page, verified current as of July 2026.

SECTION 6 · WHY THIS MATTERS

The Adoption Gap: If the Money Is Already Here, Why Isn't Every Firm AI-First?

88% of organizations use AI regularly in at least one function -- yet only 39% report any EBIT impact, and nearly two-thirds remain stuck in pilot mode rather than scaled deployment (McKinsey).

A 12,000-firm European study found AI adoption raised labor productivity by about 4% -- real, but far short of the capital being deployed against it.

The OECD projects only a 0.2 to 1.3 percentage point annual G7 productivity gain from AI over the next decade -- a modest, gradual number relative to the scale of investment covered in this module.

END OF PART 1

Investment Is Only the Beginning.

To understand whether AI spending creates durable value, we must examine the ecosystem through which that investment flows.

PART 2

The AI Industry Ecosystem

Where Value Is Created—and Who May Capture It

01

PART 2 · SECTION 1

Defining the Industry

Before we map the ecosystem, we need shared vocabulary for what an "industry" actually is.

SECTION 1 · DEFINING THE INDUSTRY

What Is an Industry?

1. Strategy definition: an industry is a group of firms producing products or services that customers view as close substitutes for one another (Michael E. Porter, Competitive Strategy: Techniques for Analyzing Industries and Competitors, Free Press, 1980).

2. The test is the customer's, not the company's: if a buyer would swap Product A for Product B in response to a price or quality change, A and B are in the same industry -- regardless of how the two companies describe themselves.

3. AI immediately complicates this test: is Nvidia in the "semiconductor industry," the "AI infrastructure industry," or the "cloud computing industry"? The honest answer is that it depends which customer decision you are analyzing.

WHY PRECISION MATTERS

Why Industry Boundaries -- and Classification Codes -- Matter

WHAT BOUNDARY-DRAWING AFFECTS

— Competitor identification -- who actually gets compared in a research note

— Market sizing and TAM/SAM/SOM forecasts

— Valuation multiples -- the comparable-company set changes the multiple

— Regulatory and antitrust analysis -- market definition often decides the case

NAICS / SIC: USEFUL BUT DATED

— NAICS (U.S. Census Bureau) and SIC (SEC) codes classify firms for statistical and disclosure purposes, revised only every five years

— Nvidia -- SIC 3674 "Semiconductors" -- despite most current revenue coming from AI data-center systems, not legacy chips

— Microsoft -- SIC 7372 "Prepackaged Software" -- despite Azure/AI infrastructure now representing a large and growing share of revenue and capex

TAKEAWAY: Classification systems built for a pre-cloud, pre-AI economy cannot cleanly capture vertically integrated AI players -- analysts must define boundaries functionally, not by looking up a code.

STRUCTURE OF RELATIONSHIPS

Upstream, Downstream, and the Integration Decision

- ↑ Upstream: further back toward raw inputs. For a hyperscaler, upstream means chip designers, foundries, and electricity generation.
- ↓ Downstream: closer to the end customer. For a chip designer like Nvidia, downstream means hyperscalers, enterprise customers, and ultimately consumers.
- V Vertical integration: owning multiple stages of the same chain -- e.g., Google designing its own TPU chips (upstream) while also operating Google Cloud and offering Gemini models (downstream) and consumer products (further downstream).
- H Horizontal integration: expanding within the same stage -- e.g., a cloud provider acquiring another cloud provider, or a foundation-model lab acquiring a competing lab.
- O Outsourcing and strategic partnership as the alternative to integration: Nvidia does not own a foundry -- it partners with TSMC (the fabless model) rather than integrating backward into manufacturing.

THE CORE STRATEGIC DISTINCTION

Value Creation vs. Value Capture

VALUE CREATION

— The total benefit delivered to the customer relative to their next-best alternative

— A firm can create enormous customer value and still earn thin margins

— Example: foundation-model API pricing has fallen roughly 50x per year on a price-per-performance basis (Epoch AI, Mar. 12, 2025) -- enormous value delivered to developers

VALUE CAPTURE

— The share of that created value the firm actually keeps as profit

— Capture depends on bargaining power, not on how much value was created

— Example: ASML captures outsized margin (FY2025 gross margin 52.8%, ASML 2025 Annual Report) because it is the sole global supplier of EUV lithography machines -- a genuine bottleneck, not because it creates more value than every other participant

TAKEAWAY: A company that controls a scarce bottleneck can capture value far out of proportion to the value it alone creates -- this single idea is the organizing logic for the rest of this module.

02

PART 2 · SECTION 2

The Complete AI Ecosystem

Thirteen layers of capital, technology, data, and services -- from power plants to end users.

THE VALUE CHAIN AT A GLANCE

Six layers connect technical capability to real-world impact

1

Enablers

Make AI technically possible

2

Infrastructure

Compute, data, storage, connectivity

3

Models

Create intelligence & general capability

4

Platforms & Tools

Make AI deployable and manageable

5

Applications

Turn capability into usable solutions

6

End Users & Impact

Outcomes, behavior change, value

LAYER 1 OF 6

1 Enablers

This layer sets AI's physical limits: speed, cost, energy use, availability, and scale.

- **GPUs & Accelerators**
- **Memory (HBM/DDR)**
- **High-speed Interconnects**
- **Power & Cooling**
- **Sensors**
- **Semiconductor Equipment**

LEADING COMPANIES

NVIDIA | AMD | TSMC | Samsung | ASML | Broadcom

TEACHING MESSAGE

Enablers determine what's physically possible before any product decision gets made. Every downstream layer inherits these constraints.

DIFFUSION CONNECTION

When enabling technology is scarce or expensive, adoption stays concentrated among Innovators and Early Adopters who can absorb the cost and risk.

LAYER 2 OF 6

2 Infrastructure

Infrastructure turns technical components into capacity an organization can actually depend on.

Cloud Infrastructure	**Data Centers**	**Storage**	**Network Fabric & Connectivity**	**Data Pipelines**	**Energy Systems**

LEADING COMPANIES

AWS	**Microsoft Azure**	**Google Cloud**	**Oracle**

TEACHING MESSAGE

Infrastructure is what lets AI move from an impressive demo to something an operations team can actually run every day.

DIFFUSION CONNECTION

Reliable, well-supported infrastructure lowers the perceived risk of adoption — helping AI cross from Early Adopters into the Early Majority.

LAYER 3 OF 6

3 Models

Models convert data and compute into intelligence — but a model alone is potential, not finished value.

Pre-training	**Model Architecture**	**Training Data**	**Evaluation**	**Optimization**	**Safety & Alignment**

LEADING COMPANIES

OpenAI | Google DeepMind | Anthropic | Meta AI | Mistral AI

TEACHING MESSAGE

A more capable model does not automatically create business value — capability still has to be packaged, deployed, and adopted before it matters.

DIFFUSION CONNECTION

Model-level relative advantage is necessary but not sufficient — Rogers' other attributes (compatibility, complexity, observability) determine whether that advantage is ever realized.

4 Platforms & Tools

Platforms and tools reduce the complexity of building, deploying, and governing AI at scale.

- **MLOps / LLMOps**
- **Model Hosting**
- **Vector Databases**
- **Evaluation & Monitoring**
- **Security & Governance**

LEADING COMPANIES

- Databricks
- Snowflake
- Hugging Face
- Pinecone
- Weights & Biases

TEACHING MESSAGE

This layer exists to make AI usable by teams who are not model researchers — the bridge between raw capability and everyday deployment.

DIFFUSION CONNECTION

Tools directly increase two of Rogers' five attributes — compatibility and trialability — among the strongest levers for adoption rate.

5 Applications

People do not adopt "AI" — they adopt a specific solution that improves a specific task.

Copilots & Assistants	**Enterprise Software**	**Automation**	**Industry Solutions**	**Content Generation**

LEADING COMPANIES

- Microsoft Copilot
- Adobe Firefly
- Salesforce Einstein
- ServiceNow
- Cursor

TEACHING MESSAGE

The application layer is where 'jobs-to-be-done' thinking matters most — success depends on fit to a real task, not general capability.

DIFFUSION CONNECTION

This is where observability lives — a visible, well-packaged application makes AI's results legible to the next potential adopter.

6 End Users & Impact

Value appears only when real people and organizations change how they actually work.

Productivity	Cost Savings	Revenue Growth	Faster Decisions	Workforce Transformation

LEADING COMPANIES

JPMorgan Chase	Pfizer	Toyota	Walmart	Disney	U.S. Government

TEACHING MESSAGE

Deployment is not the same as adoption. Adoption is not the same as impact — each step can still fail even after the previous one succeeds.

DIFFUSION CONNECTION

This is Rogers' Confirmation stage at economy-wide scale — impact is only 'real' once it survives sustained use, not just an initial rollout.

RECOMMENDED SOURCES

Where to go deeper on the AI industry value chain

Company & Market Data

Company filings & investor relations
Gartner, IDC, Forrester
CB Insights, PitchBook

Research Institutions

Stanford HAI AI Index
McKinsey Global Institute
Academic papers & arXiv

Sector-Specific Data

SEMI, Omdia, TrendForce (semiconductors)
EIA and DatacenterMap (energy & data centers)

Company names and layer positioning are illustrative, not exhaustive, and evolve rapidly — verify current figures and market claims before presenting.

MASTER ECOSYSTEM MAP

Capital, Technology, Data, and Services Flow Upstream to Downstream

1. Energy & Physical Inputs

Constellation, Vistra, NextEra, Talen, Duke, Dominion, GE Vernova

2. Data-Center Construction & Cooling

Equinix, Digital Realty, Vertiv, Schneider Electric, Trane

3-4. Semiconductor Equipment & Foundries

ASML, Applied Materials, Lam, KLA, TSMC, Samsung, Intel Foundry

5. Chip Designers

Nvidia, AMD, Google TPU, Amazon Trainium, Microsoft Maia, Broadcom

6-7. Memory & Networking

SK Hynix, Samsung, Micron, Arista, Cisco, Broadcom

8. Hyperscalers & Cloud

Microsoft Azure, AWS, Google Cloud, Oracle OCI, CoreWeave

9. Foundation Models

OpenAI, Anthropic, Google DeepMind, Meta AI, xAI, Mistral

10. AI Platforms & Data Tools

Databricks, Snowflake, Palantir, Hugging Face, Pinecone

11. Enterprise Applications

Microsoft Copilot, Salesforce Agentforce, ServiceNow, SAP, Adobe

12. Consulting & Systems Integration

Accenture, IBM, Deloitte, PwC, McKinsey, TCS, Infosys

13. Enterprise Customers & End Users

Financial services, healthcare, manufacturing, retail, government

What Is a Hyperscaler, and How Does It Make Money?

- **1** A hyperscaler is a company that operates cloud data-center infrastructure at massive, globally distributed scale and rents that capacity to other businesses rather than using it only internally.
- **IaaS** Infrastructure as a Service: renting raw compute, storage, and networking (e.g., an Azure virtual machine) -- the customer manages everything above the hardware layer.
- **PaaS** Platform as a Service: renting a managed development environment (e.g., a managed database or model-hosting service) -- the provider manages more of the stack, the customer focuses on building.
- **RPO** Backlog / Remaining Performance Obligations (RPO): contracted future revenue not yet recognized -- the metric markets now watch most closely as evidence that AI capex is backed by real, signed demand rather than speculation (this concept was introduced in Module 1's discussion of AWS's and Microsoft's RPO figures and is not repeated here).
- **%** Utilization rate: the share of a hyperscaler's built capacity actually being used and billed -- the single biggest driver of whether massive capex translates into margin or sits idle as depreciating, unmonetized assets.

Hyperscaler Competitive Profiles

Company	Business Model	How It Uses AI	Main Business Risk
Microsoft Azure	**Internal + External**	Uses AI inside Microsoft products such as Microsoft 365 and Dynamics, while also selling cloud infrastructure, models, and AI tools to outside customers.	Depends heavily on OpenAI for some of its most differentiated AI capabilities.
Amazon Web Services	**Primarily External**	Sells cloud computing, storage, and AI infrastructure to other companies. It also develops its own AI chips to lower costs and reduce dependence on outside suppliers.	Very high investment spending may pressure cash flow before customer demand fully catches up.
Google Cloud	**Internal + External**	Uses AI across Google's own products and advertising business, while also selling cloud infrastructure, Gemini models, and AI services to outside customers.	Large infrastructure spending may take time to generate enough revenue and cash flow.
Oracle Cloud Infrastructure	**Primarily External**	Sells cloud capacity and AI infrastructure to outside customers, including companies that need large amounts of computing power.	Rapid expansion is expensive and may require substantial borrowing or customer commitments.
Meta	**Primarily Internal**	Builds AI infrastructure mainly to improve its own advertising, recommendation systems, consumer products, and future AI services. It is not primarily trying to become a general-purpose cloud provider.	Heavy spending must be justified through stronger advertising performance, engagement, and future product revenue.
CoreWeave / Lambda	**External Specialist**	Rents GPU-heavy computing capacity to AI developers and companies that need specialized infrastructure.	Relies heavily on a small number of customers, outside financing, and continued demand for premium GPU capacity.

03

PART 2 · SECTION 3

Industry Economics and Competitive Strategy

SECTION 3 · FRAMEWORKS

Porter's Five Forces: A 30-Second Refresher

1. Rivalry among existing competitors -- how intensely do incumbents compete on price, features, and marketing?
2. Threat of new entry -- how easily can a new competitor show up and take share?
3. Supplier power -- can the firms that sell inputs to this industry dictate price and terms?
4. Buyer power -- can the firms' own customers dictate price and terms back to them?
5. Threat of substitutes -- is there a fundamentally different way to solve the same customer problem?

Graphic recommendation: the standard Five Forces pentagon, with 'Industry Rivalry' in the center box and the four surrounding forces as arrows pressing inward.

The GPU / AI Accelerator Industry

Rivalry	LOW-MED	Nvidia dominant; AMD and hyperscaler custom silicon (TPU, Trainium, Maia, MTIA) are rising but not yet at parity on frontier training workloads.
Threat of Entry	LOW	Extreme capital intensity, multi-year design cycles, and dependence on ASML/TSMC's own scarce capacity make new merchant entry nearly impossible.
Supplier Power	HIGH	ASML (sole EUV supplier) and TSMC (advanced-node concentration) hold real leverage over every chip designer, including Nvidia itself.
Buyer Power	MED-HIGH	A handful of hyperscalers are the overwhelming majority of demand -- and several are simultaneously building substitute in-house chips (Section 2, Layer 5).
Threat of Substitutes	MODERATE	Custom ASICs are a direct, growing substitute for merchant GPUs in specific, well-defined workloads -- rising but not yet dominant.

VERDICT: Structurally favorable overall -- low entry threat and rivalry protect margin -- but Nvidia's own suppliers (ASML, TSMC) and largest customers (hyperscalers building substitutes) both hold real counter-leverage.

Foundation Models

Rivalry	HIGH	OpenAI, Anthropic, Google DeepMind, Meta, xAI, Mistral, Cohere, and DeepSeek all compete on overlapping capability within months of one another.
Threat of Entry	MODERATE	Frontier-scale entry requires billions in compute, but open-weight releases (Llama, DeepSeek, Mistral) let smaller players enter at the mid-tier with far less capital.
Supplier Power	HIGH	Labs depend on hyperscaler compute (OpenAI-Microsoft, Anthropic-Amazon/Google) and, indirectly, on Nvidia/custom-silicon supply -- real dependency, not choice.
Buyer Power	RISING	API pricing has fallen roughly 50x per year on a price-per-performance basis (Epoch AI, Mar. 12, 2025); enterprise buyers can and do switch between model providers with modest integration cost.
Threat of Substitutes	HIGH	Open-weight models (Llama, DeepSeek-R1, Mistral) are increasingly credible substitutes for many commercial use cases, per the commoditization debate covered in Section 2.

VERDICT: Structurally the least favorable of the three segments analyzed here -- high rivalry, high supplier dependency, and a live, credible substitution threat from open-weight models all compress durable margin.

Enterprise AI Applications

Force	Rating	Rationale
Rivalry	MODERATE	Incumbents (Microsoft, Salesforce, ServiceNow, SAP, Adobe, Workday) compete mainly by embedding AI into products customers already own, not head-to-head on price.
Threat of Entry	LOW-MED	AI-native challengers can build a feature quickly, but winning enterprise trust, security review, and a net-new budget line is slow and expensive (Section 2, Layer 11).
Supplier Power	MODERATE	Application vendors depend on foundation-model APIs, but most maintain multi-model flexibility, limiting any single model supplier's leverage over them.
Buyer Power	LOW-MED	Enterprise customers face real switching costs once AI is embedded in existing workflow software they already depend on -- the opposite dynamic from the model layer.
Threat of Substitutes	LOW	Once AI is embedded in a system of record, a generic substitute cannot easily replicate the underlying workflow and data integration.

VERDICT: The most structurally favorable of the three segments -- low substitute threat and low-to-moderate buyer power let incumbents capture durable value even though the model underneath is increasingly commoditized.

FIVE FORCES ANALYSIS 4 OF 4

Semiconductor Equipment (ASML, TSMC)

Rivalry	LOW	ASML has no true rival in EUV lithography -- it is the sole global supplier; TSMC's closest foundry rivals (Samsung Foundry, Intel Foundry) trail by one or more full process generations at the leading edge.
Threat of Entry	LOW	Decades of accumulated, largely tacit process know-how plus tens of billions in R&D and fab capital make new entry at the leading edge essentially not viable within any relevant planning horizon.
Supplier Power	MODERATE	ASML itself depends on a small number of irreplaceable sub-suppliers (e.g., Zeiss for EUV optics); TSMC depends on ASML for its own EUV tools -- concentration cascades up the chain, not just down it.
Buyer Power	LOW	Demand for leading-edge capacity exceeds supply; chip designers and hyperscalers pre-commit capital and capacity years in advance, the reverse of a buyer-power dynamic (Section 2, Layer 5).
Threat of Substitutes	LOW	No viable alternative to EUV lithography exists at the leading edge today; advanced packaging and chiplet strategies reduce the need for the newest nodes at the margin, but do not replace the requirement outright.

VERDICT: The most structurally dominant of the four segments analyzed -- near-total barriers to entry, essentially no substitute technology, and enough scarcity that customers, not suppliers, absorb most of the pricing risk. This is the ecosystem's ultimate chokepoint.

THE PATTERN ACROSS SEGMENTS

Why Profitability Differs -- and the Barriers That Explain It

Segment	Five Forces Verdict	Key Barrier to Entry Protecting It
GPU / AI Accelerators	Favorable, but with real supplier (ASML/TSMC) and buyer (hyperscaler) counter-leverage	Capital requirements + IP/talent + CUDA ecosystem lock-in (learning curve + installed base)
Foundation Models	Least favorable -- high rivalry, high supplier dependency, credible substitutes	Capital + technical talent only -- weak on distribution, data, or installed base
Enterprise Applications	Very favorable -- low substitute threat, workflow lock-in suppresses buyer power	Distribution + installed base + customer trust + switching costs, not raw technology
Semiconductor Equipment (ASML, TSMC)	Most favorable - the ultimate scarce chokepoint	Extreme capital + decades of accumulated process know-how -- the highest barrier in the entire chain

04

PART 2 · SECTION 4

Who Captures the Profit Pool?

Revenue size and profit capture are not the same thing.

The Profit Pool Concept

REVENUE SIZE

— How much money flows through a layer in total

— A layer can be enormous in revenue terms and still capture very little as profit if competition, supplier power, or buyer power compress margin

— Example: foundation-model API revenue is large and fast-growing, but per-token pricing has fallen roughly 50x/year (Epoch AI, Mar. 12, 2025, per Section 2)

PROFIT CAPTURE

— How much of that revenue converts into durable, defensible margin

— Determined by the Section 3 frameworks combined: barriers to entry, switching costs, bottleneck control, and platform position

— Example: ASML's revenue is a fraction of the total AI ecosystem's revenue, yet its 52.8% gross margin (ASML 2025 Annual Report, per Section 2) reflects near-total profit capture within its own narrow bottleneck

THE PROFIT POOL CONCEPT (McKinsey, 'How to Think Strategically About Sequential Innovation,' and related profit-pool literature): map an entire industry chain by both revenue AND profit simultaneously -- the two maps often look completely different.

THE COMPARATIVE FRAMEWORK

Profit-Pool Scorecard Across the AI Value Chain

Layer	Gross Margin Profile	Expected Commoditization Risk
Energy / Utilities	Regulated: moderate, stable; Merchant (Talen/Vistra): variable, deal-dependent	Low -- physical scarcity of power persists
Data Centers / Cooling	Moderate; Vertiv Q1 2026 gross margin trending higher on AI-driven mix (Section 2)	Low-Moderate
Semiconductor Equipment (ASML)	Very high -- 52.8% gross margin (Section 2)	Very low -- among the most durable positions in the chain
Foundries (TSMC)	High -- 67.7% gross margin (Section 2)	Low -- geographically and technically concentrated
GPU / Chip Designers (Nvidia)	Very high -- ~75% gross margin (Section 2)	Moderate -- custom silicon and AMD are eroding share at the margin
Memory (HBM)	High and rising -- Micron Cloud Memory unit 66% gross margin (Section 2), a historic reversal for a traditionally cyclical, commodity business	Moderate -- memory has a long history of severe cyclicality
Hyperscalers / Cloud	Moderate and compressing near-term on capex (Module 1: negative/falling FCF across the group)	Low at the infrastructure layer; Moderate at the model-hosting layer
Foundation Models	Low-to-moderate and falling -- 50x/year price-performance decline (Section 2/3)	High -- the most commoditization-exposed layer in this scorecard
Enterprise Applications	High for incumbents embedding AI into existing products (Section 2, Layer 11)	Low -- protected by switching costs even as the model underneath commoditizes
Consulting / Systems Integration	Moderate-to-high on a labor-cost basis; largely undisclosed at the deal level (Section 2, Layer 12)	Low-Moderate -- implementation labor doesn't commoditize as fast as the model

All figures per Section 2/3 sourcing already cited in this module (company filings and named research, 2025-2026). Customer/supplier concentration and recurring-revenue detail for each layer are covered qualitatively in Sections 2-3 rather than repeated numerically here.

