

- Evans’ deck argues that generative AI is another huge platform shift (after PCs, the web, and smartphones), and that we are in an enormous capex and infrastructure boom without a clear, stable product-market shape yet.
- He thinks most near-term value is “absorb and automate the obvious” (coding, marketing, customer service), with longer-term disruption coming from unbundling existing businesses and giving everyone “infinite interns.”
- He is skeptical of strong, durable moats for model labs and thinks value will accumulate either in capital-intensive infrastructure (hyperscalers, Nvidia, data centers) or in “normal” software companies that wrap models in product, UX, and distribution.
- He is notably light on alignment, politics, labor dislocation, education, and broader social/ethical questions.

Included in this analysis below. (scrolling)

1. Top 10 takeaways.
2. Pros and cons of his framing.
3. Important perspectives that are largely missing.

1 Generative AI as the next big platform shift

Evans frames generative AI as the latest in a series of 10–15-year platform shifts:

- Mainframes → PCs → web → smartphones → generative AI (page 2).

Each shift:

- Redirects innovation and investment.
- Creates new gatekeepers and wipes out some old winners (e.g., Microsoft’s struggles in smartphones).
- Forces every industry to decide whether the new thing is:
 - Just a tool,
 - A revenue channel, or
 - An existential threat (page 3).

He stresses that we have seen this pattern before:

- Early periods are full of noise, hype, and counter-hype.
- Many “obvious” early winners turn out to be dead ends (AOL, WAP, J2ME, early portals, etc., page 6).
- It often takes a decade or more before the “native” products of a new platform become clear (page 7).

Conventional wisdom right now is that “AI is the new electricity / new internet.” Evans broadly agrees with the scale of importance, but emphasizes that:

- We don’t know yet what the canonical AI “device,” UX, or dominant usage pattern will be.
- We don’t know whether the central platform will be:
 - A small number of giant closed models,
 - Clouds and infra providers,
 - Open-source models, or
 - Some future, more decentralized pattern.

He uses this historical lens mainly to argue for:

- Intellectual humility.
- A long time horizon.
- The likelihood that many current narratives will look quaint or wrong in hindsight.

2 The capex super-cycle and AI FOMO

Evans argues that inside the tech industry, AI is—first and foremost—a capital expenditure story (pages 17–18, 26).

Key points:

- Big Tech capex (Microsoft, Alphabet, Amazon, Meta, etc.) is running at roughly US\$350–400B per year by 2025.
- This rivals or exceeds global telecom capex.
- Data-center construction is overtaking office construction in the US (page 19).
- Nvidia and TSMC cannot keep up with demand for advanced chips (pages 20–21).
- US power constraints and permitting processes are now major bottlenecks (page 22).
- OpenAI is announcing tens of gigawatts of future capacity and “circular revenue” deals with Nvidia, Oracle, and petrodollar-backed investors (pages 30–31).

The idea:

- Rational actors are “over-building” because under-investing in the new platform is more frightening than burning cash on infrastructure that might later be repurposed.

Pros:

- This is a sharp, concrete way of showing how seriously the biggest firms are taking AI.
- It shifts the conversation from abstract “AI will change everything” into:
 - Gigawatts of power,
 - Data-center footprints,
 - Chip fabrication,
 - Lease structures and capex cycles.
- It highlights that access to capital and energy might become a moat, not just “better models.”

Cons:

- Evans treats this mainly as “FOMO + rational hedging.”
- He does not deeply examine systemic risk:
 - What if infra truly is overbuilt relative to demand?
 - How big could the write-offs be (railroads, dot-com fiber as analogies)?
 - How would that affect innovation and risk appetite in the 2030s?

Time frame:

- The capex wave is very much a 2024–2028 story.
- The utility (or overhang) of that build-out will play out across 10–15 years, shaping what is possible with AI well into the 2030s.

3 Models are improving, proliferating, and converging – but moats are unclear

Evans stresses several interlocking uncertainties (page 11, 33–35):

- We do not know how much better LLMs can get on general reasoning tasks:
 - Some argue for AGI “soon.”
 - Others believe multiple further breakthroughs are needed.
- There are far more models now than in the early ChatGPT era:
 - OpenAI, Anthropic, Google, and other Western labs.
 - Chinese labs.
 - A fast-moving ecosystem of open-source models.
- New benchmarks, metrics, and acronyms appear weekly.
- On broad benchmarks like MMLU-Pro, leaders are closely clustered and leadership changes frequently (page 35).

This leads to one of his core strategic questions (page 42):

- If models are close substitutes on many tasks and lack strong network effects, where is value captured?

Candidates:

- Down the stack:
 - Capital, chips, power, data centers.
- Up the stack:
 - Distribution, brand, UX, vertical integration, proprietary data.

Conventional wisdom among many analysts now is:

- “Foundation models commoditize, value moves to applications and data.”

Evans is firmly in that camp.

He emphasizes that:

- Benchmark leadership is volatile.
- Customers will increasingly treat “AI models” as a menu of options, not as one sacrosanct vendor relationship.
- Price, latency, reliability, and integration may matter as much as raw IQ.