THE STRATEGIC MAP

Bottleneck Control vs. Ability to Capture Customer Value

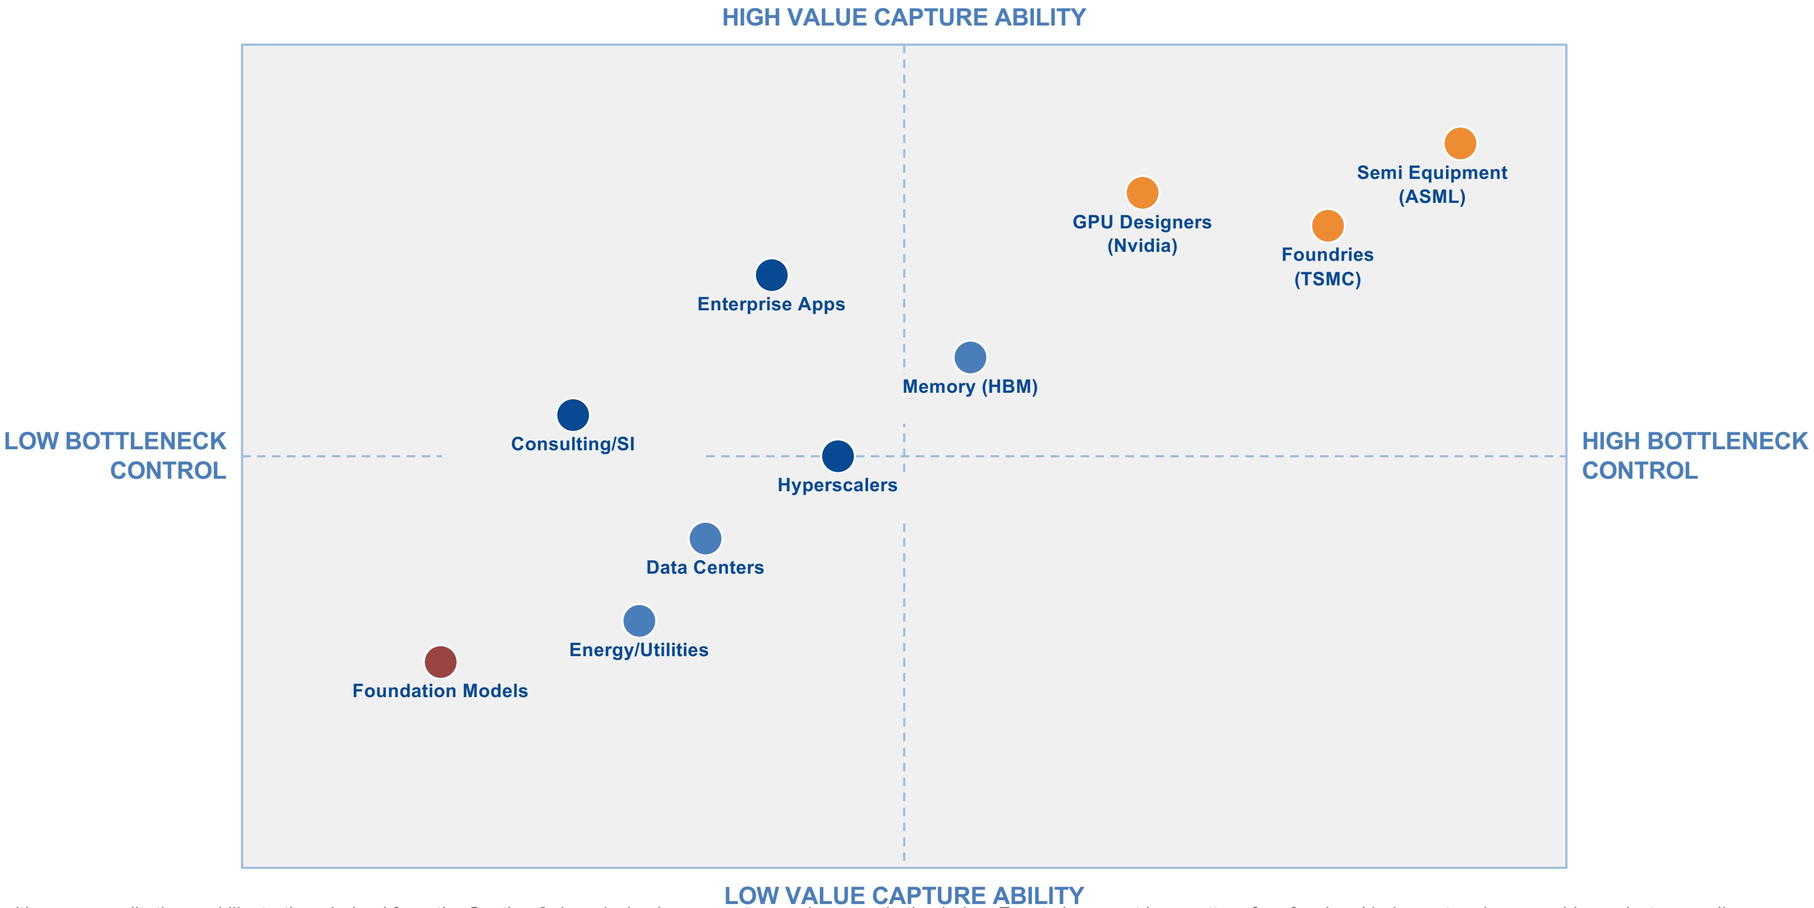

Positions are qualitative and illustrative, derived from the Section 3-4 analysis above -- not a precise quantitative index. Exact placement is a matter of professional judgment and reasonable analysts may disagree.

FRAMING THE ANSWER

The Questions Behind the Map

1. Who controls scarce assets? -- ASML and TSMC most unambiguously; increasingly, whoever controls electricity and grid interconnection capacity (Section 2, Layer 1).
2. Who owns the customer relationship? -- Enterprise-application incumbents and hyperscalers, far more than foundation-model labs, most of whose enterprise customers reach them indirectly.
3. Who owns the data? -- Enterprise customers themselves, in principle; in practice, whichever vendor's system of record captures and retains that data operationally.
4. Who controls distribution? -- Hyperscalers (cloud marketplaces) and large enterprise-software incumbents (existing seat bases), more than any single model provider.
5. Who has pricing power? -- Semiconductor equipment and foundries most clearly; foundation models least clearly, per the Section 3-4 evidence.
6. Who bears the most capital risk? -- Hyperscalers and neoclouds (CoreWeave/Lambda), given the capex and leverage levels documented in Module 1.
7. Who may become commoditized? -- Foundation models are the highest-risk candidate on current evidence, though the Section 2 commoditization debate remains genuinely open.

HISTORICAL PERSPECTIVE

"Picks and Shovels": AI Compared to Earlier Infrastructure Booms

Historical Boom	The "Picks and Shovels" Winners	The Caution
Railroads (mid-to-late 1800s U.S.)	Steel rail manufacturers and equipment suppliers often outearned many individual railroad operators, several of which overbuilt and went bankrupt in recurring panics (1873, 1893)	Not every equipment supplier won either -- overcapacity hit suppliers eventually too, once the building wave slowed
Electrification (early-to-mid 1900s)	Electrical equipment makers profited steadily for decades; economist Paul David documented a multi-decade lag between electrification and measurable productivity gains ("The Dynamo and the Computer," American Economic Review Papers & Proceedings, 1990)	The productivity payoff took roughly three to four decades to fully materialize -- a direct historical precedent for Module 1's AI investment/productivity paradox
Telecom Fiber Buildout (late 1990s-2002)	Equipment makers (Cisco, Nortel, Lucent) profited heavily during the buildout phase	Massive overbuild led to a subsequent bust -- Global Crossing and WorldCom both filed for bankruptcy (2002); fiber capacity took years to be absorbed by real demand
Internet Infrastructure / Dot-Com (1995-2001)	Some infrastructure and equipment providers survived and thrived after the crash even as many dot-com applications failed	Distinguishing genuine infrastructure winners from speculative excess was extremely difficult in real time -- an explicit caution for AI today

05

PART 2 · SECTION 5

How the AI Industry May Evolve

Six scenarios, not one forecast -- industry structure is still being written.

1

SCENARIO

Compute Remains Scarce

If chip, power, and fabrication capacity continue to lag demand, today's bottleneck owners keep their leverage.

Nvidia
Sustains premium pricing and high margin (Section 2/3) as allocation, not price, remains the binding constraint.

Foundries (TSMC)
Advanced-node capacity remains the industry's tightest chokepoint; customer bargaining power stays limited.

Memory (HBM)
High-margin HBM economics (Micron, Section 2) persist rather than reverting to historical commodity cyclicality.

Networking
Specialized interconnect suppliers (Section 2, Layer 7) retain pricing power as clusters keep scaling.

Data Centers
Colocation and construction firms retain leverage as available capacity stays the constraint, not demand.

Utilities
Power purchase agreements (Section 2, Layer 1) keep commanding premium terms from capacity-constrained buyers.

2

SCENARIO

Compute Becomes Abundant and Cheap

If chip supply, power buildout, and efficiency gains outpace demand growth, today's scarcity-based advantages erode.

Model Pricing

Foundation-model API prices, already falling ~50x/year (Section 2/3), fall further and faster.

Application Innovation

Cheap compute lowers the cost of experimentation, likely accelerating new application-layer entrants.

Enterprise Adoption

Lower cost per unit of AI capability could finally close some of Module 1's investment-versus-return gap.

Hyperscaler Margins

Infrastructure margins compress as scarcity-based pricing power fades; competition shifts toward service quality and switching costs.

New Entrants

Lower capital barriers could allow smaller players to compete meaningfully for the first time since the buildout began.

3

SCENARIO

Foundation Models Commoditize

Open-weight models and falling API prices erase differentiation at the model layer -- the trend already visible in Section 2's commoditization debate accelerates.

Open-Weight Models

DeepSeek, Llama, and Mistral (Section 2) continue narrowing the capability gap with closed frontier labs.

API Pricing

Continued rapid price declines make the model itself a shrinking share of total AI-solution cost.

Differentiation Shifts

Competitive advantage moves from 'whose model' to 'whose data, workflow, and distribution' (Section 3-4).

Winners

Enterprise applications, data platforms, and consulting/SI (Section 2, Layers 10-12) gain relative importance as the model becomes interchangeable.

4

SCENARIO

A Few Models Dominate

The opposite path: scale economies, regulation, and capital requirements consolidate the frontier-model layer into a small oligopoly.

Scale Economies

Only a handful of labs can fund frontier-scale training runs, following the economies-of-scale logic from Section 3.

Regulation

Safety and governance requirements could raise compliance costs in ways that favor large, well-resourced incumbents.

Ecosystem Lock-In

Developer tooling and enterprise integration (Section 2, Layer 10-11) increasingly build around a small number of dominant model APIs.

Concentration Risk

Regulators and enterprises alike face growing single-vendor dependency risk, raising antitrust and business-continuity questions.

5

SCENARIO

Agentic AI Changes the Software Industry

If AI systems increasingly complete multi-step tasks autonomously rather than just answering questions, enterprise software's entire pricing logic may shift.

Pricing Shift

Seat-based pricing (pay per human user) faces pressure from usage-based and outcome-based models (pay per task/result), both already emerging (Section 2, Layer 11).

Autonomous Workflows

Software increasingly executes processes rather than merely assisting a human doing them.

Labor Substitution

Some categories of knowledge work shift from human-performed to AI-performed, with humans moving toward supervision and exception-handling.

New Governance Needs

Autonomous action requires new controls: audit trails, permissioning, and accountability frameworks that don't fully exist yet.

Software Demand

Net effect on total enterprise software spending is genuinely unclear.

6

SCENARIO

Physical Infrastructure Becomes the Binding Constraint

Rather than chips or models, the ultimate limiting factor becomes the physical world: power, water, land, and permitting.

Power & Transmission

Grid interconnection queues and transmission buildout (Section 2, Layer 1) become the slowest-moving link in the entire chain.

Water & Cooling

Data-center water use for cooling draws increasing regulatory and community attention in water-stressed regions.

Permitting & Opposition

Local permitting delays and community opposition to new data centers and power plants slow buildout independent of capital availability.

Construction Labor

Skilled trades and construction-labor shortages (Section 2, Layer 2) constrain how fast new capacity can physically be built.

Semiconductor Capacity

Even with abundant capital, fab capacity (Section 2, Layer 4) still takes years to expand, creating a hard physical floor under any recovery.

SIX SCENARIOS, ONE DISCIPLINE

Hold Multiple Futures at Once

These six scenarios are not mutually exclusive, and this module does not pick a winner. Different layers of the ecosystem may experience different scenarios simultaneously -- compute could remain scarce for frontier training while becoming abundant for routine inference, in the same year, inside the same company.

The discipline that matters is not picking the right scenario in advance -- it is identifying, for your own organization, which scenario you are most exposed to, and building optionality against being wrong.

06

PART 2 · SECTION 6

Executive Synthesis

From ecosystem map to strategic action.

Module 2 Executive Summary

KEY TAKEAWAYS

— Industry structure, not technology quality alone, determines who captures value from AI

— Bottleneck owners (ASML, TSMC, Nvidia) currently capture the most durable margin in the chain

— Foundation models face the least favorable competitive structure despite the largest headlines and valuations

— Switching costs and workflow lock-in let incumbent enterprise software capture value even as the model underneath commoditizes

IMPLICATIONS & QUESTIONS TO WATCH

— Investors: revenue growth and headline excitement are not reliable proxies for durable profit capture -- use the Section 4 scorecard logic instead

— Managers: identify which layer(s) your organization actually sits in, and which bottlenecks or complements determine your own success (Strategic Diagnostic, next slide)

— Watch: does compute scarcity persist or resolve (Scenarios 1-2)? Does the model layer commoditize further (Scenario 3) or consolidate (Scenario 4)?

— Watch: which vertically integrated players (Google, Amazon, Microsoft, Meta) prove the efficiency case, and which prove the dependency/concentration-risk case?

AUDIENCE POLL

Which layer of the AI ecosystem will capture the most durable economic value over the next ten years?

A
Chips & semiconductor equipment (Nvidia, ASML, TSMC)

B
Energy & data centers (utilities, colocation)

C
Hyperscalers & cloud (Azure, AWS, Google Cloud)

D
Foundation models (OpenAI, Anthropic, DeepMind)

E
Enterprise software (Copilot, Salesforce, ServiceNow)

F
Consulting & implementation (Accenture, IBM, Big Four)

PART 3

Valuing the AI Investment Cycle

Cash Flow, Capital Allocation, and the Expectations Embedded in Market Prices

Module 3 turns from strategic structure to financial discipline: free cash flow, capital allocation, DCF valuation, market multiples, earnings calls, 10-Ks and 10-Qs, and AI investment returns -- testing whether today's prices already reflect everything this module has just established.

01

PART 3 · SECTION 1

The Financial-Statement Mechanics of AI Investment

Before we value anything, we need to see exactly where AI spending shows up -- and where it hides.

SECTION 1 · MECHANICS

Where AI Spending Appears -- and CapEx vs. OpEx

- **BS** Balance sheet: property/equipment (owned servers, GPUs, data centers), right-of-use lease assets, and, in some structures, capitalized internal-use software.
- **IS** Income statement: depreciation/amortization expense, cloud-service and inference costs (typically expensed as incurred), R&D, and compensation for AI-related headcount.
- **CF** Statement of cash flows: capital expenditures appear in investing activities; cloud/subscription AI costs typically appear in operating activities -- the same underlying 'AI cost' can hit different statement sections depending on how it's structured contractually.
- **FN** Footnotes: commitments and contingencies, lease disclosures (ASC 842), segment reporting, and depreciation-policy/useful-life disclosures often contain the most decision-relevant AI detail -- not the face financials.
- **!** IMPORTANT: whether a specific AI cost is capitalized or expensed depends on the applicable accounting standard and the specific facts and contractual terms involved. This module explains the general framework; it is not accounting or legal advice, and specific treatment should be confirmed with a qualified accountant for any real transaction.

Graphic recommendation: a simplified three-statement diagram with arrows showing a single AI investment (e.g., a GPU cluster purchase) flowing into the balance sheet (asset), income statement (depreciation), and cash flow statement (capex, investing activities) simultaneously.

A SHARPER QUESTION THAN "HOW LONG DOES IT LAST?"

GPU Obsolescence: Physical Life vs. Economic Life

PHYSICAL LIFE

— Whether the chip still functions -- runs, computes, doesn't fail

— Data-center hardware can often remain physically operational for a decade or more with proper maintenance

— Physical life is largely an engineering question, not a financial one

ECONOMIC LIFE

— Whether the chip remains cost-competitive to keep running versus replacing it with a newer generation

— New chip generations arrive with large performance and energy-efficiency gains, compressing the economic (not physical) life of the prior generation

— Resale value, software/framework compatibility, and utilization all factor into whether continuing to run an asset still makes economic sense

THE QUESTION FOR THIS AUDIENCE: an asset can be fully functional and still be uneconomic to keep running -- impairment testing exists precisely to capture this gap between physical and economic life.

STRUCTURING THE INVESTMENT

Leasing vs. Ownership: Same Compute, Different Balance Sheet

OWNERSHIP & FINANCE LEASES

— Owned data centers and finance leases: asset (or right-of-use asset) and corresponding liability both appear on the balance sheet under ASC 842, Leases

— Under ASC 842 (effective for public companies since fiscal years beginning after Dec. 15, 2018), operating leases ALSO now require balance-sheet recognition of a right-of-use asset and lease liability -- unlike the pre-2019 regime, where operating leases stayed off-balance-sheet

— Full exposure to residual value, obsolescence, and impairment risk sits with the owner/lessee

CAPACITY COMMITMENTS & COLOCATION

— Colocation agreements: renting physical space and power in someone else's data center -- infrastructure risk shifts partly to the colocation provider

— Cloud capacity reservations and take-or-pay contracts: contractual commitment to pay for capacity whether or not it is fully used -- creates a real financial obligation even without owning a physical asset

— These structures affect cash-flow TIMING (upfront capex vs. spread contractual payments) and RISK ALLOCATION, but under ASC 842 most no longer let a company avoid balance-sheet recognition entirely

TAKEAWAY: Since ASC 842, the 'lease vs. buy' decision changes cash-flow timing, risk allocation, and flexibility far more than it changes whether the obligation shows up on the balance sheet at all.

THE NUMBER EVERYONE IS WATCHING

The Free Cash Flow Bridge

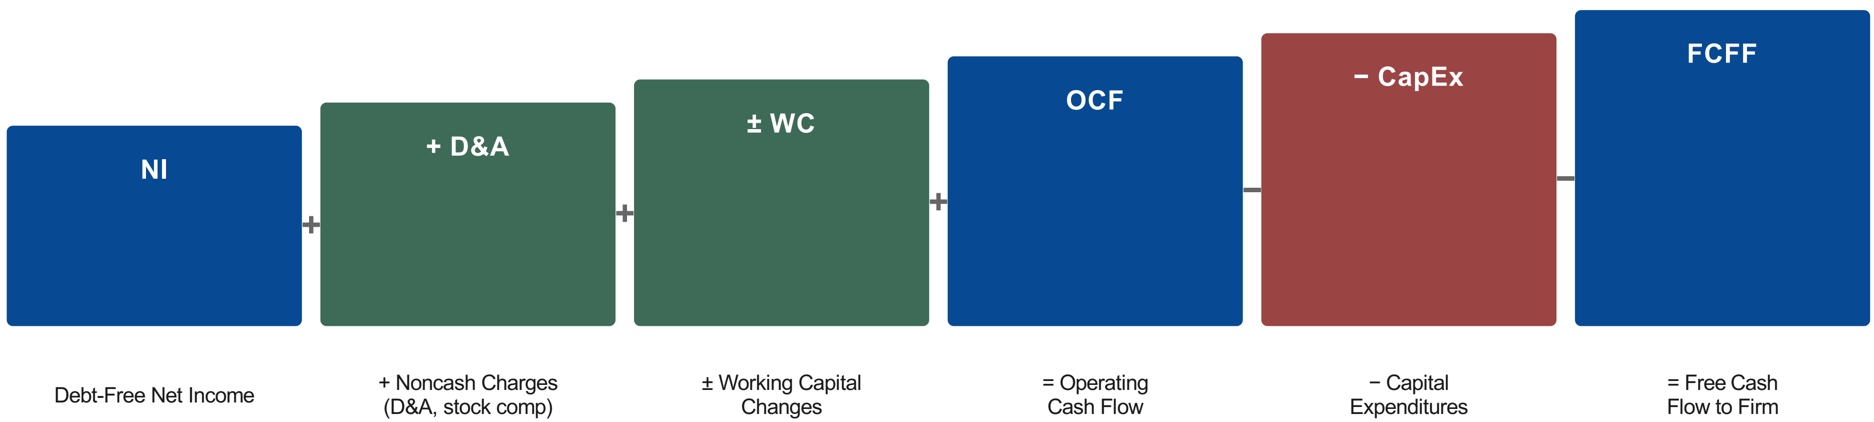

This is the most common convention (OCF less CapEx) -- Section 2 introduces two additional conventions (free cash flow to the firm and free cash flow to equity) and specifies which one is used on every subsequent slide.

SECTION 1 SUMMARY

Executive Summary: Financial-Statement Mechanics

KEY TAKEAWAYS

— AI spending touches every financial statement simultaneously -- and the footnotes often carry the most decision-relevant detail

— Useful-life assumptions are a real, scrutinizable lever on reported earnings, independent of actual cash spent

— Physical life and economic life are different questions -- impairment testing exists to bridge that gap

— Since ASC 842, leasing rarely avoids the balance sheet entirely -- it mainly changes cash-flow timing and risk allocation

IMPLICATIONS & QUESTIONS TO INVESTIGATE

— Accounting: treatment always depends on the specific standard and facts -- this module is a framework, not a substitute for professional accounting judgment

— Investors: read the commitments, leases, and depreciation-policy footnotes before trusting the income statement alone

— Questions to watch: whose useful-life assumptions are most aggressive relative to actual technology turnover? Whose off-balance-sheet-adjacent commitments are growing fastest?

02

PART 3 · SECTION 2

Free Cash Flow and the AI Capital Cycle

One number, three conventions, and why AI spending is compressing it across the industry.

DEFINITIONS DISCIPLINE

What Is Free Cash Flow? -- and Why It Matters

THREE CONVENTIONS (PICK ONE, STATE IT)

— Operating cash flow less CapEx -- the simplest, most commonly reported convention (used by Module 1's hyperscaler figures and by this module unless stated otherwise)

— Free cash flow to the firm (FCFF): cash available to ALL capital providers (debt + equity) before financing effects

— Free cash flow to equity (FCFE): cash available to equity holders specifically, after debt service -- different analysts mean different things by 'FCF,' so always confirm the convention

WHY FCF MATTERS

— Funds debt repayment and interest coverage

— Funds dividends and share repurchases

— Funds acquisitions without needing new financing

— Funds further reinvestment -- the core question for this module

— Preserves financial flexibility during downturns

— Feeds directly into DCF valuation (Section 4)

CONVENTION USED IN THIS MODULE: operating cash flow less capital expenditures (the same convention used throughout Module 1's hyperscaler analysis), unless a slide explicitly states otherwise.