He does acknowledge some possible moats, but doesn’t dwell on them:

- Safety and alignment at scale.
- Deep specialization on high-value, domain-specific data.
- Tooling and agent ecosystems that could become de facto standards.

His net message:

- Don't assume that being "the smartest model" (in a narrow benchmark sense) is a defensible business strategy.
- Expect margin pressure and commoditization in raw inference over time.

4 Engagement gap: lots of trial, less daily use

Evans points to a gap between adoption and engagement (pages 36–39):

- ChatGPT has hundreds of millions of weekly active users, but only a small minority pay for it (about 5% is his ballpark figure, page 37).
- Consumer surveys in the US and UK show:
 - A large share of people have tried generative AI.
 - A much smaller share use it daily or even weekly.

His conclusion:

- Much of today's usage is "experimentation," not deeply embedded, habitual workflow integration.

He then asks (page 40):

- How many people have jobs that naturally fit AI into their daily routine?
- How many are in roles where they must invent new workflows to extract value?
- Do we need to wrap models in domain-specific tools so "normal" people don't have to invent their own use cases?

Pros of this framing:

- It punctures the naive idea that "everyone now lives in chatbots all day."
- It pushes companies and policymakers to distinguish between:
 - One-off curiosity,
 - Heavy professional usage,
 - Deep infrastructure embedding.

Limitations:

- Evans doesn't fully differentiate between:
 - Consumers,

- Knowledge workers,
- Developers,
- Front-line workers.
- In some segments (coding, research, content teams), daily use is already the norm.
- In many others, AI is still an occasional tool, not a core part of work identity or practice.

The bigger story is one of uneven diffusion:

- Heavy adoption in certain cognitive and creative domains.
- Slower, patchier adoption elsewhere, waiting on better integration into existing tools and workflows.

5 “Absorb” phase: AI as a feature for obvious use cases

Evans borrows a useful three-stage pattern for tech adoption (pages 50, 61):

1. Absorb
 - Use the new tech as a feature inside existing products.
 - Automate obvious processes and low-hanging fruit.
2. Innovate
 - Build genuinely new products that couldn’t exist before.
 - Experiment with new bundles and business models.
3. Disrupt
 - Rethink the problem definition.
 - Restructure entire industries around the new capabilities.

He argues that we are mostly still in the “absorb” phase:

- AI copilots for coding, marketing, and customer support.
- Summarization and search tools.
- Simple automations inside office suites (email drafting, slide building, notes).

Coding gets special emphasis:

- He calls AI coding “the new AWS” for software development (page 52).
- It acts as an abstraction layer that:
 - Reduces boilerplate.
 - Speeds up prototyping.
 - Allows more “vibe coding” (describing intent rather than hand-crafting every detail).

Conventional wisdom now is that:

- First phase of any new platform is to bolt it onto existing products.
- The truly AI-native products will appear later, when developers and users internalize what's actually possible.

Evans aligns with this, and his message to enterprises is:

- You should be in active “absorb” mode now.
- Don't wait for perfect AI-native products to fall from the sky.
- Start by upgrading existing workflows with AI capabilities where they obviously fit.

6 The Jevons paradox and “infinite interns”

Evans leans on Jevons' paradox (pages 65–66):

- When something becomes cheaper and more efficient, total usage often increases, not decreases.

Classic historical example:

- Steam power and industrial mechanization made labor vastly more productive.
- In the UK, by 1900, mechanization was equivalent to roughly 5× the human labor force in effective output (page 66).

Applied to cognitive work and AI:

- LLMs can make many forms of white-collar work cheaper and faster.
- That may not mean “same work, fewer people.”
- It may mean “more work, new kinds of work, same number (or more) of people.”

He uses the phrase:

- “Infinite interns” (page 63).

Questions he raises:

- Do you use AI to shrink headcount?
- Or do you use AI to expand the scope and ambition of what your organization tackles?
- How does this change the economics of small vs large firms?

Pros of this framing:

- It reframes AI from a narrow “job killer” narrative toward a broader “capability multiplier” story.
- It connects AI to earlier industrial transitions, emphasizing continuity as well as novelty.

Limitations:

- Jevons paradox speaks to total output, not distribution of gains.
- Evans does not spend much time on:
 - Who gets displaced and how fast.
 - Which roles or regions are most vulnerable.
 - How retraining, reskilling, or social support structures might work.

So while the economic capacity story is optimistic, the social and political story is largely left unexplored.

7 Human in the loop and persistent error rates

Evans is clear that there is no sign of model error rates going to zero (page 64).

He emphasizes:

- LLMs hallucinate and make mistakes.
- Some classes of error can be detected automatically; others require human judgment.

So design questions shift from “Can we make this perfect?” to:

- In which tasks do errors matter a lot (surgery, mission-critical control, legal filings)?
- Can we automate verification in those domains?
- Where must humans remain in the loop to:
 - Check outputs,
 - Approve actions,
 - Handle edge cases?
- How much software “wrapping” is required to make LLMs:
 - Safe enough,

- Reliable enough,
- Legally and ethically usable?

In practice, this means:

- Extensive guardrails,
- Domain-specific constraints,
- Logging and audit,
- Workflow orchestration around the model.