THE MECHANISM

Why AI Spending Compresses Free Cash Flow

1. Data-center construction, servers, GPUs, networking equipment, and energy infrastructure all require large upfront cash outflows with long lead times -- often 18-36 months from commitment to revenue-generating capacity.
2. Capacity is built BEFORE the associated revenue exists -- hyperscalers are spending against expected future demand, not billing for capacity already earning revenue (Module 1: AWS backlog of $364B, Microsoft RPO of $678B, +84% YoY, per Module 1/2 sourcing).
3. Working-capital requirements (vendor deposits, prepaid capacity, inventory of chips/components) add further near-term cash drag ahead of revenue recognition.
4. Lease and take-or-pay purchase commitments (Section 1) create contractual cash obligations timed independently of when revenue actually materializes.
5. Net effect, industry-wide in 2026: Amazon's trailing-twelve-month FCF fell from $25.9B to $1.2B; Alphabet posted its first-ever negative quarterly FCF (-$5.9B); Meta's FCF fell 81% YoY (all per Module 1) -- while revenue grew robustly across all three.

THE CENTRAL INTERPRETIVE QUESTION

Negative Free Cash Flow: Warning Sign or Growth Signal?

POTENTIALLY CONSTRUCTIVE

— High-return reinvestment opportunity, not just spending for its own sake

— Backed by contracted demand (growing backlog/RPO), not speculation

— Genuine capacity constraints -- the company could sell more if it had more capacity

— Strong balance sheet able to fund the gap without excessive leverage

— Rising utilization trajectory and durable competitive advantage

POTENTIALLY DESTRUCTIVE

— Speculative capacity built ahead of any confirmed demand signal

— Weak capital discipline -- spending because competitors are spending

— Low or falling utilization of existing capacity

— Rapid technological obsolescence risk (Section 1) undermining the asset base

— Excess leverage funding the gap, or reliance on vendor financing

— Thin or deteriorating disclosure quality around what the capital is actually buying

THE ANALYST'S JOB: negative FCF is a symptom, not a verdict -- the diagnosis depends entirely on backlog quality, balance-sheet strength, and utilization trajectory, all covered later in Sections 4 and 6.

THE CENTRAL PRINCIPLE OF THIS ENTIRE MODULE

ROIC > WACC

Growth creates value only when the incremental return on invested capital exceeds the cost of that capital. Growth below the cost of capital destroys value, no matter how fast revenue grows.

ROIC = NOPAT (net operating profit after tax) ÷ invested capital. WACC = weighted average cost of debt and equity capital. Economic profit = NOPAT − (WACC × invested capital) -- when this is negative, the business is growing its way into value destruction, not value creation, regardless of headline revenue growth. Apply this test directly to every hyperscaler capex figure from Module 1 before assuming bigger spending is automatically better.

03

PART 3 · SECTION 3

Capital Budgeting for AI Investments

From company-level ROIC to a single project decision: should we make this specific AI investment?

NET PRESENT VALUE

NPV = present value of future net benefits minus initial investment.

INTERNAL RATE OF RETURN (IRR)

The discount rate at which a project's NPV equals exactly zero -- the project's own break-even rate of return.

PAYBACK

How long until cumulative cash inflows recover the initial investment, but ignores the time value of money and any cash flows after the payback point.

TWO KINDS OF BENEFIT, ONE FULL COST PICTURE

Efficiency vs. Growth Use Cases -- and Total Cost of Ownership

- **EFF** Efficiency use cases (labor savings, automation, reduced errors, faster cycle time, lower support costs) are generally easier to measure -- they show up as a countable reduction in an existing cost line.
- **GRO** Growth use cases (new products, better pricing, improved conversion, personalization, retention, faster innovation) are harder to measure because they require isolating AI's specific contribution from everything else affecting revenue at the same time -- a preview of Section 7's counterfactual problem.
- **TCO** Total cost of ownership extends far beyond the software license: model/token usage, cloud compute, data storage, integration, cybersecurity, governance, human review, training, ongoing maintenance, vendor management, compliance, and change management all add real, recurring cost.
- **!** The software or model license itself is frequently a SMALL fraction of total cost of ownership -- the surrounding organizational and operational cost is usually larger, and easy to underestimate in an initial business case (intangible benefits like better decisions, employee satisfaction, and risk reduction matter too, but are difficult to place into a DCF with precision -- treat them as a qualitative supplement to the financial case, not a substitute for it).

04

PART 3 · SECTION 4

Discounted Cash Flow Valuation and AI

Scaling from a single project decision to valuing an entire AI-driven enterprise.

SECTION 4 · THE FRAMEWORK

The DCF Framework -- and Why AI Makes It Unusually Hard

THE FRAMEWORK

— Enterprise value = present value of forecast free cash flow (forecast period) + present value of terminal value

— Equity value = enterprise value − net debt

— Requires: a forecast period, projected free cash flow each year, a discount rate, and a terminal value assumption

WHY AI ADDS UNUSUAL UNCERTAINTY

— Revenue growth, pricing, and market share are all genuinely unsettled (Module 2's commoditization debate)

— Margins, CapEx, useful life, and compute/model costs are all moving targets (Sections 1-2 of this module)

— Terminal economics -- what the business looks like once growth normalizes -- is especially unclear for a technology this early in its diffusion curve

TAKEAWAY: DCF is still the right conceptual framework -- but every input is unusually uncertain right now, which is exactly why Sections 4's sensitivity and scenario tools matter more here than in a typical valuation.

BUILDING THE FORECAST, PART 1

Revenue and Margin Forecasting

REVENUE: BREAK IT INTO COMPONENTS

— Number of customers × adoption rate (what share of potential customers actually adopt) × usage per customer × price per unit

— Plus retention and expansion revenue from existing customers -- often the most reliable component to forecast

— Adoption rate is not a finance input alone -- it is fundamentally a Diffusion-of-Innovation question (next module) about how fast a population moves from innovators through laggards

MARGIN: WILL AI RESEMBLE...

— Software? (high gross margin, low marginal cost per additional user)

— Cloud infrastructure? (moderate margin, high capital intensity, Module 2)

— Consulting? (labor-driven margin, scales with headcount, Module 2)

— Utilities? (regulated, capital-intensive, low but stable margin)

OPEN QUESTION FOR THE CLASS: different AI business models plausibly resemble different rows on the right -- and the correct answer may differ by company, and by layer of Module 2's value chain.

Capital Expenditure Forecasting: Growth vs. Maintenance

Growth CapEx: spending to add NEW capacity to serve NEW demand -- this is the spending investors most readily associate with an exciting growth story.

Maintenance CapEx: spending required just to REPLACE aging or obsolete capacity and keep existing revenue flowing -- server refresh cycles, GPU refresh cycles, network upgrades, and power/cooling maintenance.

Given the pace of chip-generation turnover and the physical-vs-economic-life gap established in Section 1, maintenance CapEx for AI infrastructure may become a much larger ongoing burden than investors currently assume once the initial growth-driven buildout matures.

A critical forecasting question rarely made explicit in AI models: once the current buildout wave slows, what fraction of a hyperscaler's total CapEx becomes pure replacement spending just to stand still -- rather than growth spending funding new revenue?

Graphic recommendation: a stacked bar chart over a 10-year forecast horizon showing the mix of growth CapEx (shrinking share over time) versus maintenance CapEx (growing share over time) as the buildout matures.

PRICING THE RISK

The Discount Rate: Higher or Lower for AI?

THE CASE FOR A HIGHER RATE

— Genuine uncertainty across revenue, margin, CapEx, and regulation simultaneously (this section) justifies a larger risk premium than a typical mature business

— Technological obsolescence risk (Section 1) adds a risk dimension most DCFs don't need to price explicitly

— Competitive intensity in several ecosystem layers (Module 2's Five Forces analysis of foundation models) argues for a higher required return

THE CASE FOR A LOWER RATE

— Dominant scale and market position in some layers (Module 2's profit-pool analysis: semiconductor equipment, foundries, GPU designers) may justify a LOWER risk premium than the uncertainty alone would suggest

— Diversified scope advantages (Module 2: Microsoft, Amazon, Alphabet, Meta spreading AI investment across many revenue lines) reduce company-specific risk relative to a single-product AI-native firm

— Large, growing backlog/RPO (Module 1/2) provides a degree of revenue visibility unusual for an early-stage technology

PRACTICAL RESOLUTION: WACC still blends a risk-free rate, equity risk premium, beta, cost of debt, and capital structure -- the AI-specific judgment call is mainly in the beta/risk-premium inputs, and reasonable analysts can and do land on different answers for the same company.

THE ASSUMPTION THAT DOMINATES THE MODEL

Terminal Value

In most DCF models -- and especially for a company still early in its growth/investment cycle -- terminal value can represent the majority of total enterprise value.

SCENARIO ANALYSIS

Three Coherent Futures, Not One Point Estimate

BULL CASE	BASE CASE	BEAR CASE
— Rapid adoption across the customer base	— Gradual, uneven adoption	— Slow adoption relative to capacity built
— Strong pricing power sustained	— Moderate pricing pressure over time	— Commoditization compresses pricing (Module 2)
— High utilization of built capacity	— Mixed organizational returns across firms	— Excess capacity built ahead of real demand
— Durable, defensible margins	— High but manageable capital expenditure	— Regulatory constraints slow deployment
— Falling compute/inference cost per unit	— Partial workflow redesign success	— Weak utilization of installed base
— Successful, completed workflow redesign	— Margins settle between software and infrastructure norms	— Rapid asset obsolescence (Section 1)

05

PART 3 · SECTION 5

Market Multiples and Investor Expectations

The market's shorthand for everything Section 4 just modeled explicitly.

Common Valuation Multiples Across the AI Ecosystem (Illustrative, as of 7/30/26)

Company (Business Model)	P/E (TTM/Fwd)	EV/Revenue	FCF Yield
Nvidia (GPU designer)	29.1x / 19.1x	18.3x	2.55%
Microsoft (hyperscaler/software)	25.4x / 23.4x	10.1x	1.98%
Palantir (AI software/analytics)	138.3x / 77.4x	55.0x	0.91%
CoreWeave (GPU neocloud)	n/m (net loss)	11.9x	-25.8%
Snowflake (data platform)	n/m / 133.8x fwd	19.5x	1.17%
ServiceNow (enterprise software + AI)	72.3x / 25.6x	7.75x	4.07%
Oracle (database/cloud infrastructure)	20.2x / 14.6x	7.05x	-6.98%
Arista Networks (AI networking)	54.3x / 41.7x	20.9x	2.46%

Snapshot as of July 30, 2026 (~1:28 PM ET); multiples move daily with stock price. Source: stockanalysis.com company overview/statistics pages, underlying fundamentals from SEC filings/earnings releases.

READING REVENUE MULTIPLES CORRECTLY

Why High-Growth AI Firms Use Revenue Multiples and Comparable-Company Discipline

WHY REVENUE MULTIPLES DOMINATE HERE

— Negative or unstable earnings make P/E meaningless for many high-growth AI firms (CoreWeave, Snowflake, per the previous slide)

— Heavy reinvestment (Sections 1-4) depresses near-term profit even when the underlying business is healthy

— Early lifecycle stage means limited comparable operating history to anchor an earnings-based multiple

— RISK: a revenue multiple can conceal genuinely poor unit economics -- fast revenue growth funded by deeply negative FCF (Oracle, CoreWeave) looks identical on an EV/Revenue basis to fast, profitable growth

BUILDING A REAL COMPARABLE SET

— Match on business model, customer type, growth rate, margin profile, capital intensity, revenue durability, risk, geography, and competitive position -- not just a shared 'AI' label

— Module 2's ecosystem layers are the right starting filter: a foundation-model lab, a GPU designer, and an enterprise-software incumbent are NOT comparable companies just because all three are 'AI companies'

— Labeling every company 'AI' does not make it comparable -- this is one of the most common analytical errors in the current market

DECODING THE NUMBER

What Is Embedded in a High Multiple?

1. High expected future growth -- the multiple is effectively a bet on a growth trajectory, not a reflection of current-year fundamentals alone.
2. High expected future margins -- often assuming a business scales toward software-like profitability even if it doesn't operate that way today.
3. A long assumed competitive-advantage period -- the market is pricing in durable differentiation for many years, not just the next four quarters.
4. Low perceived risk and low expected future capital intensity -- an implicit bet that today's heavy reinvestment phase (Sections 1-2) eventually fades.
5. Strong customer retention and low expected competitive erosion -- an implicit bet against the commoditization dynamics Module 2 flagged as a live, open risk for several ecosystem layers.

Graphic recommendation: an iceberg diagram -- the visible multiple above the waterline, with these five assumptions listed as the much larger hidden mass below it.

06

PART 3 · SECTION 6

Reading Earnings Calls, 10-Ks, and 10-Qs

Where the numbers in Sections 1-5 actually come from -- and how to read them critically.

THE SINGLE MOST IMPORTANT IDEA IN THIS SECTION

Earnings vs. Expectations: Same Week, Opposite Reactions

ALPHABET, JULY 22-23, 2026 -- "BEAT, BUT STOCK FELL"

— Google Cloud revenue +82% YoY to $24.8B, beating ~64% consensus growth (LSEG data); total revenue also beat estimates

— But: adjusted EPS of $2.85 missed the $2.89 consensus, 2026 capex guidance raised to $195-205B (above the ~$188B Visible Alpha consensus), and Alphabet posted its first-ever negative quarterly FCF (-$5.9B)

— Shares fell over 3.5% (some reports cite ~7%) despite the strong headline growth

MICROSOFT, JULY 29-30, 2026 -- "BEAT, AND STOCK JUMPED"

— Revenue $90B (vs. $87.42B est.), EPS $4.81 (vs. $4.20 est.), Azure growth 43% (vs. ~40% est.), Azure crossed $100B in annual revenue for the first time

— Critically: Microsoft did NOT raise CapEx guidance -- FY2027 Q1 capex guided to $50B (below the $56B Street estimate), CY2026 capex guided DOWN to $175B (from a prior ~$190B estimate)

— FCF of $19.6B beat the ~$13.4B estimate; shares rose ~3.5% after-hours and were up 16.7% intraday the next day

TEACHING TAKEAWAY: both companies grew cloud revenue robustly and beat headline estimates in the SAME WEEK -- the market moved them in opposite directions by double digits based entirely on forward CapEx/FCF guidance, not the reported quarter itself.

Why Earnings Calls Matter -- and Their Anatomy

WHY THEY MATTER

— Management's own narrative and emphasis reveal priorities the raw numbers alone don't show

— Analyst questions probe exactly the areas of greatest market uncertainty

— Guidance sets the expectations baseline Section 5 showed the market actually trades against

— Tone and terminology shifts over time can signal changing management confidence before the numbers fully reflect it (next slide)

ANATOMY OF AN EARNINGS CALL

— 1. Safe-harbor statement (forward-looking-statements disclaimer under the Private Securities Litigation Reform Act of 1995, Section 21E of the Securities Exchange Act of 1934, 15 U.S.C. §78u-5)

— 2. CEO remarks (strategy, narrative)

— 3. CFO remarks (financial detail, guidance)

— 4. Segment performance

— 5. Guidance for future periods

— 6. Analyst Q&A (often the most informative section)

— 7. Closing comments

THE LISTENING CHECKLIST

What to Listen for in an AI Earnings Call

CapEx Guidance
Is it rising, falling, or steady versus the prior quarter's guidance? (Section 5's Alphabet/Microsoft example)

Capacity & Utilization
Constrained, balanced, or excess? Watch for softening language over time (next slide)

Backlog / RPO
Growing, shrinking, or changing definition? (Module 1/2, and this section's next slides)

AI Revenue / Bookings
Is the metric newly disclosed, newly discontinued, or redefined? (Module 2 flagged several companies doing exactly this)

Model / Inference Costs
Rising or falling per-unit cost -- directly affects margin forecasting (Section 4)

Gross Margin & Depreciation
Any useful-life or capitalization-policy change disclosed? (Section 1)

Customer Concentration
Reliance on a small number of large AI customers or partners

Power Availability & Timelines
Data-center construction/power delays (Module 2, Layer 1-2)

Free Cash Flow & Debt
Direction of FCF, and any new debt issuance to fund the gap (Section 2)

INTERPRETIVE SKILL

Reading Between the Lines -- and Management Credibility

LANGUAGE PATTERNS WORTH NOTICING

— General patterns to listen for over time (illustrative language patterns, not attributed to any specific real company's actual transcript language, unless independently verified): softer, more hedged phrasing replacing previously confident, definitive statements about demand or capacity

— New non-GAAP measures appearing, or previously disclosed metrics quietly disappearing (Module 2's Accenture/IBM example, previous slide, is a real, sourced instance of this exact pattern)

— Changes in how 'AI revenue' itself is defined from one quarter to the next -- a red flag for comparability even absent any wrongdoing

MANAGEMENT CREDIBILITY CHECKLIST

— Track record of prior guidance accuracy -- has guidance been reliably hit, consistently missed, or consistently beaten (sandbagged)?

— Consistency of message across consecutive quarters

— Capital-allocation track record -- disciplined, or chasing every new narrative?

— Disclosure transparency and stability of non-GAAP measures used

— Management incentive structure and compensation design, and any notable insider-transaction activity

07

PART 3 · SECTION 7

Measuring AI Return on Investment

The formula is simple. Measuring its inputs honestly is the hard part.

SECTION 7 · WHY THIS IS GENUINELY HARD

Why AI ROI Is Difficult to Measure -- and a Working Formula

WHY IT'S HARD

— AI investments often serve multiple objectives and share infrastructure across many use cases simultaneously, making cost allocation to any single initiative genuinely ambiguous

— Productivity gains may show up as time reallocated to other tasks, not as a directly countable output

— Quality improvement, employee learning effects, and workflow-redesign effects (Section 2's J-curve) are all real but hard to isolate from an AI tool's direct effect

— Employee adoption varies, and any productivity gain can be partly offset by cannibalizing existing revenue rather than creating new value

A WORKING FORMULA

— (Incremental benefits − implementation and operating costs) ÷ total investment

— This simplified formula is a useful starting point, but it can miss TIMING (when benefits actually arrive, per the J-curve) and RISK (the range of possible outcomes, per Section 4's scenario analysis)

— A more complete answer requires NPV/IRR treatment (Section 3) applied specifically to the AI initiative in question, not just a single-period ratio

A FULLER MEASUREMENT SYSTEM

AI KPI Categories: Financial, Operational, Customer, Risk, Adoption

Financial
Revenue growth, cost savings, gross margin, free cash flow, ROIC, payback, NPV (Sections 2-3)

Operational
Cycle time, throughput, error rate, automation rate, resolution time, utilization

Customer
Conversion, retention, satisfaction, response time, personalization

Risk & Control
Fraud detection, compliance exceptions, audit findings, security incidents, hallucination rate, human-review rate

Adoption
Active users, frequency, workflow penetration, repeat usage, department coverage, training completion

A FREQUENTLY BLURRED DISTINCTION

Productivity Gains vs. Headcount Reduction

1 Productivity gains can show up as MORE output at the same headcount, BETTER quality at the same cost, FASTER delivery, or HIGHER capacity to take on new work -- none of which necessarily reduce headcount at all.

2 Productivity gains can also show up as reduced HIRING NEEDS going forward, REDEPLOYMENT of existing staff to higher-value work, or, in some cases, direct headcount reduction.

3 These outcomes are NOT economically equivalent, and conflating them produces bad forecasts: a company that redeploys staff to revenue-generating work captures very different economic value than one that simply cuts headcount and captures pure cost savings -- even if both started from an identical productivity gain.

! NBER Working Paper 31161 (previous slide) is a useful real example precisely because it measured productivity gain WITHOUT assuming headcount reduction -- the study measured issues resolved per hour and customer sentiment, not job losses.

08

PART 3 · SECTION 8

Risks, Red Flags, and Executive Judgment

Careful, neutral language for a genuinely important skill: knowing when to ask harder questions.

Financial Red Flags Worth Investigating Further

Category	Signals Worth Investigating
Capacity & Demand	CapEx growing faster than demand signals; falling utilization; weak backlog quality (cancellable or loosely defined commitments)
Balance Sheet & Leverage	Rising debt without a clearly disclosed repayment plan; negative free cash flow without a credible, evidenced path to positive FCF (Section 2)
Disclosure Quality	Repeated guidance changes; expanding or shifting non-GAAP adjustments; metrics that appear, then quietly disappear (Section 6)
Accounting Judgment	Aggressive useful-life assumptions relative to observed technology turnover (Section 1); rapidly increasing capitalized costs
Related-Party & Structural	Related-party transactions; vendor financing arrangements; circular commercial relationships (e.g., a chip supplier financing its own customer's purchases) -- worth scrutinizing structurally, without presuming improper intent

IMPORTANT: these are neutral, structural signals worth investigating further -- not allegations of wrongdoing. Any specific real-company concern should be evaluated against authoritative evidence (filings, disclosures) rather than inference alone.

BEYOND THE NUMBERS

Strategic Red Flags

1. No clear business use case -- the project exists because AI is fashionable, not because it solves a defined problem.
2. Technology-first implementation with no accompanying workflow redesign (Section 2's J-curve stalling at stage 3-4).
3. Weak underlying data quality, no clear executive owner, and low employee adoption despite the tool being available.
4. Heavy vendor lock-in with no governance framework and no measurement system in place to know whether it's working.
5. No plan for scaling a successful pilot -- and, just as importantly, no plan for discontinuing a project that has clearly failed.

END OF PART 3 · THE PIVOT POINT OF THIS COURSE

Every Valuation Contains an Adoption Forecast.

Revenue growth, margins, productivity, and terminal value all depend on how rapidly AI moves from experimentation to sustained use.

PART 4

The Diffusion of AI Innovation

Why Adoption Spreads Unevenly Across People, Organizations, Industries, and Markets

01

PART 4 · SECTION 1

Why Diffusion of Innovation Matters for AI

Every valuation model built in Module 3 contains a hidden forecast about how fast AI will be adopted

THE BRIDGE

From Valuation to Adoption: Closing the Loop from Module 3

WHAT MODULE 3 ESTABLISHED

— Every DCF requires a revenue-growth and margin forecast

— Every multiple embeds a market view of future growth

— Every capital-budgeting decision (NPV/IRR/real options) requires an assumption about when benefits materialize

— Module 3 closed with: 'Every DCF assumes an adoption curve'

ROGERS' DEFINITION OF DIFFUSION

— "Diffusion is the process by which an innovation is communicated through certain channels over time among members of a social system" (Rogers, 2003, p. 5)

— Diffusion is a special type of communication in which the messages are about a new idea

— Diffusion can be planned or spontaneous

This module makes the hidden adoption-curve assumption explicit — and gives finance and strategy professionals a rigorous framework for interrogating it.

ROGERS' FRAMEWORK

The Four Main Elements of Diffusion

1. Innovation
An idea, practice, or object perceived as new by an individual or unit of adoption. AI: generative AI, agentic systems, ML-based forecasting.

2. Communication Channels
The means by which messages about the innovation move from one individual to another. AI: vendor demos, LinkedIn, internal town halls, peer conversations.

3. Time
The innovation-decision process, relative time of adoption, and rate of adoption all involve a time dimension. AI: months from pilot to scaled deployment.

4. Social System
A set of interrelated units engaged in joint problem-solving toward a common goal. AI: a firm, an industry, a profession, or society.

FOUNDATIONS

Invention vs. Innovation — and Why Newness Creates Uncertainty

INVENTION vs. INNOVATION

— Invention: the original creation of a new idea, process, or technology (e.g., the transformer architecture, Vaswani et al., 2017)

— Innovation (Rogers): "An idea, practice, or object that is perceived as new by an individual or other unit of adoption" — newness is perceptual, not chronological

— A technology invented years earlier can still function as a Rogers 'innovation' the first time a given adopter encounters it

WHY NEWNESS CREATES UNCERTAINTY

— Diffusion is a special type of communication in which the messages concern a new idea (Rogers, 2003, p. 5)

— Because there is newness, there is uncertainty about the innovation's expected consequences (Rogers, 2003, p. 5)

— Uncertainty triggers information-seeking and information-processing activity — the entire innovation-decision process (Section 2) exists to reduce it

For most enterprise adopters, generative AI is not new technology — it is a new perceived option, which is exactly what makes it a Rogers 'innovation' rather than merely a product launch.

EARLY STUDIES

Historical Perspective: The Empirical Roots of Diffusion Theory

Gabriel Tarde (1903)

French magistrate and early sociologist; developed the 'laws of imitation,' a precursor to diffusion theory. Asked: 'why, given one hundred different innovations conceived at the same time... ten will spread abroad while ninety will be forgotten' (Tarde, 1903, cited in Rogers, 2003, p. 40). Also documented that adoption over time follows an S-shaped curve (Rogers, 2003, p. 41).

Georg Simmel (1908)

Introduced the concept of 'the stranger' — someone 'who is a member of a system but not strongly attached to it' (Simmel, 1908/1964, cited in Rogers, 2003, p. 45). This work seeded the later concepts of cosmopoliteness and heterophily central to understanding innovators who adopt outside prevailing norms (Rogers, 2003, pp. 45-47).

Ryan & Gross (1943)

The hybrid-corn adoption study among Iowa farmers — the first empirical diffusion paradigm. Farmers first learned of hybrid corn from salesmen, but peer-to-peer exchange among neighboring farmers was the most persuasive factor in the final adoption decision (Rogers, 2003, pp. 49-51).

02

PART 4 · SECTION 2

The Innovation-Decision Process

Five stages every individual and organization moves through — Knowledge, Persuasion, Decision, Implementation, Confirmation

ROGERS' ORIGINAL DIAGRAM

The Full Model: Prior Conditions, Channels, Attributes, and Outcomes

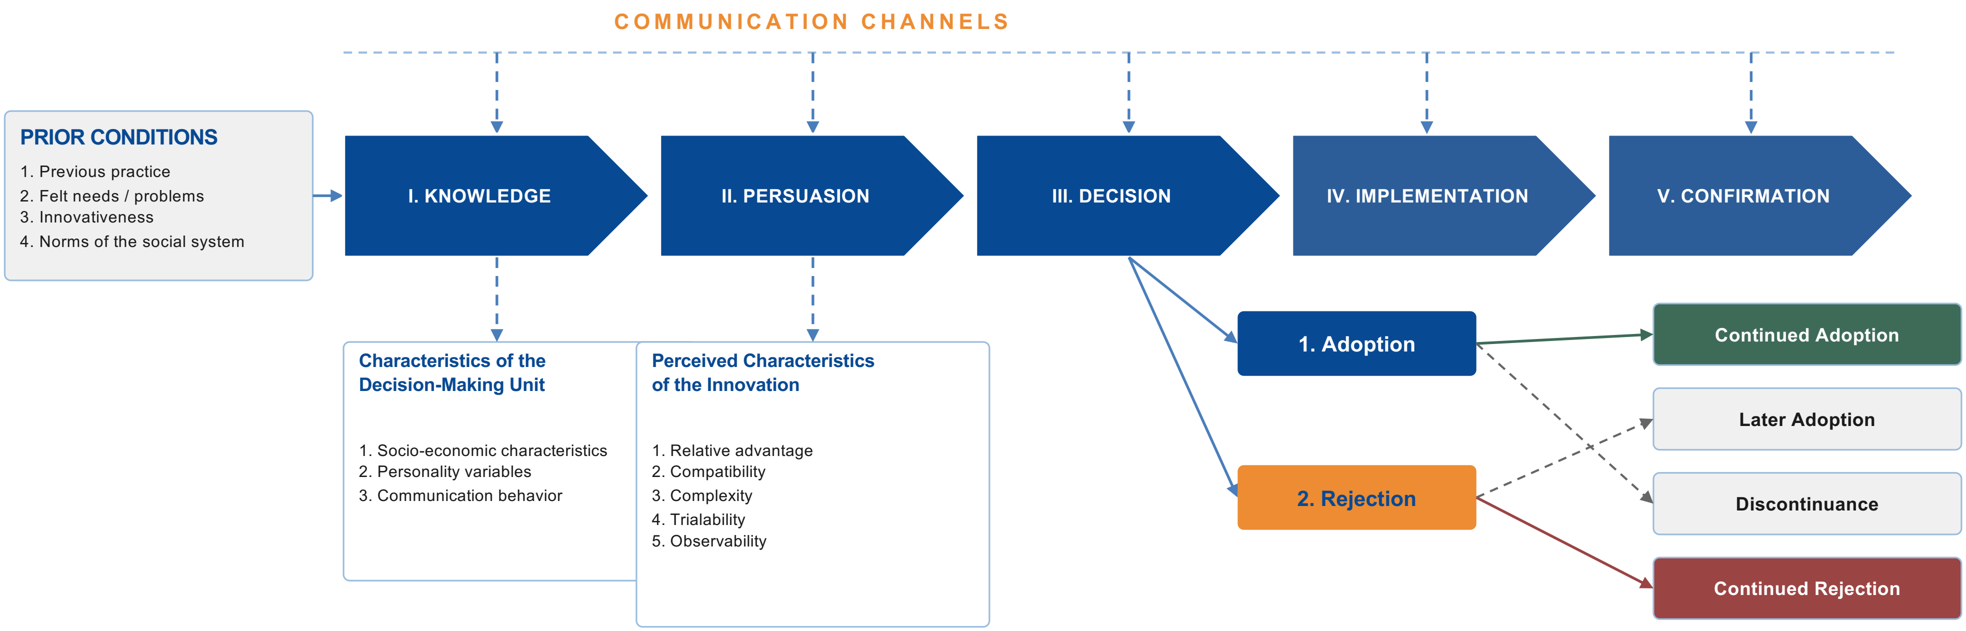

Figure 5-1. A model of stages in the innovation-decision process.

Rogers, Everett M. 2003. Diffusion of Innovations, 5th ed. New York: Free Press. (Figure 5-1).

?

AUDIENCE POLL

Where would you place your organization's current AI adoption position?

A
Knowledge stage — aware of AI options, still building understanding of how they work

B
Persuasion/Decision — actively evaluating and forming attitudes, early pilots underway

C
Implementation — AI tools are in active use in at least one workflow or function

D
Confirmation — AI use is established and being reinforced, expanded, or reconsidered

OVERVIEW

What Precedes Knowledge: The Four Prior Conditions

Previous Practice	Felt Needs / Problems	Innovativeness	Norms of the Social System
The customary, already-established way this task or decision was handled before the innovation appeared (p. 172)	A gap the potential adopter already perceives between the current state and a desired one (p. 172)	How early this individual or organization tends to adopt new ideas, relative to others in its social system (p. 172)	Established behavior patterns and expectations that can either encourage or resist a proposed change (p. 172)

OVERVIEW

Who Is Deciding: Characteristics of the Decision-Making Unit

Socioeconomic Characteristics	Personality Variables	Communication Behavior
Education, income, organizational role, and the resources available to act on a decision (p. 173)	Risk tolerance, dogmatism, and the capacity to handle abstraction and ambiguity (p. 173)	Social participation, media exposure, and connectedness to opinion leaders in the network (p. 173)

OVERVIEW

The Five-Stage Innovation-Decision Process

1. Knowledge

Individual is exposed to the innovation's existence and gains understanding of how it functions (Rogers, 2003, p. 172).

2. Persuasion

Individual forms a favorable or unfavorable attitude toward the innovation (Rogers, 2003, p. 176).

3. Decision

Individual engages in activities leading to a choice to adopt or reject (Rogers, 2003, p. 179).

4. Implementation

Individual puts the innovation into use; reinvention often occurs here (Rogers, 2003, pp. 181-183).

5. Confirmation

Individual seeks reinforcement for the decision already made, or may reverse it (Rogers, 2003, p. 189).

Reference: Rogers, Everett M. 2003. Diffusion of Innovations, 5th ed. New York: Free Press. This is a process, not a single event — organizations can stall at any stage.

THE INNOVATION-DECISION PROCESS

Stage 1 — Knowledge: Three Types, Uneven Access

1. The process begins when an individual 'is exposed to the innovation's existence and gains some understanding of how it functions' (Rogers, 2003, p. 172).

2. Knowledge is not one thing: awareness-knowledge (knowing that an innovation exists), how-to knowledge (knowing how to use it correctly), and principles-knowledge (understanding the underlying functioning principles) (Rogers, 2003, pp. 173-174).

3. AI application: nearly universal awareness-knowledge of ChatGPT/Copilot by 2026, but how-to knowledge (effective prompting, workflow integration) and principles-knowledge (how a model actually generates output, where it can fail) remain unevenly distributed.

4. Knowledge Gaps: Rogers found individuals differ in how quickly they encounter innovations based on socioeconomic status, personality, and social-network position (Rogers, 2003, p. 175). AI evidence: knowledge gaps by occupation (legal/medical/technical staff ahead of administrative staff), by age, and by industry are consistently documented in 2025-2026 adoption surveys.

5. Discussion question (from the instructor's original material): has anyone in the room not heard of AI at this point? The near-universal 'no' response demonstrates how saturated awareness-knowledge already is — the real gap is in how-to and principles knowledge.

Stage 2 — Persuasion: Attitude Formation, FOMO, and Adoption Types

1. In the persuasion stage, 'the individual forms a favorable or unfavorable attitude toward the innovation' (Rogers, 2003, p. 176). This is affective, not purely cognitive — emotion, peer influence, and social norms all play key roles (Rogers, 2003, p. 177).

2. The five perceived attributes of the innovation (Section 3) strongly influence persuasion outcomes and, taken together, explain 49-87% of the variance in adoption rate (Rogers, 2003, pp. 221-222).

3. Discussion question (original material): is there a FOMO (fear of missing out) component to AI adoption right now? Competitive anxiety — 'our competitors are already using this' — is a documented persuasion-stage driver distinct from a rational cost-benefit case.

4. Adoption is frequently symbolic rather than strategic in this stage: organizations may adopt visible AI tools (a public chatbot, an 'AI strategy' slide) primarily to signal legitimacy to investors, regulators, or customers — before any operational case has been built (see also Decision-stage symbolic adoption, next slide).

5. Distinguish adoption motives that surface at this stage: defensive (avoiding falling behind competitors), strategic (a deliberate capability build), symbolic (signaling legitimacy), and experimental (structured, bounded learning) — each implies a different governance and measurement approach.

THE INNOVATION-DECISION PROCESS

Stage 3 — Decision: Choosing to Adopt or Reject

1. The decision stage occurs when an individual 'engages in activities that lead to a choice to adopt or reject the innovation' (Rogers, 2003, p. 179).
2. Adoption is not purely rational — it is shaped by social confirmation and perceived risk, not solely by expected economic return.
3. Symbolic adoption often crystallizes at this stage: a formal decision to announce or pilot AI can precede any operational decision about how the tool will actually be used — the decision to be seen adopting and the decision to operationally deploy are not the same decision.
4. Rejection is a legitimate outcome of this stage, not a failure state — Rogers treats active rejection (after genuine consideration) as distinct from simple non-adoption due to lack of knowledge.
5. Discussion question (original material): is there a FOMO component driving premature or purely symbolic 'yes' decisions before the underlying use case has been validated?

Stage 4 — Implementation and the Pilot-to-Scale Gap

1. Implementation occurs when 'an individual puts an innovation into use' (Rogers, 2003, p. 181). This phase often reveals unanticipated problems, leading to reinvention — modifying the innovation to fit local needs (Rogers, 2003, pp. 182-183).

2. Organizational settings require coordination, resource allocation, and formal support to sustain implementation (Rogers, 2003, p. 184) — this is precisely where most enterprise AI programs currently stall.

3. The Pilot-to-Scale Gap: Project NANDA's 'The GenAI Divide: State of AI in Business 2025' (MIT Media Lab, reported Aug. 2025; 150 leader interviews, 350-employee survey, 300 public deployments) found approximately 95% of enterprise generative-AI pilots fail to reach measurable P&L impact, while purchased/vendor tools succeed roughly 67% of the time vs. about one-third for internal builds.

4. Critically, individual-level consumer tool use (e.g., ChatGPT) frequently succeeds where the same organization's formal enterprise workflow deployment stalls — do not conflate the two when assessing 'AI adoption' at a company (Project NANDA, 2025).

5. Discussion question (original material): is coordination, resource allocation, and formal support an issue with AI adoption right now in your organization?

A NOTE ON IMPLEMENTATION

Reinvention: The Innovation Rarely Survives Contact Unchanged

1. "Reinvention is the degree to which an innovation is changed or modified by a user in the process of its adoption and implementation" (Rogers, 2003).

2. Reinvention is not a failure of fidelity — Rogers treats it as a normal and often adaptive part of implementation, allowing an innovation to fit local conditions.

3. AI application: large language models are, in a sense, built for reinvention — the same base model is 'reinvented' into wildly different configurations (custom instructions, retrieval pipelines, fine-tuning, agentic workflows) by each adopting organization.

4. This is potentially what we are seeing with LLMs at scale: rather than one canonical 'correct' way to deploy generative AI, thousands of organizations are independently reinventing the same underlying technology for their own workflows (echoing the instructor's original observation).

5. Strategic implication: vendors who under-anticipate reinvention (rigid, non-configurable tools) diffuse more slowly into heterogeneous enterprise environments than vendors who design explicitly for local adaptation.

THE INNOVATION-DECISION PROCESS

Stage 5 — Confirmation, and Its Opposite: Discontinuance

STAGE 5 — CONFIRMATION

— Individuals 'seek reinforcement for the innovation-decision already made, or they may reverse this decision if exposed to conflicting messages' (Rogers, 2003, p. 189)

— Continued adoption depends on satisfaction, peer validation, and consistent results

— Contradictory feedback can produce discontinuance (Rogers, 2003, p. 190)

— Discussion question (original material): are we getting ROI on AI adoption yet?

DISCONTINUANCE — TWO TYPES

— Replacement discontinuance: the innovation is rejected in order to adopt a better idea that supersedes it — a positive, forward-looking form of abandonment

— Disenchantment discontinuance: the innovation is rejected because the adopter is dissatisfied with its performance

— AI application: an employee dropping one AI tool for a better one (replacement) is a very different signal than an employee abandoning AI use entirely after a bad experience (disenchantment)

INNOVATION-DECISION PROCESS RECAP

Process Summary: The Five Stages at a Glance

Stage	Key Question	Common AI Barrier	Suggested Metric
1. Knowledge	Do people know this exists and how it works?	Awareness saturated; how-to and principles knowledge lag	How-to / principles training completion rate
2. Persuasion	Do people believe this is worth adopting?	FOMO-driven symbolic attitudes without validated use case	Attitude/sentiment survey; pilot opt-in rate
3. Decision	Has a real choice to adopt been made?	Symbolic 'yes' decisions with no operational follow-through	Approved budget converted to active projects
4. Implementation	Is the tool actually in use in workflows?	Pilot-to-scale gap; lack of coordination/support	Active weekly users; workflow integration rate
5. Confirmation	Is continued use being reinforced?	Disenchantment discontinuance; unmeasured ROI	Retention rate; realized ROI vs. Module 3 case

Reference: Rogers, Everett M. 2003. Diffusion of Innovations, 5th ed. New York: Free Press., pp. 168-218 (Innovation-Decision Process chapter).

03

PART 4 · SECTION 3

The Five Perceived Attributes of AI

Together these attributes explain 49-87% of the variance in adoption rate (Rogers, 2003, p. 221)

OVERVIEW

Innovation Characteristics: The Five Perceived Attributes

Relative Advantage	Compatibility	Complexity	Trialability	Observability
Better than the idea it supersedes (p. 229)	Consistent with values, past experience, needs (p. 240)	Difficult to understand and use (p. 257)	Can be experimented with on a limited basis (p. 258)	Results are visible to others (p. 258)

ATTRIBUTE I

Relative Advantage: The Strongest Predictor of Adoption — and a Value Proposition

1. "The degree to which an innovation is perceived as being better than the idea it supersedes" (Rogers, 2003, p. 229) — measurable in economic gain, prestige, convenience, or satisfaction, and the strongest single predictor of adoption rate (Rogers, 2003, p. 229).

2. From a firm's perspective, relative advantage is a value proposition: it incorporates strategic positioning, customer utility, and elements of value — directly connected to 'Jobs-to-Be-Done' theory (per the instructor's original framing).

3. Jobs-to-Be-Done framing for AI: the practical question is never 'is AI impressive?' but 'better than what, at what job?' — an AI drafting tool must be compared against the specific task (a junior associate's first-draft memo, a paralegal's document review) it is displacing, not against a generic baseline.

4. AI-specific complication: relative advantage is often asymmetric — clearly positive for repetitive, high-volume cognitive tasks (document summarization, code scaffolding) and much less clear for judgment-intensive, high-stakes tasks (final audit sign-off, clinical diagnosis) where error cost is high.

ATTRIBUTE I — APPLIED

Relative Advantage and ROI: A Task-Level Scorecard

Document drafting / summarization	**HIGH**	Clear time savings on first drafts; human review still required before finalization.
Code generation / scaffolding	**HIGH**	Strong relative advantage on boilerplate; lower on architecture-level decisions.
Data analysis / forecasting support	**MEDIUM**	Useful for exploration; relative advantage depends heavily on data quality and domain fit.
Customer-facing judgment calls	**LOW-MED**	Relative advantage constrained by trust, liability, and error-cost concerns.
Final professional sign-off (audit, legal, clinical)	**LOW**	Regulatory and liability structures limit relative advantage regardless of technical capability.

Relative advantage is not a property of the technology alone — it is a property of the technology matched to a specific, bounded task.

ATTRIBUTE II

Compatibility: Fitting Existing Values, Workflows, and Professional Identity

1. "The degree to which an innovation is perceived as consistent with the existing values, past experiences, and needs of potential adopters" (Rogers, 2003, p. 240). Incompatible ideas meet resistance regardless of technical merit.

2. Historical example: the water-boiling campaign in Peru failed because it clashed with cultural beliefs about health and cooking practices (Rogers, 2003, p. 242) — technically sound, culturally incompatible.

3. Compatibility across industries is uneven: software/tech and marketing functions report high compatibility with AI-assisted workflows; heavily regulated, judgment-and-liability-driven professions (audit, law, medicine, education) report lower compatibility with existing values around accountability and sign-off.

4. Professional identity and compatibility: auditors, physicians, attorneys, professors, and analysts derive professional identity partly from exercising independent judgment — an AI tool perceived as substituting for that judgment (rather than supporting it) is perceived as incompatible with professional identity itself, triggering resistance independent of the tool's actual accuracy.

ATTRIBUTE III

Complexity — and the AI-Specific 'Complexity Paradox'

COMPLEXITY (ROGERS)

— "The degree to which an innovation is perceived as relatively difficult to understand and use" (Rogers, 2003, p. 257)

— 'New ideas that are simpler to understand are adopted more rapidly than innovations that require the adopter to develop new skills and understandings' (Rogers, 2003, p. 257)

— Complexity slows diffusion through cognitive and behavioral barriers

THE COMPLEXITY PARADOX

— Personal-use AI (a chat interface) has extremely low perceived complexity — near-zero learning curve for basic use

— Enterprise-scale AI deployment (governance, data integration, workflow redesign, risk controls) has very high actual complexity

— This asymmetry — trivially easy to try personally, genuinely hard to deploy at scale — is a defining feature of AI diffusion not present in most prior Rogers-era case studies

Low personal-use complexity accelerates individual adoption (and Shadow AI, next section); high enterprise complexity is a primary driver of the Pilot-to-Scale Gap.

ATTRIBUTE IV — TRIALABILITY, AND SHADOW AI

78%

of AI users bring their own AI tools to work ("Bring Your Own AI"), rising to 80% at small and mid-sized companies

Source: Microsoft/LinkedIn, "2024 Work Trend Index Annual Report — AI at Work Is Here. Now Comes the Hard Part," May 8, 2024 (Edelman Data & Intelligence survey, 31,000 workers, 31 countries). Trialability — "the degree to which an innovation may be experimented with on a limited basis" (Rogers, 2003, p. 258) — is especially important for high-cost or high-risk innovations (Rogers, 2003, p. 259); AI tools like ChatGPT and Claude are inexpensive and trivially trialable, which is precisely why unsanctioned 'Shadow AI' use has become so widespread ahead of any formal governance.

ATTRIBUTE V

Observability and 'The Observability Problem' for AI ROI

1. "The degree to which the results of an innovation are visible to others" (Rogers, 2003, p. 258). 'The easier it is for individuals to see the results of an innovation, the more likely they are to adopt' (Rogers, 2003, p. 258).

2. Visibility accelerates diffusion by enabling social learning, imitation, and peer influence — directly related to ROI for AI, per the instructor's original framing.

3. The Observability Problem: many of AI's highest-value applications (better judgment, faster analysis, avoided errors) produce outcomes that are inherently hard to observe directly — a saved hour of drafting time or an avoided audit misstatement rarely shows up as a visible, attributable event the way a completed physical task does.

4. This creates a diffusion paradox: the AI use cases with the clearest relative advantage (back-office efficiency, per MIT NANDA's 2025 finding) are often the least observable, while the most visible AI use cases (public-facing chatbots, flashy demos) are not always where the greatest relative advantage lives.

APPLIED EXERCISE

Attribute Comparison: Scoring AI Against the Five Attributes

Attribute	Score	Description
Relative Advantage	HIGH	Strong and well-documented for bounded, high-volume cognitive tasks; weaker for judgment-intensive sign-off tasks.
Compatibility	MIXED	High in tech/marketing functions; low where professional identity is tied to independent judgment (audit, law, medicine).
Complexity	MIXED	Very low at the individual-use level; very high at the governed, enterprise-deployment level (the Complexity Paradox).
Trialability	HIGH	Near-zero cost of individual experimentation; drives both fast adoption and unsanctioned Shadow AI use.
Observability	LOW-MED	Highest-value use cases (back-office efficiency) are often the least visible; visible use cases are not always highest-value.

AI's attribute profile explains its own diffusion pattern: fast individual adoption (advantage, trialability strong) alongside slow, uneven enterprise scaling (compatibility, complexity, observability weaker or mixed).

04

PART 4 · SECTION 4

Adopter Categories and the AI Adoption Curve

Not everyone adopts at the same rate — and an organization is rarely just one category

SECTION 4

Rogers' Adopter Categories: The Standard Distribution

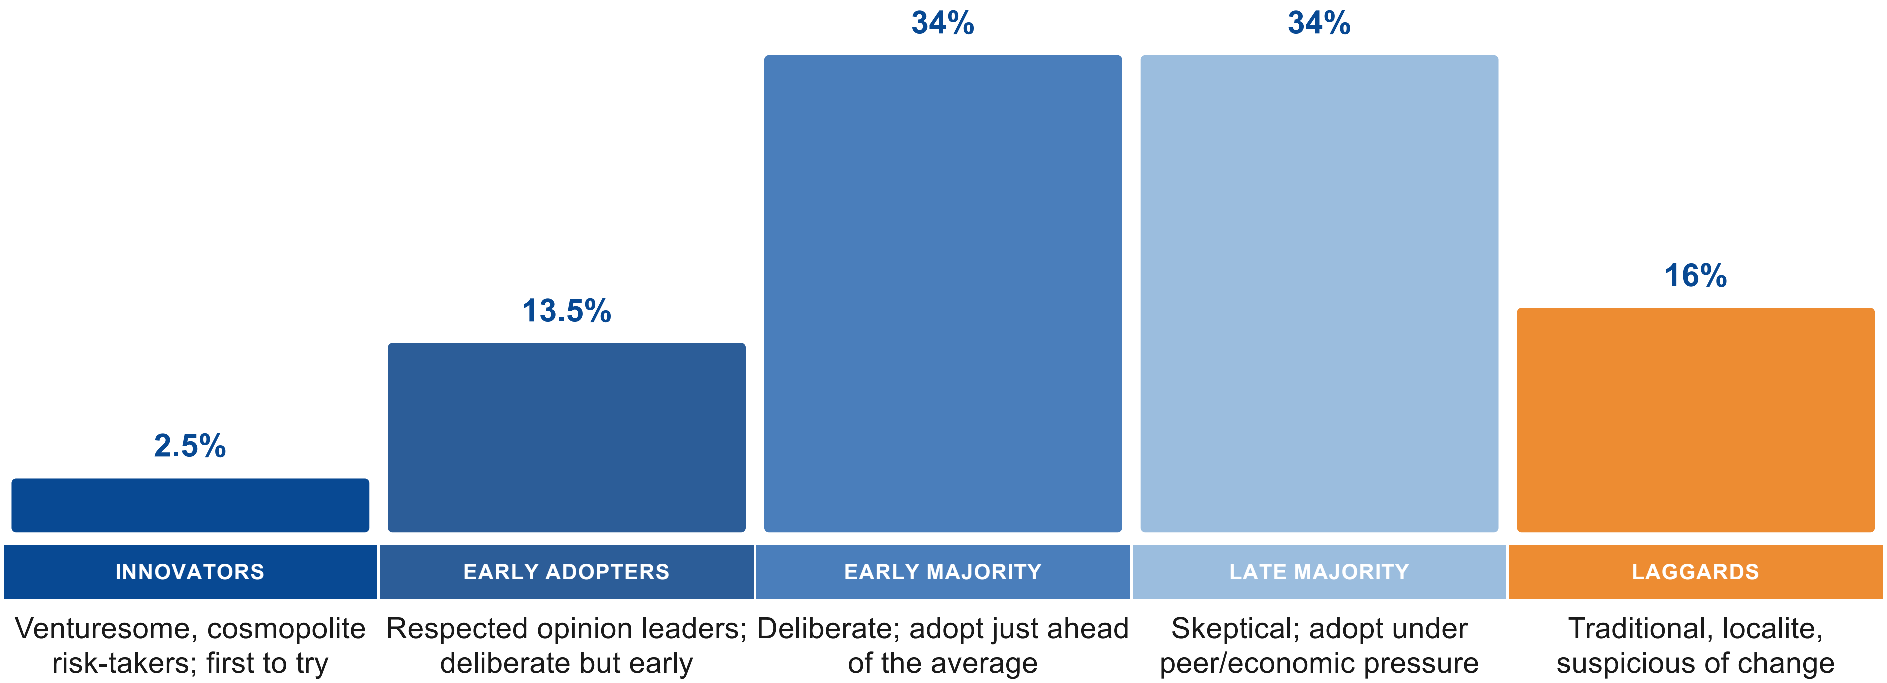

Categories and standard percentages (2.5% / 13.5% / 34% / 34% / 16%) are the widely-cited Rogers classification derived from a normal distribution of adoption time (Rogers, 2003, chapter on adopter categories, pp. 267-299).

SECTION 4

Adopter Categories: AI-Specific Characteristics and Organizational Examples

Innovators (2.5%)	Early Adopters (13.5%)	Early Majority (34%)	Late Majority (34%)	Laggards (16%)
Build custom AI agents/pipelines unprompted; comfortable with technical risk and failure. Example: an engineer independently wiring an internal LLM API into a workflow before any formal policy exists.	Respected internally; often the first formally sanctioned pilot users and de facto opinion leaders for the rest of the organization.	Adopt once a credible peer case exists and tools are packaged/supported; deliberate, not risk-averse but not first-movers.	Adopt largely due to competitive or economic pressure and peer-group norms, often after vendor tools mature and cost falls.	Localite orientation, higher skepticism of external change agents; may never adopt absent mandate (an Authority decision, Section 5).

SECTION 4

Individual vs. Organizational Adopter Categories

INDIVIDUAL CATEGORIZATION

— Rogers' original framework classifies individual persons by relative innovativeness within a social system

— An individual's category can vary by innovation — someone can be an Innovator for one technology and a Laggard for another

— AI example: a tax partner might be an Early Adopter of AI research tools but a Laggard on AI-assisted client communication

ORGANIZATIONAL HETEROGENEITY

— A single organization is virtually always heterogeneous — containing individuals from every adopter category simultaneously

— "The organization's AI adoption rate" is really an aggregate of many individual trajectories, not one trajectory

— Strategic implication: a firm can be an industry 'Early Adopter' in its public positioning while most of its workforce remains Early/Late Majority in actual daily use — a gap directly relevant to earnings-call claims examined in Module 3

SECTION 4 — SYNTHESIS

Why Adopter Categories Matter for AI Strategy and Change-Agent Targeting

1. Rogers identifies Early Adopters — not Innovators — as the most strategically important category: they are respected, socially integrated opinion leaders whose example the Early Majority watches closely.

2. Practical implication for AI rollout sequencing: target visible, credible Early Adopters for the first sanctioned pilots, rather than only the most technically enthusiastic Innovators, who may be viewed by peers as atypical.

3. The Late Majority and Laggards frequently require different decision structures entirely (see Section 5) — persuasion-based strategies that work on the Early Majority often fail on these groups, who may need an Authority decision (mandate) to move.

4. This sets up Section 8's Change Agent framework: Rogers' change-agent strategies explicitly include 'early adopter targeting' as a named acceleration tactic (Rogers, 2003, pp. 379-382).

05

PART 4 · SECTION 5

Types of AI Adoption Decisions

Who decides — and how — shapes both the speed of adoption and its legitimacy

DECISION TYPES I & II

Optional and Collective Decisions

OPTIONAL DECISIONS

— "Choices to adopt or reject an innovation made by an individual independent of others in the system" (Rogers, 2003, p. 404)

— Produce the fastest diffusion rates — no coordination or authorization required

— Driven mainly by personal perception of the five attributes

— AI example: an individual employee independently adopting ChatGPT for personal productivity (Rogers, 2003, pp. 404-405; cf. mobile-banking-app uptake)

COLLECTIVE DECISIONS

— "Choices to adopt or reject an innovation made by consensus among members of a system" (Rogers, 2003, p. 404)

— Require discussion, negotiation, and agreement — slower diffusion, but often greater legitimacy once adopted

— Common in cooperatives, schools, and community/committee-based structures

— AI example: a firm's practice group collectively agreeing on shared AI usage standards (Rogers, 2003, pp. 405-406; cf. school curriculum reform)

DECISION TYPES III & IV

Authority and Contingent Decisions

AUTHORITY DECISIONS

— "Choices to adopt or reject an innovation made by relatively few individuals possessing power, status, or technical expertise" (Rogers, 2003, p. 406)

— Can yield rapid initial adoption, but risk resistance if not socially endorsed; sustainability depends on communication between decision-makers and implementers (Rogers, 2003, p. 407)

— AI example: a hospital director mandating electronic health-record AI features produced immediate adoption but uneven staff compliance (Rogers, 2003, p. 407)

CONTINGENT DECISIONS

— "The adoption of one innovation depends on prior adoption of another innovation" (Rogers, 2003, p. 408)

— Diffusion speed depends on how widely the prerequisite innovation has already spread

— AI example: agentic AI workflow adoption is contingent on prior adoption of cloud data infrastructure and API integration — firms without that prerequisite cannot adopt agentic tools regardless of enthusiasm (Rogers, 2003, pp. 408-409; cf. online tax filing depending on prior home-computer/internet adoption)

SECTION 5 SYNTHESIS

Decision-Type Comparison: Speed, Control, Legitimacy, Risk, and Governance

Decision Type	Speed	Legitimacy	Governance Risk	Best Suited For
Optional	Fastest	Low (individual)	High — ungoverned Shadow AI risk	Low-stakes personal productivity tools
Collective	Slow	High (consensus)	Low — vetted before rollout	Cross-functional standards, shared workflows
Authority	Fast (initial)	Variable — depends on endorsement	Medium — compliance may be superficial	Compliance-critical, regulated deployments
Contingent	Depends on prerequisite	Depends on prerequisite adoption	Low if prerequisite is governed	Infrastructure-dependent tools (agentic AI, data-fed systems)

Reference: Rogers, Everett M. 2003. Diffusion of Innovations, 5th ed. New York: Free Press., pp. 403-410 (Types of Innovation-Decisions). Power dynamics — pay structure, culture, and organizational hierarchy — materially shape how Authority decisions are actually received.

06

PART 4 · SECTION 6

Communication Channels, Networks, and Opinion Leaders

Channels control how fast awareness spreads — but interpersonal trust controls whether it converts to adoption

SECTION 6

Mass Media vs. Interpersonal Channels

MASS MEDIA CHANNELS

— "Means of transmitting messages that involve a mass medium, such as radio, television, newspapers, and the Internet" (Rogers, 2003, p. 199)

— Most influential at the Knowledge stage — rapidly informing large audiences but providing little interpersonal reinforcement

— AI example: vendor press releases, LinkedIn, and industry conferences build awareness-knowledge broadly and quickly

INTERPERSONAL CHANNELS

— "More effective in persuading individuals to adopt new ideas, especially when such channels link near peers" (Rogers, 2003, p. 199)

— Crucial during Persuasion and Decision stages — enable trust, clarification, and observation of others' experience

— The degree of homophily (similarity between communicator and receiver) strongly influences credibility (Rogers, 2003, p. 306)

Communication channels link those who know about the innovation to those who do not, controlling how fast awareness spreads (Rogers, 2003, p. 199) — but Ryan & Gross showed peer-to-peer exchange, not the initial mass channel, drove final adoption.

SECTION 6

Homophily and Heterophily in AI Diffusion

HOMOPHILY

— "Communication is most effective when source and receiver are homophilous — that is, similar in certain attributes such as beliefs, education, and social status" (Rogers, 2003, p. 305)

— AI example: a peer accountant demonstrating an AI audit tool is more persuasive to other accountants than the same demonstration from an outside IT consultant

HETEROPHILY

— Effective diffusion networks balance homophily and heterophily — enough similarity for credibility, enough difference for genuinely new information (Rogers, 2003, p. 306)

— A source too similar to the receiver may not have novel information to transmit; a source too different lacks credibility

— AI example: cross-functional 'AI office hours' pairing technical staff with domain experts deliberately introduces productive heterophily

SECTION 6

Cosmopolite and Localite Channels

COSMOPOLITE CHANNELS

— Connect individuals "to sources outside their own social system" (Rogers, 2003, p. 199)

— Early adopters tend to be more cosmopolite, maintaining external contacts that introduce new ideas into the system (Rogers, 2003, pp. 200-201)

— AI example: attending an outside AI industry conference, following external researchers, engaging with vendor communities

LOCALITE CHANNELS

— Transmit information within the social system itself

— Dominant channel for the Late Majority and Laggards, who rely primarily on internal, familiar sources rather than external ones

— AI example: an internal AI users' group or internal wiki that circulates locally-validated practices

Diffusion strategy should combine cosmopolite scouting (bringing new AI practice in from outside) with localite legitimation (validating it through trusted internal channels) before it reaches the Majority.

SECTION 6

Opinion Leaders and Networks: Why the CIO May Not Be the Most Influential Voice

1. Opinion leadership is 'the degree to which an individual is able to influence others' attitudes or overt behavior informally' (Rogers, 2003, p. 27) — a function of network position and trust, not formal title.

2. The structure and density of communication networks affect diffusion rate: clustered, siloed systems slow adoption, while highly connected systems accelerate it (Rogers, 2003, p. 308).

3. AI-specific implication: a firm's CIO or Chief AI Officer holds formal authority but may not hold genuine opinion leadership if staff do not view IT leadership as a credible peer voice on day-to-day workflow judgment — formal authority and informal influence are distinct resources.

4. The actual opinion leaders for AI adoption are often mid-level staff who are early, credible, visible users respected by their immediate peer group — precisely the Early Adopters identified in Section 4, not necessarily anyone in a named 'AI' role.

SECTION 6

AI Champions: Design Principles and Common Failure Modes

WHAT WORKS

— Selecting champions for genuine peer credibility (homophily) and network centrality — not just technical enthusiasm

— Pairing a homophilous internal champion with heterophilous outside expertise (Section 6)

— Giving champions real time and organizational support, not an unpaid extra duty

— Making champion results observable (Section 3) through internal show-and-tell, not just private success

COMMON FAILURE MODES

— Choosing a champion for technical enthusiasm (an Innovator) rather than peer credibility (an Early Adopter)

— Treating the champion role as an unfunded volunteer add-on with no protected time

— Champion isolation: a single champion with no network of aides (Rogers' term, Section 8) cannot scale their own influence

— No feedback loop back to leadership — champion insights about real workflow friction never reach decision-makers

SECTION 6 SYNTHESIS

A Practical Communication Strategy for AI Diffusion — Seven Steps

Step	Action	Rogers Concept Applied
1	Build broad awareness through mass channels (email, town halls, intranet)	Mass media — most effective at the Knowledge stage
2	Identify genuine opinion leaders through informal network mapping, not org charts	Opinion leadership as a network property
3	Recruit Early Adopters (not just Innovators) as visible pilot users	Adopter categories; change-agent early-adopter targeting
4	Pair homophilous internal champions with heterophilous outside expertise	Homophily / heterophily balance
5	Translate cosmopolite (external) learning into localite (internal) validation	Cosmopolite / localite channels
6	Make results observable through internal case studies and dashboards	Observability (Section 3)
7	Build a feedback loop from champions and users back to governance and leadership	Change-agent feedback loops (Section 8)

Reference: Rogers, Everett M. 2003. Diffusion of Innovations, 5th ed. New York: Free Press., pp. 199-201, 300-320 (Communication Channels and Diffusion Networks chapters).

07

PART 4 · SECTION 7

The Social System and Organizational AI Adoption

Structure, culture, incentives, governance, and trust set the boundaries within which diffusion occurs

SECTION 7

The Social System: Structure and Its Effect on Diffusion

1. A social system is 'a set of interrelated units that are engaged in joint problem-solving to accomplish a common goal' (Rogers, 2003, p. 23) — it provides the context for communication structure, norms, and diffusion boundaries.

2. 'The structure of a social system is the patterned arrangements of the units within it' (Rogers, 2003, p. 24). Structural properties — hierarchy, network density, role differentiation — determine how efficiently new ideas flow.

3. Formal organizations (corporations, bureaucracies) often exhibit slower diffusion due to rigid structures, while informal networks enable flexibility (Rogers, 2003, p. 229).

4. AI application: a highly centralized, hierarchical firm (many approval layers between an idea and a decision) will structurally diffuse AI more slowly than a flat, cross-functional organization, independent of either firm's enthusiasm for the technology.

SECTION 7

Organizational Culture: Supportive vs. Inhibiting AI Diffusion

SUPPORTIVE CULTURE CHARACTERISTICS

— Rogers: norms are 'the established behavior patterns for members of a social system' (Rogers, 2003, p. 26); systems with open communication and tolerance for deviance adopt more quickly (Rogers, 2003, pp. 232-233)

— Psychological safety to experiment and report AI failures without punishment

— Visible leadership modeling of AI use (not just mandating it for others)

— Reward structures that credit efficient outcomes, not just visible hours worked

INHIBITING CULTURE CHARACTERISTICS

— Highly conformist systems resist change (Rogers, 2003, pp. 232-233)

— Fear of visible error (a hallucinated figure, a wrong citation) in high-accountability professions discourages experimentation

— Leadership that mandates AI use for staff while not adopting it themselves

— Reward structures tied to billable hours or visible activity that AI-driven efficiency directly threatens (see Incentives, next slide)

SECTION 7

Incentives: Billable Hours, Status, and the Threat to Autonomy

1. Professional-services incentive structures built around billable hours create a direct, quantifiable disincentive to adopt time-saving AI tools — an efficiency gain that shortens billable time can reduce individual compensation under an hours-based model.

2. Status incentives matter as much as compensation: professionals whose status derives from being the visible source of expertise (the person who 'knows the answer') may resist tools that appear to substitute for that expertise, independent of firm-level financial incentives.

3. Autonomy threats: AI tools that structure or monitor how work is done (workflow AI, not just assistive AI) can be perceived as reducing professional autonomy — a compatibility issue (Section 3) with direct incentive consequences.

4. Strategic implication: firms that do not redesign incentive structures alongside AI rollout (e.g., shifting toward value-based or outcome-based billing) are asking staff to adopt tools that work against their own compensation and status interests.

SECTION 7

Leadership and Governance as a Diffusion Mechanism

LEADERSHIP'S ROLE

— 'A social system's norms are established and maintained by its leadership, who often act as gatekeepers of new ideas' (Rogers, 2003, p. 233)

— The most adaptive systems have leaders who act as bridges between change agents and local members (Rogers, 2003, p. 233)

— Visible executive use of AI tools (not just mandates) models compatibility and reduces perceived professional-identity threat

GOVERNANCE AS DIFFUSION MECHANISM

— Governance is often framed only as a control/risk function — but well-designed governance also reduces the uncertainty that Rogers identifies as diffusion's core obstacle (Rogers, 2003, p. 5)

— Clear, published guardrails (what data can be used, what review is required) reduce the ambiguity that otherwise slows Persuasion and Decision stages

— Poorly designed governance (unclear, slow, or purely prohibitive) increases uncertainty and pushes adoption underground into Shadow AI (Section 3)

Reframe governance from 'the thing slowing adoption down' to 'the mechanism that reduces the uncertainty diffusion theory identifies as the core barrier to adoption.'

SECTION 7

Trust and Industry/Professional Norms

1. Trust operates at two levels relevant to AI diffusion: trust in the technology itself (will it give reliable output) and trust in the organization's governance of that technology (will misuse be caught and addressed fairly).

2. Professional and industry norms often predate and constrain AI-specific policy: an audit profession's norms around independent verification, or a medical profession's norms around clinical judgment, shape what AI use is considered acceptable regardless of internal firm policy.

3. AI application: professional bodies (e.g., AICPA guidance for accounting, medical-board guidance for clinicians) increasingly issue AI-specific professional-conduct guidance that functions as an external social-system norm layered on top of any single employer's internal policy.

4. Strategic implication: firms operating in professions with strong external normative bodies should align internal AI governance with emerging professional-association guidance rather than developing policy in isolation, both for legitimacy and for legal/regulatory defensibility.

SECTION 7 SYNTHESIS

Social System Diagnostic: Scoring Your Organization's AI Readiness

Structure (approval layers)	**MEDIUM**	Fewer approval layers for low-risk pilots = faster diffusion; audit your own approval chain length.
Culture (tolerance for deviance)	**MIXED**	Psychological safety to report AI errors without career penalty is the key marker.
Incentives (billable hours/status)	**LOW-MED**	Time-based compensation structures often actively disincentivize efficient AI use.
Leadership modeling	**MIXED**	Executives who visibly use AI themselves reduce compatibility resistance faster than mandates alone.
Governance clarity and speed	**MEDIUM**	Fast, clear guardrails reduce uncertainty; slow or absent guardrails push use into Shadow AI.
Trust and professional-norm alignment	**MEDIUM**	Internal policy should explicitly reference relevant external professional-association guidance.

Use this as a live diagnostic: score your own organization on each row before assuming the barrier to AI diffusion is purely technological.

08

PART 4 · SECTION 8

Change Agents and Responsible AI Diffusion

The individuals and functions whose deliberate work accelerates — and governs — AI adoption

SECTION 8

Change Agents: Definition and the Credibility Requirement

1. A change agent is 'an individual who influences clients' innovation-decisions in a direction deemed desirable by a change agency' (Rogers, 2003, p. 368). Their task is to reduce uncertainty, translate technical concepts into practical terms, and link external sources of innovation with local users.

2. 'A change agent's success is positively related to the degree to which clients trust the agent' (Rogers, 2003, p. 372). Effective agents demonstrate empathy, reliability, and communication competence.

3. AI application: an internal AI transformation office, a designated AI champion network, or an external consulting team can all function as Rogers-style change agents — but only to the degree staff actually trust them, which is a function of demonstrated competence and genuine responsiveness, not formal authority.

4. 'Aides are near peers of the clients and can therefore bridge the heterophilous gap between professional change agents and local people' (Rogers, 2003, p. 375) — this is the formal basis for pairing external change agents with internal peer aides (echoing Section 6's champion-pairing recommendation).

SECTION 8

Change-Agent Functions Applied to AI Implementation

Rogers' Function	AI Implementation Application
Developing a need for change	Diagnose specific workflow pain points AI can address — not a generic 'we need an AI strategy' mandate
Establishing an information-exchange relationship	Build two-way channels (office hours, feedback forms) — not one-way policy announcements
Diagnosing problems	Identify the actual stage (Knowledge/Persuasion/Decision/Implementation/Confirmation) where a specific team is stuck
Creating intent to change	Use credible Early Adopter demonstrations and observable results (Section 3), not mandates alone
Translating intent into action	Provide packaged, supported tools and workflows — reduce the complexity barrier (Section 3)
Stabilizing adoption and preventing discontinuance	Monitor for disenchantment discontinuance (Section 2) and respond with support before staff quietly revert
Achieving a terminal relationship (self-sufficiency)	Transition ownership to internal teams so the change agent's involvement is no longer required for continued use

Reference: Rogers, Everett M. 2003. Diffusion of Innovations, 5th ed. New York: Free Press., p. ~369-371 (seven change-agent roles/functions) applied to AI-specific practice.

SECTION 8

Change-Agent Strategies: Demonstration Projects and Early-Adopter Targeting

DEMONSTRATION PROJECTS

— Reduce uncertainty through observability (Rogers, 2003, pp. 379-382; connects directly to Section 3's Observability attribute)

— A bounded, visible pilot with measured, publicized results does more to persuade the Majority than abstract argument

— AI example: a single, well-instrumented department-level AI pilot with published before/after metrics

EARLY-ADOPTER TARGETING

— Uses social influence by concentrating initial change-agent effort on credible Early Adopters, not the broadest possible audience (Rogers, 2003, pp. 379-382; connects to Section 4)

— Early Adopters' visible, successful use persuades the Early Majority far more efficiently than direct persuasion of the Majority itself

— AI example: equipping a small number of respected, credible staff first, rather than a simultaneous firm-wide rollout

SECTION 8

Training and Support: Five Distinct Types, Not One

Tool Training	Workflow Training	Principles Training	Risk Training	Manager Training
Basic mechanics — how to access and operate a specific AI tool or interface.	How the tool fits into an actual end-to-end process, not just standalone use.	How the underlying model works, where and why it can fail (Section 2's principles-knowledge).	Data handling, confidentiality, hallucination risk, and professional-liability implications.	How to supervise, review, and evaluate AI-assisted work product — a distinct skill from using the tool oneself.

SECTION 8 SYNTHESIS

Reinforcement, Feedback Loops, and a Responsible Diffusion Framework

1. Rogers' change-agent strategies include 'feedback loops (continuous monitoring and support)' as a distinct, ongoing function, not a one-time training event (Rogers, 2003, pp. 379-382).

2. 'The extent and quality of change agent efforts mediate between innovation attributes and social system readiness' (Rogers, 2003, p. 394) — sustained effort, not a single launch event, determines outcomes.

3. Responsible AI diffusion requires deliberately balancing competing goals: speed of adoption, realized value, risk management, user trust, fairness, data privacy, security, and accountability — optimizing for speed alone at the expense of the others creates downstream failures (compliance incidents, disenchantment discontinuance, reputational harm).

4. Practical framework: before accelerating any AI diffusion initiative, explicitly score it against each of these eight dimensions rather than defaulting to 'how fast can we roll this out.'

09

PART 4 · SECTION 9

Forecasting AI Diffusion

Adoption curves are financial assumptions in disguise — connecting diffusion theory back to Module 3's valuation models

SECTION 9

Determinants of the Rate of Adoption

1. Rogers identifies the rate of adoption as a function of: the five perceived attributes (Section 3); the type of innovation-decision (Section 5); communication channels used (Section 6); the nature of the social system (Section 7); and the extent of change-agents' promotional efforts (Section 8).

2. AI-specific determinants layered on top of Rogers' general framework: capital intensity of deployment (Module 3), regulatory/compliance environment by industry, data readiness and infrastructure prerequisites (Contingent decisions, Section 5), and labor-market structure (billable-hour vs. salaried professions, Section 7).

3. No single variable dominates — rate of adoption is the joint product of all of these factors interacting, which is why simple, single-driver AI adoption forecasts (e.g., 'cost will fall, so adoption will rise') are consistently unreliable predictors on their own.

SECTION 9

AI Diffusion Across Levels of Analysis

Individual / Team

Fastest-moving level; Optional decisions dominate; high trialability drives rapid personal uptake and Shadow AI (Section 3).

Firm / Organization

Slower; Collective and Authority decisions dominate; bounded by social system, incentives, and governance (Section 7).

Industry

Slower still; shaped by competitive dynamics, regulatory posture, and Contingent-decision infrastructure prerequisites (Section 5).

Consumer Market (B2C)

Distinct dynamics from enterprise B2B diffusion — driven by price, convenience, and network effects rather than governance and liability.

Society

Slowest and broadest level; involves labor-market adjustment, education systems, and public policy — largely outside this module's scope.

Do Not Conflate

Rapid individual/consumer adoption does not imply rapid enterprise/industry adoption — the Pilot-to-Scale Gap (Section 2) demonstrates these levels diverge sharply.

SECTION 9

B2C vs. B2B AI Diffusion: Different Games Entirely

B2C (CONSUMER) DIFFUSION

— Optional decisions dominate; individual price sensitivity and convenience drive adoption

— Very high trialability (free tiers, low switching cost) accelerates diffusion sharply

— Network effects (a tool becomes more useful/normal as more peers use it) reinforce rapid uptake

— Example dynamic: consumer generative-AI chat tools reaching very high usage rates within roughly two years of public release

B2B (ENTERPRISE) DIFFUSION

— Collective and Authority decisions dominate; procurement, security review, and governance gate adoption

— Lower effective trialability at the organizational level — a 'free trial' for one user is not the same as a governed enterprise deployment

— Compatibility with professional identity, liability exposure, and incentive structures (Section 7) create friction largely absent at the consumer level

— Example dynamic: the Pilot-to-Scale Gap (Project NANDA, 2025) showing roughly 95% of enterprise pilots failing to reach P&L impact despite high individual-level usage

SECTION 9

Industry AI-Adoption Readiness: An Analytical Framework, Not a Ranking

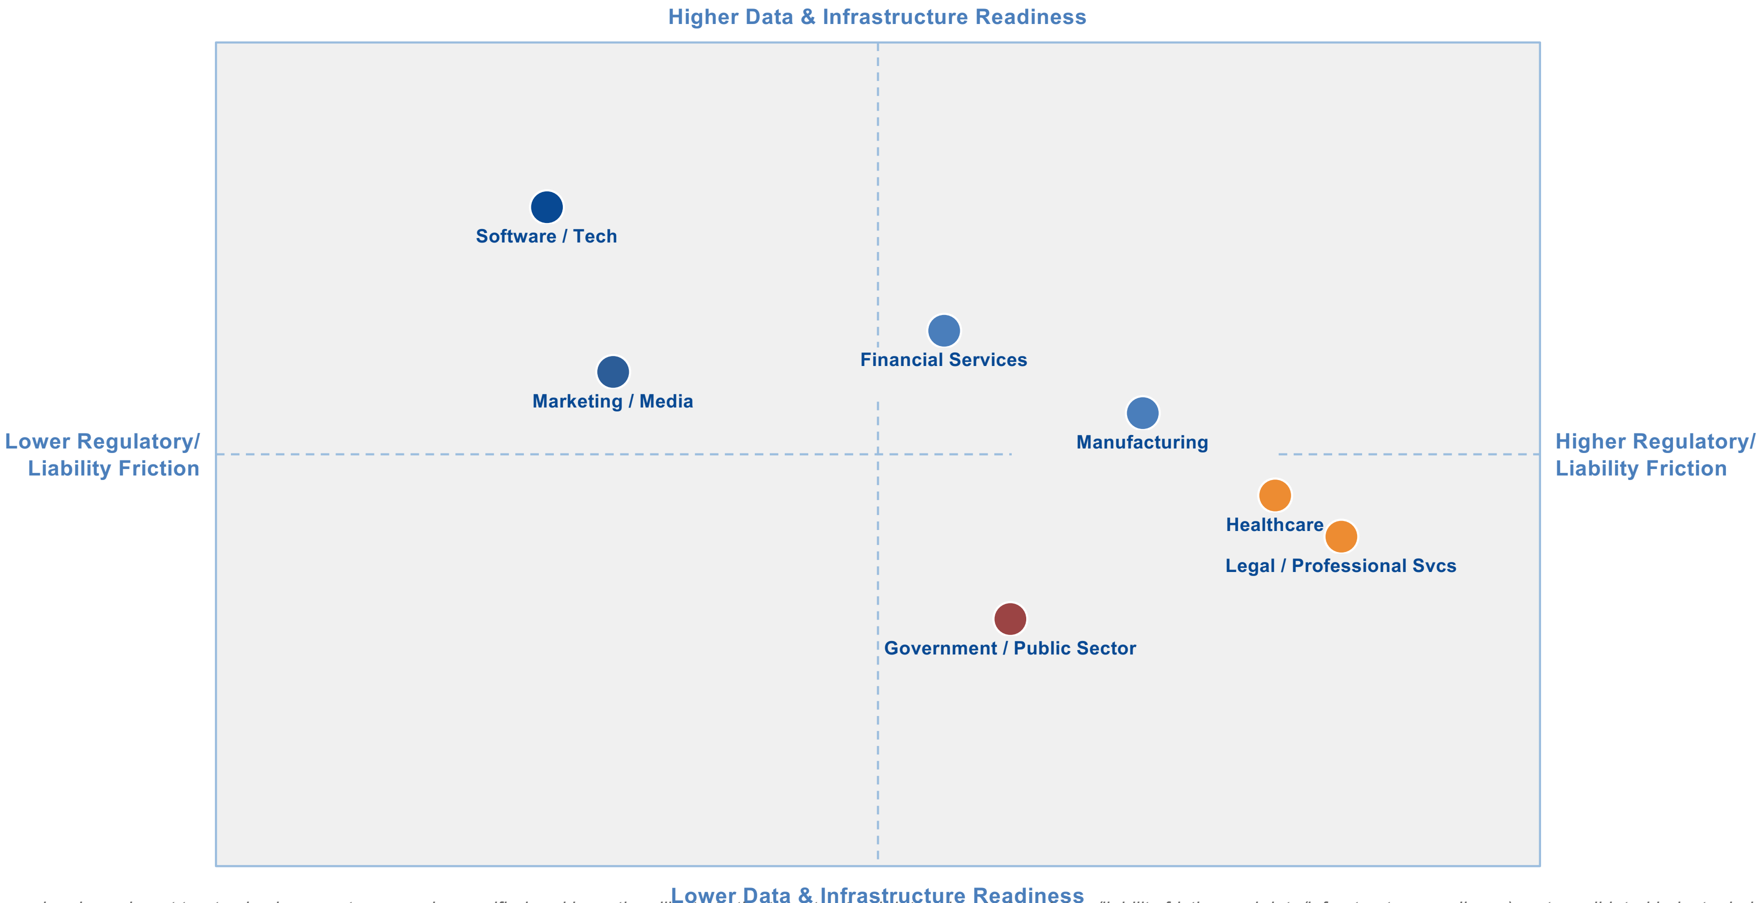

Framework only — do not treat axis placements as precise, verified rankings; they illustrate the two structural determinants (regulatory/liability friction and data/infrastructure readiness), not a validated industry index. Positions should be refreshed with current, named benchmark data before precise claims are made.

SECTION 9

Leading Indicators of AI Diffusion Worth Tracking

1. Stage-specific indicators: training completion and how-to/principles-knowledge assessment scores (Knowledge stage); pilot opt-in and sentiment survey rates (Persuasion); approved-budget-to-active-project conversion (Decision); active weekly users and workflow-integration depth (Implementation); retention and realized-ROI tracking versus original business case (Confirmation).

2. Structural indicators: number of approval layers for AI pilots (Section 7); presence and clarity of published AI governance guardrails; whether incentive/compensation structures have been adjusted alongside AI rollout (Section 7).

3. Change-agent indicators: existence of a credible Early Adopter-led demonstration project; presence of a homophily/heterophily-balanced champion network (Section 6); whether feedback loops from users back to governance actually function (Section 8).

4. Financial indicators (bridging to Module 3): whether earnings-call and 10-K disclosures distinguish pilot-stage from scaled-deployment AI spending, and whether management provides workflow-level (not just budget-level) evidence of Implementation-stage progress.

RECONNECTING TO MODULE 3

Adoption Curves Are Financial Assumptions in Disguise

Every DCF, every multiple, every capital-budgeting decision in Module 3 contained an implicit diffusion forecast — Rogers' framework makes that forecast explicit, testable, and falsifiable.

Adoption Does Not Occur in a Vacuum

Diffusion theory explains how individuals and organizations evaluate and adopt innovation. But adoption also depends on whether the surrounding physical, legal, economic, and social systems can support it.

A willing adopter may still be unable to adopt.

PART 5

The Constraints and Consequences of AI Diffusion

Energy, Regulation, Labor, Trust, and the Limits of Adoption at Scale

What could prevent AI from diffusing as quickly—or as broadly—as markets currently expect?

SECTION 1 · ADOPTION DOES NOT OCCUR IN A VACUUM

A Willing Adopter May Still Be Unable to Adopt

1. Technology readiness -- the tool itself must function reliably enough for the intended use case, at an acceptable error rate and cost.

2. Organizational readiness -- the firm must have the data, workflows, budget, management support, and change capacity to actually deploy and sustain it.

3. System readiness -- the wider environment must permit and support adoption: electricity and data-center capacity, legal permission, professional standards, cybersecurity, labor-market capacity, and public legitimacy.

Rogers explains how innovations move through a social system. This module examines the wider constraints acting on that social system.

AI Diffusion May Stall Even When the Technology Works

1. Physical Capacity
Electricity, data centers, chips, cooling, connectivity, construction lead times.

2. Economic Feasibility
Cost of compute, energy, and implementation relative to measurable benefit.

3. Institutional & Regulatory Permission
Law, liability, professional standards, data and intellectual-property rules.

4. Organizational & Workforce Capability
Skills, workflow redesign, training, management readiness.

5. Social Legitimacy & Trust
Employee, customer, community, and public acceptance.

AI diffusion may stall even when the technology works and users perceive an advantage.

SECTION 2 · PHYSICAL CONSTRAINTS

AI May Feel Digital, but Its Constraints Are Physical

Electricity Generation
Large models require large, dependable power supplies to train and run.

Transmission Capacity
Power must reach the facility, not just exist somewhere on the grid.

Data Centers
Physical buildings, racks, and space take years to plan and build.

Cooling
Dense compute generates heat that must be continuously removed.

Chips & Memory
Specialized processors and high-bandwidth memory are supply-constrained.

Network Connectivity
High-capacity fiber links facilities to users and to each other.

Construction Lead Times
Permitting, materials, and skilled labor all take real time to secure.

Software can diffuse quickly. The physical systems supporting it often cannot.

SECTION 2 · PHYSICAL CONSTRAINTS

Electricity Availability May Shape Where AI Can Scale

1. AI training and inference facilities require large, dependable power supplies -- not just any electricity, but consistent, high-capacity, low-interruption supply.

2. New generation and transmission capacity take years to plan, permit, and build -- far longer than the timeline of a typical AI product cycle.

3. Local power limitations -- an already-constrained grid, limited substation capacity, or slow interconnection queues -- can delay new AI capacity regardless of capital available.

4. Energy cost is a direct input to the economics of training and operating AI systems, and therefore affects which use cases are economically viable in which locations.

DISCUSSION QUESTION

Could access to electricity become a source of competitive advantage?

SECTION 2 · PHYSICAL CONSTRAINTS

A Global Innovation Can Depend on a Small Number of Suppliers

1. Advanced chip design is concentrated among a small number of specialized firms with the engineering expertise to design leading-edge processors.

2. Leading-edge fabrication -- actually manufacturing the most advanced chips -- is concentrated among an even smaller number of foundries capable of it.

3. The specialized manufacturing equipment those foundries depend on is itself produced by a small number of highly specialized suppliers.

4. Advanced packaging and high-bandwidth memory may become bottlenecks in their own right, independent of raw chip-fabrication capacity.

Supply disruptions or capacity limits can slow diffusion throughout every downstream layer.

SECTION 2 · PHYSICAL CONSTRAINTS

Not Every Form of AI Requires the Same Resources

Lower-Resource Uses	Higher-Resource Uses
Document summarization	Frontier-model training
Narrow prediction tools	Large-scale video generation
Embedded workplace assistants	Autonomous agent networks
Small local models	High-volume global inference

Teaching message: Resource constraints may shape which forms of AI diffuse first -- lower-resource uses face fewer physical bottlenecks and may scale faster than higher-resource uses.

SECTION 3 · REGULATION, LAW, AND PROFESSIONAL PERMISSION

Rules Do More Than Restrict Adoption

REGULATION MAY SLOW ADOPTION

— Compliance costs that raise the price of implementation

— Regulatory uncertainty that delays investment decisions

— Approval delays before a use case can be deployed

— Liability exposure for adopting organizations

— Conflicting jurisdictions with inconsistent requirements

REGULATION MAY ENABLE ADOPTION

— Clearer standards that reduce ambiguity about what is permitted

— Greater trust from customers, employees, and the public

— Defined accountability when something goes wrong

— Safer implementation through required controls and testing

— Reduced uncertainty that unlocks previously stalled investment

A lack of rules can be as limiting as overly restrictive rules.

AI Does Not Diffuse at One Speed Across All Industries

Sector	Typical Pace of AI Adoption
Advertising / Media	Faster -- low consequence of error, high tolerance for experimentation
Software Development	Faster -- testable in controlled environments, rapid feedback loops
Financial Services	Moderate -- meaningful benefit, but high regulatory intensity and data sensitivity
Healthcare	Slower -- high consequence of error and significant professional liability
Legal Services	Slower -- high professional liability and data confidentiality concerns
Education	Moderate -- mixed regulatory intensity, uneven institutional readiness
Government	Slower -- high public accountability, cautious procurement and testing
Transportation	Moderate to slower -- physical safety consequences, heavy regulatory intensity

Adoption speed depends on: consequences of error; regulatory intensity; data sensitivity; professional liability; ability to test safely; clarity of measurable benefit. Illustrative characterizations, not a ranking or forecast.

Uncertainty About Data and Output Ownership Can Delay Adoption

Training-Data Disputes
Open questions about what data may lawfully train a model.

Ownership of Generated Material
Who owns content an AI system produces.

Licensing
Terms governing use of underlying models and data.

Confidential Information
Handling of proprietary or sensitive inputs.

Attribution
Whether and how sources must be credited.

Use of Protected Internal Data
Rules for using a firm's own protected data in AI systems.

Organizations may delay adoption when the rights associated with inputs and outputs remain unclear.

The Same Data That Improves AI Can Reduce Trust

MORE DATA ENABLES

— Better personalization of outputs and recommendations

— More accurate prediction and decision support

— Faster, more relevant responses to individual needs

— Improved model performance over time

MORE DATA CAN RAISE

— Privacy concerns about what is being collected

— Surveillance concerns about how data is being used

— Loss of a sense of anonymity or control

— Distrust of the organizations holding the data

Discussion prompt: When does personalization begin to feel like monitoring?

AI Expands Both Defensive Capability and Attack Surface

Faster Threat Detection
AI can identify anomalies faster than manual review.

Automated Security Analysis
AI can process far more security signals at scale.

More Convincing Fraud & Phishing
AI can also make attacks more realistic and harder to detect.

Data Leakage Through AI Tools
Sensitive data may be exposed through AI tool use.

Model Manipulation
Systems can potentially be deceived or manipulated by bad actors.

Dependence on Outside Vendors
Security now depends partly on third-party AI providers.

Access-Control Problems
Who can query, train on, or extract sensitive information.

Security concerns can slow adoption, but AI may also become necessary to manage AI-enabled threats.

AI Usually Reaches Tasks Before It Reaches Entire Jobs

Routine Analytical Tasks
Structured analysis with clear, repeatable steps.

Communication Tasks
Drafting, summarizing, and responding to routine messages.

Judgment
Weighing ambiguous, high-stakes, or novel situations.

Relationship Management
Building and sustaining trust with people over time.

Accountability
Owning outcomes and answering for decisions made.

Physical Activity
Tasks requiring in-person, physical presence and action.

Coordination
Aligning people, priorities, and resources across a team.

Adoption is more likely to reorganize bundles of tasks before eliminating complete occupations.

AI Can Affect Work in Three Different Ways

1. Automation
AI performs an existing task that a person previously performed.

2. Augmentation
AI helps a person perform a task better or faster; the person remains central.

3. Creation
AI enables a task, role, or product that did not previously exist at all.

DISCUSSION QUESTION

Which of these is most likely in your profession over the next three years?

A Faster Task Does Not Necessarily Produce a More Productive Organization

Additional Review
Someone must check the AI-assisted output before it is used.

Verification Requirements
New processes to confirm accuracy and appropriateness.

Output of Questionable Value
Faster production of work that still needs revision.

Workflow Fragmentation
The task speeds up but the surrounding process does not.

Coordination Costs
New handoffs and dependencies introduced by the tool.

Employee Resistance
Time and effort spent managing pushback or reluctance.

New Security Procedures
Added steps required to use the tool safely.

Time Correcting Errors
Time spent fixing mistakes the tool introduced.

Task-level time savings must survive the entire workflow before becoming organizational productivity.

Benefits That Cannot Be Measured May Not Be Sustained

1. **AI Tool Use**

   What is being measured?

2. **Task Improvement**

   Against what baseline?

3. **Workflow Improvement**

   Over what period?

4. **Business Outcome**

   Who receives the benefit?

5. **Financial Result**

   Who bears the cost?

AI May Build Some Skills While Weakening Others

POTENTIAL SKILL DEVELOPMENT

— Faster learning through instant, on-demand explanation

— Broader access to expertise that was previously scarce

— More consistent feedback and coaching on work product

— Support for more advanced analysis than before

POTENTIAL DESKILLING

— Reduced practice performing the underlying task manually

— Growing dependence on the tool being available and correct

— Weaker independent judgment when the tool is absent

— Gradual loss of foundational knowledge over time

Adoption can change not only productivity, but also the future capability of the workforce.

What Happens When AI Performs the Work Through Which Beginners Learn?

1. Junior financial analysis -- the first-draft models and comparisons that used to teach analysts how markets and companies actually behave.
2. Legal research -- the case-law and precedent searches that used to build a young attorney's substantive knowledge.
3. Coding exercises -- the debugging and implementation work that used to build a new developer's practical fluency.
4. Document review -- the detailed reading that used to build familiarity with real-world documents and edge cases.
5. Marketing drafts -- the early writing assignments that used to build a junior marketer's voice and judgment.
6. Basic diagnostic support -- the pattern-recognition repetitions that used to build a new clinician's or technician's intuition.

DISCUSSION QUESTION

If entry-level tasks are automated, how will future experts develop experience and judgment?

Productivity Gains and Economic Gains May Go to Different Groups

Workers Those whose tasks are automated or augmented.	**Customers** Those who may benefit from lower prices or better service.	**Managers** Those who oversee redesigned workflows.	**Shareholders** Those who own the firm capturing efficiency gains.
Technology Vendors Those who sell the AI tools and platforms.	**Infrastructure Providers** Those who supply power, compute, and data-center capacity.	**Highly Skilled Professionals** Those whose judgment becomes more valuable, not less.	**Owners of Proprietary Data** Those who control data that makes AI more valuable.

AI may create aggregate value while distributing that value unevenly.

SECTION 5 · PUBLIC ACCEPTANCE, TRUST, AND SOCIAL RESISTANCE

A Technology Can Work and Still Lack Permission to Scale

Legal Permission Formal law and regulation permitting the use case.	**Managerial Approval** Internal sign-off to deploy within the organization.	**Professional Acceptance** Endorsement by the relevant professional community.
Employee Acceptance Willingness of the workforce to actually use the tool.	**Customer Acceptance** Willingness of customers to interact with or rely on it.	**Community Acceptance** Willingness of the surrounding community to host it.

Adoption requires more than usefulness. It requires legitimacy.

Resistance Is Not Always Irrational

Job Security Concern about losing work to automation.	**Fairness** Concern about who benefits and who bears the cost.	**Privacy** Concern about personal data being collected or used.	**Loss of Autonomy** Concern about reduced control over one's own decisions.	**Safety** Concern about errors with real physical or financial harm.
Reduced Human Contact Concern about losing valued personal interaction.	**Opacity** Concern about not understanding how a decision was made.	**Concentration of Power** Concern about a small number of firms gaining outsized control.	**Environmental Impact** Concern about energy and resource consumption.	**Distrust of Vendors/Leadership** Concern rooted in past experience, not the tool itself.

Resistance may contain useful information about risks the implementation team has overlooked.

SECTION 5 · PUBLIC ACCEPTANCE, TRUST, AND SOCIAL RESISTANCE

AI Diffusion Can Produce Local Costs and Distant Benefits

Electricity Demand Local grid strain from new data-center load.	**Water Consumption** Local water use for cooling large facilities.	**Land Use** Conversion of local land to industrial infrastructure.	**Noise** Ongoing operational noise affecting nearby residents.
Construction Temporary disruption during buildout.	**Ratepayer Effects** Possible upward pressure on local electricity rates.	**Tax Incentives** Local tax revenue forgone to attract the facility.	**Limited Local Employment** Facilities are often capital-intensive, not labor-intensive.

The benefits of AI may be broadly distributed while infrastructure costs are geographically concentrated.

People Do Not Have One General Level of Trust in AI

HIGHER COMFORT DELEGATING TO AI

— Music and media recommendations

— Grammar and writing correction

— Calendar and schedule organizing

— Routine, low-stakes task reminders

LOWER COMFORT DELEGATING TO AI

— Illness diagnosis

— Employment evaluation

— Credit determination

— Legal interpretation

— Military decisions

Adoption forecasts must specify the task, consequence, and decision context.

Not Every AI Use Case Requires the Same Controls

Lower Risk	**LOWER**	Drafting, brainstorming, internal summarization -- light-touch controls, fast experimentation.
Moderate Risk	**MODERATE**	Customer interaction, operational recommendations, employee support -- defined review and monitoring.
Higher Risk	**HIGHER**	Healthcare decisions, credit, hiring, legal rights, public benefits, physical safety -- strict controls and human oversight.

Excessive controls can block low-risk experimentation; weak controls can make high-risk use unacceptable.

Move Fast and Govern Well Is Harder Than Either One Alone

	WEAK GOVERNANCE	STRONG GOVERNANCE
FAST ADOPTION	**Uncontrolled Exposure** Speed without control -- incidents, backlash, and intervention become more likely.	**TARGET QUADRANT** **Responsible Scaling** Speed and durability together -- the quadrant most organizations are aiming for.
SLOW ADOPTION	**Stagnation** Neither speed nor control -- value creation stalls without a corresponding safety benefit.	**Cautious Learning** Safe and deliberate, but may be too slow to capture the opportunity.

Excellent governance without speed forfeits opportunity. Excellent speed without governance forfeits durability. Both are genuinely hard to sustain together.

SECTION 6 · RESPONSIBLE ADOPTION AT SCALE

What Should We Watch Before Financial Results Appear?

Pilots to Sustained Use Movement beyond one-off pilots into ongoing use.	**Active Use vs. Licenses** Actual usage, not just purchased seats or licenses.	**Workflow Integration** The tool is embedded in real day-to-day processes.	**Employee Training** Structured investment in building employee capability.
Management Accountability Named owners responsible for outcomes.	**Documented Use Cases** Specific, written descriptions of what is being solved.	**Incident Rates** Frequency and severity of problems encountered.	**User Trust** Willingness of employees and customers to rely on it.
Measurable Task Outcomes Concrete evidence of improved task-level results.	**Expansion Across Functions** Use spreading to new teams and business functions.	**Discontinuation Rates** How often initiatives are stopped or abandoned.	**Infrastructure Utilization** How much of built capacity is actually being used.

These are leading indicators, not lagging financial metrics -- track them before quarterly results confirm or refute an adoption thesis.

SECTION 7 · IMPLICATIONS FOR MARKETS AND STRATEGY

The Adoption Curve Is Not Predetermined

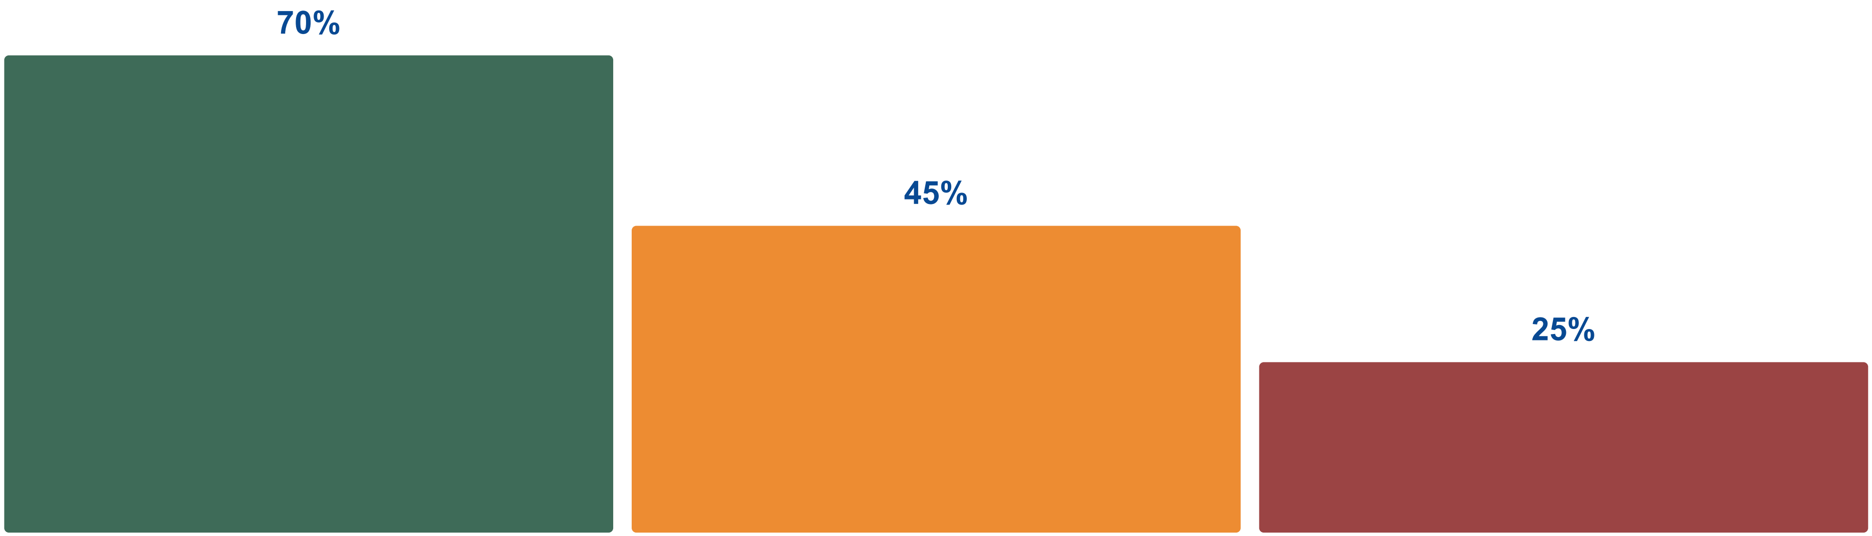

Rapid Diffusion	Delayed but Eventual Diffusion	Limited or Fragmented Diffusion
Early takeoff, steep slope, high ultimate penetration across most sectors.	Slower start, gradual acceleration once constraints ease over time.	Uneven uptake, concentrated in a few sectors, tasks, or geographies.

Strong Interest Does Not Guarantee Fast Revenue Growth

Available Capacity Power, compute, and data-center capacity to serve demand.	**Implementation Cycles** Time required to actually deploy and embed a solution.	**Budget Approval** Internal capital-allocation and procurement processes.	**Integration** Connecting new tools to existing systems and data.
Regulation Legal and compliance clearance for the use case.	**Workforce Readiness** Training and organizational capacity to use the tool.	**Customer Trust** Willingness of customers to actually rely on the outcome.	**Measurable Return** Evidence the investment is producing a real benefit.

Market forecasts often compress a long chain of dependencies into a single growth rate.

The Most Valuable Companies May Remove Adoption Friction

Reduce Energy Use More efficient chips, models, or cooling systems.	**Provide Secure Infrastructure** Reliable, trustworthy compute and data environments.	**Make Integration Easier** Tools that connect AI to existing systems with less effort.
Support Compliance Products that help organizations meet regulatory requirements.	**Verify Outputs** Tools that check AI accuracy and reliability before use.	**Protect Data** Security and privacy safeguards built around AI use.
Redesign Workflows Services that rebuild processes around AI capability.	**Train Employees** Programs that build organizational AI capability.	**Measure Outcomes** Tools that prove whether AI investment is working.

DISCUSSION QUESTION

Which constraints are likely to become profitable markets themselves?

A

THREE POSSIBLE PATHS FOR AI DIFFUSION

Rapid, Broad Diffusion

Infrastructure and governance constraints ease faster than expected, and productivity gains become broadly visible across sectors.

Infrastructure

Power, data-center, and chip capacity expand quickly enough to meet demand.

Cost

Training and inference costs decline, widening the set of economically viable use cases.

Governance

Regulatory and organizational governance frameworks stabilize rather than stall adoption.

Productivity

Measurable productivity and financial benefits become visible within a normal investment horizon.

B

THREE POSSIBLE PATHS FOR AI DIFFUSION

Uneven Diffusion

AI adopts quickly in some tasks and industries while remaining slow in regulated or high-risk settings, concentrating value in a smaller set of firms.

Pace

Strong adoption in low-risk, low-regulation tasks; slow adoption in healthcare, law, government, and safety-critical settings.

Value Concentration

Economic value concentrates in firms that successfully remove adoption friction rather than spreading evenly.

Geography

Adoption clusters where infrastructure, talent, and permissive regulation coincide.

Persistence

Uneven diffusion can persist for years rather than converging quickly to a single pace.

C

THREE POSSIBLE PATHS FOR AI DIFFUSION

Delayed Diffusion

Infrastructure, trust, and implementation problems persist longer than expected, and financial returns arrive later than current market expectations imply.

Infrastructure Lag

Power and data-center capacity build-out takes longer than planned, capping usable capacity.

Trust & Governance

Unresolved trust, legal, and governance questions slow enterprise deployment decisions.

Capital Reassessment

Markets reassess the pace of capital spending as expected returns arrive later than modeled.

Timing Risk

Valuations that assumed rapid diffusion face the greatest downside repricing risk.

Questions Investors Should Ask

Q1 What adoption rate is embedded in the valuation?	**Q2** Which constraints could delay revenue?	**Q3** Who must change behavior for value to appear?	**Q4** Does the company remove adoption friction, or merely provide capability?
Q5 Are benefits measurable?	**Q6** Who bears infrastructure and implementation costs?	**Q7** Could regulation increase trust -- or restrict the market?	**Q8** How sensitive is the valuation to slower diffusion?

Questions Managers Should Ask

Q1 Is the use case worth solving?	**Q2** Does the organization have the required data and infrastructure?	**Q3** Who owns the workflow change?	**Q4** What risks require control?
Q5 How will employees participate?	**Q6** How will success be measured?	**Q7** What would cause the organization to stop or redesign the initiative?	**Q8** What external constraints are outside management's control?

END OF MODULE 5

Constraints Shape Adoption, but Financial Results Reveal Its Scale

Even when AI is available, permitted, trusted, and used, the investment cycle is not financially validated until adoption begins to appear in sales, utilization, margins, cash flow, and returns on capital.

Where should we expect AI diffusion to become visible in company results?

PART 6

Measuring AI Diffusion Through Financial Performance

How Adoption Appears in Revenue, Margins, Cash Flow, and Returns on Capital

How can investors and managers tell whether AI adoption is actually becoming economically meaningful?

Adoption Becomes Visible in Stages -- Not All at Once

Usage → Commercial Demand → Revenue Growth → Capacity Utilization → Margin Improvement → Free Cash Flow → Return on Invested Capital

Revenue may confirm that customers are adopting. Margins and cash flow confirm whether that adoption is economically attractive.

SECTION 1 · FROM ADOPTION TO FINANCIAL RESULTS

Adoption Is Observed Operationally First, Commercially Second, and Financially Last

1. Technical Availability
The tool functions and is accessible for the intended use case.

2. Trial Use
Users experiment with it on a small scale to test whether it fits.

3. Sustained Usage
Users keep coming back after the novelty of trying it wears off.

4. Paid Usage
Customers pay for it, converting usage into recognized revenue.

5. Profitable Usage
The revenue it generates exceeds the cost of delivering it.

6. Value Creation
Profitable, sustained usage compounds into durable financial performance.

Adoption is observed operationally first, commercially second, and financially last.

SECTION 1 · FROM ADOPTION TO FINANCIAL RESULTS

Checking the Evidence Against the Stage

Public reporting on AI is heavily weighted toward the first three stages of the progression on the previous slide -- announcements of new features, user counts, and engagement statistics are easy to produce and easy to report. Evidence for stages 4 through 6: paid usage, profitable usage, and durable value creation is slower to appear and requires more disciplined reading of financial disclosures.

DISCUSSION QUESTION

For a company you follow closely, which of the six stages does the public evidence currently reach -- and which single additional stage of evidence would most change your view of that company's AI story?

SECTION 2 · DIFFERENT BUSINESS MODELS PRODUCE DIFFERENT SIGNALS

The Right Signal to Watch Depends on the Business Model

External AI Sellers

AWS, Oracle -- Primarily external. CoreWeave / Lambda -- Specialist external providers. Sell AI capability and infrastructure directly to other companies.

Internal AI Users

Meta -- Primarily internal use. Uses AI mainly to improve its own products, engagement, and advertising.

Hybrid Firms

Microsoft, Google -- Internal + external. Both sell AI/cloud capability externally and use AI to improve their own products.

Reading the wrong signal for a company's business model produces a false read on its adoption story.

SECTION 2 · EXTERNAL AI SELLERS (AWS, ORACLE, COREWEAVE)

External AI Sellers -- Look for Commercial Traction

1. Cloud/AI revenue growth -- is the AI-related portion of revenue actually accelerating, not just growing off a large existing base?
2. Customer contracts and backlog -- backlog is revenue a customer has already committed to but that has not yet been delivered or recognized; is it growing, and are commitments getting larger and longer?
3. Consumption growth -- are existing customers using more compute over time, not merely signing new contracts?
4. Capacity utilization -- is newly built infrastructure actually being used, or does it sit idle after construction?
5. Pricing and renewals -- are prices holding rather than being discounted away, and are customers renewing rather than churning?
6. Gross margins -- gross margin is revenue minus the direct cost of delivering the service; is there room for profit once that direct cost is covered?

For a seller of AI infrastructure, revenue growth without utilization and margin is a warning sign, not a validation.

Internal AI Users -- Look for Product and Efficiency Gains

1. Advertising effectiveness -- are AI-driven targeting and ranking improvements raising the return advertisers get on their ad spend?
2. Engagement -- are users spending more time on the platform, or returning more often, because of AI-driven features?
3. Recommendation quality -- are AI-powered feeds and recommendations measurably better at holding attention and relevance?
4. Revenue per user -- is monetization per user rising as AI improves targeting, content ranking, and ad matching?
5. Cost efficiency -- is AI lowering the cost of content moderation, operations, or content production at the same output level?
6. New product monetization -- are new AI-native features converting into genuinely new revenue streams, not just improving existing ones?

An internal AI user rarely reports a separate 'AI revenue' line -- adoption shows up inside existing metrics instead.

SECTION 2 · HYBRID FIRMS (MICROSOFT, GOOGLE)

Hybrid Firms -- Look for Signals on Both Sides

1. External cloud/AI sales -- cloud consumption growth, backlog, and AI-service revenue, exactly as tracked for a pure external seller.

2. Internal productivity and product improvement -- AI embedded in existing products (search, productivity software, workspace tools) improving engagement or efficiency, exactly as tracked for a pure internal user.

3. Capital allocation -- how infrastructure spending is split between building AI capability to sell externally and improving internal products.

4. Segment disclosure -- how much detail the company provides that lets an outside observer separate the two signal sets rather than reading one blended headline number.

Hybrid firms give the clearest overall picture of adoption -- but only if both signal sets are tracked separately, not blended into one headline number.

SECTION 2 · DIFFERENT BUSINESS MODELS PRODUCE DIFFERENT SIGNALS

Classify Before You Measure

Every company discussed in this module -- and every AI-exposed company participants are likely to analyze afterward -- falls into one of the three categories just covered. Misclassifying a company (treating a hybrid firm as a pure seller, or an internal user as if it should report a dedicated AI revenue line) is one of the most common ways an otherwise careful analysis goes wrong.

DISCUSSION QUESTION

Pick one AI-exposed company you follow. Which of the three categories does it belong to, and are you currently tracking the signals that actually apply to that category -- or the ones that apply to a different one?

Rising Revenue Is Suggestive, Not Sufficient

MAY INDICATE REAL ADOPTION

— More customers are buying the product or service.

— More workloads are running on the platform.

— Existing customers are using it more deeply.

— Customers are spending more per account over time.

— Pilots are converting into full production deployments.

MAY MISLEAD WHEN IT COMES FROM

— Steep discounts used to win short-term contracts.

— Prepaid commitments recognized before real usage occurs.

— Revenue concentrated in a small number of large customers.

— Customers reserving capacity they are not yet using.

— Revenue created by transfers between related internal entities.

— Growth achieved at an unprofitable price.

Revenue growth becomes a credible adoption signal only when it is paired with evidence that the items on the right are not what is actually driving it.

Capacity Must Be Filled Before Adoption Shows Up in Margins

Infrastructure Spending Occurs First → New Capacity Initially Sits Underused → Customer Adoption Fills Capacity → Fixed Costs Spread Across More Usage → Unit Economics Improve → Margins May Rise

Adoption improves margins only when new demand uses capacity efficiently and at attractive prices.

SECTION 4 · UTILIZATION AS THE BRIDGE TO MARGINS

Why Utilization Determines Whether Adoption Pays Off

A simplified conceptual illustration -- no company-specific figures -- showing how the same fixed infrastructure cost behaves differently depending on how much of it is actually being used.

Item	Detail
Infrastructure cost	**Largely fixed once built -- power, hardware, and depreciation do not fall just because usage is low**
Low utilization	**The fixed cost is spread across a small amount of usage, so unit cost per workload stays high**
Rising utilization	**The same fixed cost is spread across a larger amount of usage, so unit cost per workload falls**
Pricing discipline	**The unit-cost improvement only reaches the bottom line if prices are not discounted away at the same time**

NET EFFECT ON MARGIN

Depends on Utilization

Rising use of existing capacity, at held prices, is what allows margins to improve -- not revenue growth by itself.

Weak Margins Do Not Automatically Mean Weak Adoption

Depreciation
Large infrastructure investments are expensed over time, weighing on margins throughout the buildout phase.

Energy Costs
Running AI infrastructure at scale consumes significant, continuous electricity.

Staffing
Specialized AI, infrastructure, and safety talent commands a premium and adds to cost.

Model-Training Expenses
Developing and retraining models is expensive and continues even after a product launches.

Customer Acquisition
Winning and onboarding new customers carries real, up-front cost.

Price Competition
Rivals competing for the same customers can compress the pricing a company can achieve.

New Data-Center Openings
Each newly opened facility reintroduces a fresh round of underutilized capacity.

Inefficient Early Workloads
Early deployments are often less optimized than mature, tuned ones.

Product Bundling
AI features bundled into existing products may not carry a separate, visible price at all.

This list exists specifically to prevent a common mistake: treating a margin dip during heavy investment as proof that adoption itself is weak.

The Investment Thesis Is Not Validated Until Adoption Produces Cash

1. Operating costs -- the ongoing cost of running the business, which revenue from adoption must first cover.

2. Continuing capital expenditures -- money that must keep being spent on infrastructure, not just the initial build.

3. Financing costs -- interest and other costs of the debt or capital used to fund the investment.

4. Replacement investment -- money needed to replace equipment as it wears out or becomes technologically obsolete.

Adoption is not just becoming widespread. It is generating enough cash to clear all five of these hurdles, in order.

Reading the Signals by Business Model

Illustrative, conceptual comparisons -- not specific financial figures.

Company	Business Model	How AI Creates Value	Main Financial Question
Microsoft	**Internal + external**	Sells Azure and Copilot services while also using AI across its own products and operations	Will new revenue and internal efficiency justify the infrastructure investment?
AWS	**Primarily external**	Sells cloud computing and AI capacity to outside customers	Will customer usage grow fast enough to fill capacity profitably?
Google	**Internal + external**	Sells Cloud and Gemini services while using AI in advertising, search, and operations	Can both channels generate enough revenue and efficiency to support spending?
Oracle	**Primarily external**	Sells cloud and AI infrastructure capacity to outside customers	Will contracted demand arrive quickly enough to justify expansion and financing?
Meta	**Primarily internal**	Uses AI to improve advertising, recommendations, engagement, productivity, and future products	Will internal business improvements justify the spending?
CoreWeave / Lambda	**Specialist external providers**	Rent specialized AI computing capacity to outside customers	Is demand durable, diversified, and profitable enough to support financing needs?

For teaching purposes only. Signals are qualitative and directional, not company-reported figures, and are not investment recommendations.

Some Signals Move Early. Others Confirm Later.

LEADING INDICATORS

- Active users
- Paid seats
- Workloads running on the platform
- Usage volume
- Backlog
- Customer expansion
- Pilot-to-production conversion
- Capacity utilization

LAGGING INDICATORS

- Revenue
- Margins
- Free cash flow
- Return on invested capital
- Productivity gains
- Sustained retention

Leading indicators tell you adoption may be starting. Lagging indicators tell you whether it was worth it.

SECTION 9 · WHAT WOULD VALIDATE THE CYCLE?

Three Scenarios -- No Single Outcome Is Certain

A. Adoption Without Profitability	B. Profitable but Narrow Adoption	C. Broad and Profitable Adoption
Strong usage and revenue; weak margins and cash flow.	Strong economics, but limited to a small number of use cases.	Expanding usage, improving margins, strong cash generation and returns.

None of these is presented as the certain outcome -- the discipline is knowing which evidence would confirm each one.

A

WHAT WOULD VALIDATE THE CYCLE? THREE SCENARIOS

Adoption Without Profitability

Usage and revenue are strong and continue to grow, but margins and cash flow remain weak -- adoption is real but not yet economically attractive.

Usage
Active users, workloads, and revenue all continue to grow at a healthy pace.

Margins
Gross and operating margins stay compressed or fail to expand as expected.

Cash Flow
Free cash flow remains weak relative to the scale of capital being invested.

Market Read
Markets must judge whether this is an early, temporary phase or a more permanent structural outcome.

B

WHAT WOULD VALIDATE THE CYCLE? THREE SCENARIOS

Profitable but Narrow Adoption

Strong economics appear, but only in a limited number of use cases or customer segments, rather than broadly across the business.

Economics

Margins and returns are attractive wherever adoption has actually taken hold.

Breadth

Adoption remains concentrated in a small set of high-value use cases or customers.

Value Concentration

Financial benefit accrues to firms that found the right narrow use case, not to the market broadly.

Persistence

Whether adoption broadens to new use cases without repeating the same favorable economics is unresolved.

C

WHAT WOULD VALIDATE THE CYCLE? THREE SCENARIOS

Broad and Profitable Adoption

Usage expands broadly, margins improve, cash generation strengthens, and returns on invested capital become attractive across a wide range of use cases.

Usage

Expanding usage across many customers, workloads, and use cases, not just one or two.

Margins

Sustained margin improvement as utilization rises and pricing holds.

Cash Generation

Free cash flow becomes strong enough to fund continuing investment internally.

Returns

Return on invested capital exceeds the cost of the capital deployed, across the business.

Questions These Signals Should Force Investors to Ask

Q1
Are valuations assuming revenue growth only?

Q2
Are they also assuming margin expansion?

Q3
How much capacity utilization is required to reach that margin?

Q4
How much continuing capital expenditure is necessary to sustain it?

Q5
What happens to the valuation if adoption is slower than assumed?

Q6
What happens if adoption turns out to be broad but low margin?

Q7
Which firms benefit mainly from internal productivity rather than external sales?

MODULE CLOSE · INTEGRATED SYNTHESIS

Adoption Is Not Fully Proven Until It Produces Returns.

Use → Revenue → Utilization → Margins → Cash Flow → Return on Capital

THE MARKET'S ACTUAL BET

The market is not merely betting that people will use AI. It is betting that usage will become sustained, paid, profitable, and large enough to justify the capital invested.

Revenue Confirms Adoption. Margins and Cash Flow Confirm Whether It Was Worth It.

Every measure in this module exists to answer one question:

Is AI adoption becoming large enough, profitable enough, and durable enough to justify the capital markets have already committed?

"Are Markets Overestimating the Speed of AI Adoption?"

THE FINAL COURSE QUESTION

Are Markets Overestimating the Speed of AI Adoption?

A disciplined debate, not a verdict -- reasonable, well-informed people land in different places.

REASONS ADOPTION MAY BE SLOWER THAN EXPECTED

- Infrastructure takes time to build.
- Workflow redesign is difficult.
- Trust and governance remain unresolved.
- Professional norms may slow use.
- Customers may remain in pilots.
- Benefits may be difficult to measure.
- Revenue may not produce attractive margins.
- Complementary investments are required.

REASONS ADOPTION MAY MOVE FASTER THAN EXPECTED

- AI capabilities are improving.
- Costs may decline.
- Distribution is already widespread.
- Competitive pressure may force adoption.
- AI may be embedded in existing software.
- Internal benefits may appear before separate AI revenue.
- Platforms may reduce implementation difficulty.
- Successful applications may spread rapidly through professional networks.

FINAL TAKEAWAYS

Six Ideas to Carry Forward

1. AI investment and AI adoption are not the same thing.
2. Industry structure shapes who may capture value.
3. Every valuation contains an adoption forecast.
4. Diffusion depends on technology, organizations, and social systems.
5. External constraints may alter the speed and shape of adoption.
6. Adoption is financially validated only when use becomes durable, profitable, and sufficient to justify the capital invested.

FINAL DISCUSSION

What Evidence Would Change Your View?

What evidence over the next 12-24 months would lead you to conclude that AI adoption is moving faster—or slower—than financial markets currently expect?

Enterprise usage	**Pilot-to-production conversion**	**Customer spending**
Hyperscaler revenue	**Capacity utilization**	**Margins**
Free cash flow	**Productivity**	**Regulation**
Labor-market response	**Trust**	**Discontinuation rates**

Thank You!

APPENDICES

References & Source Verification

Full source detail for every citation used throughout the course.

Academic & Theoretical Foundations

Diffusion of Innovation Theory

Rogers, Everett M. 2003. Diffusion of Innovations, 5th ed. New York: Free Press. (Primary theoretical source for Part 4 -- innovation-decision process, five perceived attributes, adopter categories, communication channels, social systems, change agents, Figure 5-1; also underlies Part 5's opening adoption-readiness framework.)

Moore, Geoffrey A. 1991 (rev. 1999, 2014). Crossing the Chasm: Marketing and Selling High-Tech Products to Mainstream Customers. New York: HarperBusiness. (Chasm concept, contrasted with Rogers' continuous adoption curve.)

Corporate Finance & Valuation

Damodaran, Aswath. Investment Valuation: Tools and Techniques for Determining the Value of Any Asset. Wiley. (DCF, WACC, terminal value, reverse-DCF/"expectations investing," capital-intensity/FCF-yield frameworks throughout Part 3.)

Brealey, R., Myers, S., and Allen, F. Principles of Corporate Finance. McGraw-Hill. (NPV, capital budgeting, WACC, real-options frameworks.)

Koller, T., Goedhart, M., and Wessels, D. Valuation: Measuring and Managing the Value of Companies. McGraw-Hill/Wiley (McKinsey & Company). (DCF and terminal-value methodology.)

Strategy & Industrial Organization

Porter, Michael E. 1979. "How Competitive Forces Shape Strategy." Harvard Business Review, Mar.-Apr. 1979.

Porter, Michael E. 1980. Competitive Strategy: Techniques for Analyzing Industries and Competitors. New York: Free Press.

Cusumano, M., Gawer, A., and Yoffie, D. 2019. The Business of Platforms. New York: Harper Business.

Adner, Ron. 2012. The Wide Lens: A New Strategy for Innovation. New York: Portfolio/Penguin.

Katz, Michael L., and Carl Shapiro. 1985. "Network Externalities, Competition, and Compatibility." American Economic Review 75(3).

Arthur, W. Brian. 1994. Increasing Returns and Path Dependence in the Economy. Ann Arbor: University of Michigan Press.

Technology-Adoption & Productivity Research

David, Paul A. 1990. "The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox." American Economic Review Papers and Proceedings 80(2). (Historical technology-diffusion-lag analogy.)

Brynjolfsson, E., Li, D., and Raymond, L. R. 2023 (rev. Nov. 2023). "Generative AI at Work." NBER Working Paper 31161. nber.org/papers/w31161.

Company Filings, Earnings Calls & Investor Relations

Hyperscalers & Cloud

Microsoft Corp. FY2026 Q4 earnings call and press release, July 29, 2026 (CFO Amy Hood); FY2024 Q2 earnings call transcript, Jan. 30, 2024; SEC lease disclosures ($329.1B not-yet-commenced data-center leases).

Amazon.com Inc. Q1/Q2 2026 shareholder letters (Andy Jassy); AWS Trainium product page, updated July 15, 2026.

Alphabet Inc. Q2 2026 earnings coverage and CFO Anat Ashkenazi earnings-call commentary, July 22-26, 2026.

Oracle Corp. FY2026 Q4 8-K, June 10, 2026.

Meta Platforms Inc. Q2 2026 earnings release, investor.atmeta.com, July 29, 2026.

CoreWeave Inc. Q1 2026 results (StockTitan coverage, May 7, 2026).

Energy & Data-Center Infrastructure

Constellation Energy Corp., Form 8-K Exhibit 99.1, filed Sept. 20, 2024 (Crane Clean Energy Center / Three Mile Island Unit 1).

Talen Energy Corp., Form 8-K Exhibit 99.1, filed Aug. 26, 2024 (AWS co-located campus).

GE Vernova Inc., Form 8-K Exhibit 99, Q2 2026 earnings release, filed July 22, 2026.

Vertiv Holdings Co, Form 8-K Exhibit 99.1, filed Apr. 22, 2026.

Equinix Inc., Form 10-K for FY2025, filed Feb. 11, 2026.

Semiconductors & Equipment

NVIDIA Corp., "NVIDIA Announces Financial Results for First Quarter Fiscal 2027," May 20, 2026; SEC EDGAR filing history, CIK 0001045810.

AMD, "AMD Reports First Quarter 2026 Financial Results," May 5, 2026, and AMD IR press releases, July 20 & 22, 2026.

ASML Holding, 2025 Annual Report, published Feb. 25, 2026 (asml.com/en/investors/annual-report/2025).

TSMC, "TSMC Reports Second Quarter EPS of NT$27.25," July 16, 2026; TSMC "Fab Capacity" page, tsmc.com, accessed July 30, 2026.

Applied Materials, Q4/FY2024 earnings release, Nov. 14, 2024.

Micron Technology Inc., Q1 Fiscal 2026 earnings release, Dec. 17, 2025.

Cisco Systems, "Cisco Reports Third Quarter Earnings," May 13, 2026.

Enterprise Applications & Consulting

ServiceNow, Form 8-K, filed July 22, 2026; Salesforce Agentforce pricing page, accessed July 30, 2026; Adobe Q2 FY2026 release, June 11, 2026; Workday Q1 FY2027 release, May 21, 2026.

Tata Consultancy Services press release, July 9, 2026; Infosys IFRS-USD press release, July 23, 2026.

Accenture Q1 FY2026 (Dec. 18, 2025) and Q2 FY2026 (Mar. 19, 2026) results; IBM Q4 2025 8-K (Jan. 28, 2026) and subsequent 2026 releases.

Deloitte Global Newsroom, Sept. 30, 2025; PwC Global Annual Review 2025, Oct. 28, 2025; BCG press release, Apr. 2, 2025.

Industry Research & Market Data Providers

McKinsey & Company

"The State of AI in 2025," Nov. 5, 2025.

"The Cost of Compute: A $7 Trillion Race to Scale Data Centers," Noffsinger, Patel, Sachdeva; April 2025.

Company & Market Data Providers (Value Chain Section)

Company filings & investor-relations disclosures.

Gartner, IDC, Forrester (analyst market-share and category coverage).

CB Insights, PitchBook (private-company valuation and funding data).

SEMI, Omdia, TrendForce (semiconductor-sector data).

EIA and DatacenterMap (energy & data-center capacity data).

DeepSeek & Model-Commoditization Debate

Amodei, Dario. "On DeepSeek and Export Controls." Jan. 2025.

DeepSeek-AI. "DeepSeek-V3 Technical Report." arXiv:2412.19437, Dec. 27, 2024.

Other Named Industry Research

PwC, "2026 AI Business Predictions."

Stanford HAI, 2026 AI Index Report (Economy chapter).

a16z, "Who Owns the Generative AI Platform?" Bornstein, Appenzeller, Casado; Jan. 19, 2023, updated Jan. 2026.

SemiAnalysis, "Google: We Have No Moat, And Neither Does OpenAI," May 4, 2023; "DeepSeek Debates," Jan. 31, 2025.

Epoch AI, "LLM inference prices have fallen rapidly but unequally across tasks," Mar. 12, 2025 (epoch.ai).

Challapally, A., Pease, C., Raskar, R., Chari, P. Project NANDA, MIT Media Lab. "The GenAI Divide: State of AI in Business 2025," July/Aug. 2025.

Deloitte AI Institute, "State of AI in the Enterprise," 2025.

Aldasoro, I., et al. 2026. 12,000-firm European AI-adoption/productivity study (as cited in Part 1).

OECD (Filippucci, F., et al.). 2025. G7 AI productivity-impact projection.

American Medical Association, "2 in 3 physicians are using health AI -- up 78% from 2023," Feb. 26, 2025.

Microsoft/LinkedIn, "2024 Work Trend Index Annual Report," May 8, 2024.

Government, Regulatory, Accounting & Legal Standards

SEC Rules & Regulation

SEC Regulation S-K, 17 CFR §229.101 (Item 101, Business); §229.105 (Item 105, Risk Factors, renumbered per 85 FR 63759, Oct. 8, 2020); §229.303 (Item 303, MD&A). Verified via eCFR, ecfr.gov/current/title-17/chapter-II/part-229.

SEC EDGAR filing histories, including NVIDIA Corp. (CIK 0001045810) and Microsoft Corp. (CIK 0000789019); SEC Standard Industrial Classification (SIC) Code List, sec.gov.

U.S. Government & Classification Standards

U.S. Census Bureau, 2022 NAICS Manual (census.gov); Census Bureau 2027 NAICS revision notice, Jan. 6, 2026.

U.S. Bureau of Industry and Security, export-control guidance, bis.gov, accessed July 30, 2026.

Securities Litigation

Private Securities Litigation Reform Act of 1995 (PSLRA), Section 21E of the Securities Exchange Act of 1934, 15 U.S.C. §78u-5. Verified via Cornell Legal Information Institute, law.cornell.edu/uscode/text/15/78u-5.

FASB Accounting Standards Codification

ASC 360, Property, Plant, and Equipment (impairment testing for long-lived assets held and used).

ASC 842, Leases, effective for public business entities for fiscal years beginning after Dec. 15, 2018 (per IAS Plus/Deloitte summary); private-company effective date generally fiscal years beginning after Dec. 15, 2021, per ASU 2019-10.

ASC 606, Revenue from Contracts with Customers (remaining performance obligation / RPO disclosure requirement, commonly cited as ASC 606-10-50-13).

Financial/Market Data Sources & Image/Logo Sourcing Note

Market Data & Financial Press

stockanalysis.com, company overview/statistics pages (NVDA, MSFT, PLTR, CRWV, SNOW, NOW, ORCL, ANET), accessed July 30, 2026; underlying fundamentals from SEC filings/earnings releases.

slickcharts.com, S&P 500 Total Returns by Year and Nasdaq-100 Companies by Weight, accessed July 2026.

Reuters / Deutsche Bank, 2023 Magnificent Seven return analysis; Reuters, "Meta's AI splurge lays bare its compute conundrum," July 30, 2026.

Invezz: Alphabet Q2 2026 earnings coverage (Vatsala Gaur, July 23 & 30, 2026); Microsoft Q4 earnings coverage (Ananthu C U, July 29, 2026); Oracle CLSA research-note coverage, July 20, 2026.

Fortune: Alphabet Q2 2026 coverage (July 23 & 26, 2026); "MIT report: 95% of generative AI pilots at companies are failing" (Sarah Estrada, Aug. 18, 2025); "Exclusive: Vanta hits $300 million ARR as 'shadow AI' explodes across corporate America" (Apr. 29, 2026).

Barron's, Oracle/Meta/Alphabet CDS commentary, July 29, 2026; CNBC, "Amazon Leads Big Tech's $1 Trillion Wipeout," Feb. 6, 2026.

Analyst commentary: Moody's (2026 hyperscaler capex/ROI), Citi Research (AI GDP-contribution estimate), Bank of America (AWS/Microsoft RPO analysis), JPMorgan (hyperscaler debt-issuance projection), Apollo/Torsten Sløk (bond coverage-ratio analysis), Wells Fargo/Ohsung Kwon (semiconductor sector note), D.A. Davidson/Gil Luria (analyst commentary), all as reported mid-2026 via financial-press aggregation.

Jensen Huang, Axios interview, as reported mid-2026; Andy Jassy and Satya Nadella, public earnings-call and conference commentary, 2026.

Standards Bodies & Technical Sources

Ultra Ethernet Consortium, ultraethernet.org, Specification v1.0.3, July 2026.

modelcontextprotocol.io, "Introduction" page, verified current as of July 2026.

Image & Logo Sourcing Methodology

Company "logos" shown throughout the AI Industry Value Chain section and its recap slide are rendered as brand-color name badges (company name set on that company's approximate public brand color), not reproductions of any company's actual logo mark, wordmark typeface, or trademarked graphic.

This approach was used for all 30 companies referenced across the six value-chain layers, ensuring consistent visual treatment and avoiding any use of unlicensed trademark artwork.

No other images, photographs, or AI-generated graphics are used anywhere in this deck; all visual elements are original vector shapes built directly in the presentation.

Items Requiring Final Source Verification

Every item below is preserved with its caution language intact and has not been silently confirmed as fact -- verify before citing as a specific number in a professional setting.

SLIDE / TOPIC	WHAT NEEDS VERIFICATION	CATEGORY
Rogers Figure 5-1 (Part 4)	Exact page number for Figure 5-1 in Rogers (2003) has not been independently confirmed.	*page-number*
Reinvention (Part 4, Innovation-Decision Process)	Exact Rogers (2003) page number for the reinvention definition/quotation is pending.	*page-number*
Confirmation & Discontinuance (Part 4)	Rogers (2003) page citations for replacement discontinuance and disenchantment discontinuance are pending.	*page-number*
Adopter Categories: The Standard Distribution (Part 4)	Exact Rogers (2003) page for the 2.5/13.5/34/34/16% adopter-category table is pending.	*page-number*
Geoffrey Moore's "Chasm" (Part 4)	Exact source edition and page number for the Rogers quotation disputing the chasm concept needs confirmation.	*quotation*
Industry AI-Adoption Readiness quadrant (Part 4)	Quadrant positions are illustrative; should be refreshed with current, named benchmark data before precise competitive claims are made.	*market-claim*
Magnificent Seven return contribution (Part 1)	A precise 2025/2026 statistic for the Mag7's specific contribution to S&P 500 total return has not been sourced; only the 2023 figure (Reuters/Deutsche Bank) is confirmed.	*figure*

APPENDIX · ITEMS REQUIRING FINAL SOURCE VERIFICATION (2 OF 4)

Items Requiring Final Source Verification (continued)

SLIDE / TOPIC	WHAT NEEDS VERIFICATION	CATEGORY
Energy & Physical Foundation, power-generation deals (Part 2)	Specific power-generation deal details drawn from a16z/McKinsey sourcing should be re-verified against primary sources before final use.	*market-claim*
Data-Center Construction & Cooling, Equinix (Part 2)	Equinix FY2025 10-K's specific revenue and AI-bookings figures were not independently re-verified beyond confirming the filing exists.	*figure*
Data-Center Construction & Cooling, interconnection queue (Part 2)	LBNL "Queued Up" report series is the standard source for grid-interconnection-queue data but was not independently re-verified this session.	*market-claim*
Chokepoint Layers, TSMC (Part 2)	TSMC's exact current Arizona investment total/phased timeline and precise Nvidia/Apple customer-concentration percentages (from the 20-F) were not independently re-verified.	*figure*
Chip Designers, Nvidia GPU share (Part 2)	A specific current data-center GPU market-share percentage from a named research firm (e.g., IDC, Mercury Research) was not independently verified.	*market-claim*
Memory & Networking, HBM market share (Part 2)	The exact three-way HBM market-share split between SK Hynix, Samsung, and Micron was not independently verified (TrendForce is the standard source to pull before citing a number).	*market-claim*

Items Requiring Final Source Verification (continued)

SLIDE / TOPIC	WHAT NEEDS VERIFICATION	CATEGORY
Foundation Models, OpenAI revenue (Part 2)	OpenAI's specific revenue run-rate figures beyond the cited valuation event were not independently verified; consumer subscription prices should be re-checked given how quickly they change.	*figure*
AI Platforms & Data Tools, valuations (Part 2)	Hugging Face ($4.5B, Aug. 2023) and Pinecone ($750M, Apr. 2023) valuations are 2023-dated and likely stale; refresh before citing as current.	*date*
Consulting & Systems Integration (Part 2)	Accenture's Q3 FY2026 AI-bookings figure was not yet published at the time of research; McKinsey and Bain do not publish firm-wide or AI-specific revenue figures at all and no number should be asserted for either.	*disclosure*
Historical Perspective / technology-adoption analogy (Part 2)	The 1873/1893 railroad-panic and 2002 Global Crossing/WorldCom analogy needs a specific named source (e.g., a business-history textbook or Federal Reserve historical retrospective) if a more precise citation is required.	*quotation*
GPU Obsolescence: Physical vs. Economic Life (Part 3)	No named-company GPU redeployment/write-down example is currently cited as fact; one would need to be sourced and confirmed before presenting, if required.	*market-claim*
Depreciation & Asset Impairment (Part 3)	Exact FASB ASC 360 Codification paragraph citation should be confirmed against a live FASB subscription.	*accounting-reference*

Items Requiring Final Source Verification (continued)

SLIDE / TOPIC	WHAT NEEDS VERIFICATION	CATEGORY
Leasing vs. Ownership (Part 3)	The private-company ASC 842 effective date (generally fiscal years beginning after Dec. 15, 2021, per ASU 2019-10) should be confirmed against a live FASB source.	*accounting-reference*
Common Valuation Multiples Across the AI Ecosystem (Part 3)	The entire multiples table is a point-in-time snapshot as of July 30, 2026 (~1:28 PM ET) and will be stale by presentation date; refresh every figure immediately before delivery.	*market-claim*
Earnings vs. Expectations (Part 3)	Exact share-price-move percentages (Microsoft, Oracle, CoreWeave, Alphabet) should be reconfirmed as current at time of delivery.	*figure*
10-K/10-Q Sections and Reading Backlog Correctly (Part 3)	The ASC 606-10-50-13 paragraph citation for RPO disclosure should be independently confirmed against a live FASB Codification subscription.	*accounting-reference*
Earnings-Call Comparison worked example (Part 3)	This exercise uses July 2026 results as its working example; pull the then-current quarter's figures fresh if delivered later.	*date*
Sector Pace of AI Adoption table (Part 5)	Qualitative fast/moderate/slower pace characterizations for 8 sectors are illustrative and not tied to a specific named benchmark or survey; substantiate with named source data before citing as an authoritative ranking.	*market-claim*
AI Diffusion Path Scenarios chart -- Rapid / Delayed / Limited (Part 5)	The 70%/45%/25% bar heights are illustrative shapes, not sourced forecasts or probability estimates; should not be cited as specific numeric predictions.	*market-claim*