Evans treats this mainly as an operational and legal-risk question, not a deep philosophical or alignment issue:

- How do we design systems so that fallible models are used safely in high-stakes contexts?
- How do we support the people who must supervise them?

He does not go far into:

- What happens as systems become more agentic.
- Strategic behavior by AI agents.
- Long-term alignment with human values at very high capability levels.

But as a framing for the next 3–5 years of deployment, “human in the loop plus heavy software wrapping” is central in his analysis.

8 Adtech, recommendation, and value capture

Evans zooms in on advertising and recommendation systems as a major value pool (pages 67–71, 75):

- Global advertising spend is around US\$1T (page 68).
- A large chunk is already captured by:
 - Search (Google, etc.).
 - Social and feed-based platforms.
 - E-commerce platforms and marketplaces.

AI’s current role:

- AI already sits in the loop optimizing bids, creative, and placement for large platforms (page 69).
- Google and Meta have public statements about the uplift from AI-driven optimization.

Generative AI’s impact on creative:

- Marginal cost of creative assets drops dramatically.
- A marketer can turn a short brief into dozens or hundreds of ad variants (page 71).
- Video production, historically expensive, becomes much cheaper and more accessible (page 70).

Recommendation systems today:

- He notes that existing recommender systems depend heavily on:
 - User activity data,
 - Network effects,
 - Platform-specific engagement histories (page 67).

He then asks:

- Can LLMs recommend better with less data, because they can “understand” content more deeply?
- Can they recommend well even without a giant user base and long history of clicks?
- Could AI agents act as intermediaries between users and platforms, shifting power away from incumbent gatekeepers?

Evans also hints at a merging of:

- A system that “sees all of it” (content, products, offers).
- A system that “sees us” (our preferences, history, behavior).

Such a system could be extraordinarily powerful as both:

- A recommender, and
- A profiler.

Pros of this framing:

- Correctly identifies adtech and recommendation as enormous, near-term value pools.
- Suggests that better semantic understanding of content could:
 - Reduce dependence on pure click history,
 - Change which players capture how much value.

Limitations / missing ethics: This is one of the areas where a deeper ethical and governance perspective is conspicuously absent.

- Evans treats this primarily as an optimization and value-capture problem.
- He says relatively little about:
 - Manipulation of attention.
 - Autonomy and informed choice.
 - Impacts on democratic discourse, polarization, and civic life.

9 Most of AI strategy is still just “normal tech deployment”

Evans emphasizes that many AI “failures” are actually classic IT failures (page 56):

- Security and privacy problems.
- Data quality and integration issues.
- Poorly defined use cases.
- Vendor risk and lock-in concerns.
- Organizational resistance and change-management failures.

He cites CIO surveys showing that (page 57):

- Roughly a quarter of enterprises have at least one LLM project in production.
- A large segment plans no meaningful deployment until 2026 or later.

He also points out that even public cloud adoption is still only about 30% of total enterprise IT spend (page 58):

- It took over a decade for cloud to move from experimental pilots to core infrastructure.
- AI is unlikely to move vastly faster across the whole economy, regardless of how fast labs ship new models.

The lesson:

- Much of what organizations need to do with AI is “boring”:
 - Clean data.
 - Clarify workflows.
 - Integrate with existing systems.
 - Train staff and adjust processes.

Evans uses this to argue against both:

- Panic (“We must AI-ify everything this year”).
- Complacency (“This will blow over like a fad”).

Instead, he encourages:

- Steady, incremental deployment.
- Learning via pilots and then scaling.
- Treating AI as another major IT transition, even if the capabilities feel qualitatively different.

10 It's early, and we don't yet know what gets unbundled

The deck closes on questions rather than answers (pages 62–63, 73, 76–77).

Evans notes that:

- The internet unbundled physical retail, linear television, and the newspaper bundle.
- We didn't fully see how that would play out in 1995.

He suggests that AI could similarly unbundle existing constructs:

- We hire departments, agencies, and institutions to “do X.”
- AI may allow us to:
 - Automate parts of X,
 - Combine pieces of X and Y in new ways,
 - Or restructure the definition of X entirely.

Key questions he leaves open:

- What problems are we really trying to solve with AI?
- What do we actually care about (beyond cost savings or marginal growth)?
- Is AI:
 - Mainly a tool for CIOs and CMOs?
 - A general-purpose infrastructure like cloud and mobile?
 - A new industry that forces deep changes in institutions and social structures?

He highlights that:

- The web has been repeatedly declared “dead” or “finished” since the late 1990s (page 73).
- In reality, it continues to evolve and spawn new forms (social networks, streaming, SaaS, etc.).

His closing message is essentially:

- We should recognize both the scale of the shift and our ignorance about the end state.
- We should avoid prematurely freezing a narrative about “what AI is” or “what it will do.”

11 Pros of Evans' framing

Evans' presentation has several clear strengths:

11.1 Historically grounded, de-sensationalized

- He situates AI within a century-long story of technology waves, from steam and electrification through PCs, web, and mobile.
- This tempers both utopian hype and apocalyptic doom.
- It encourages realistic timeframes for adoption and change.

11.2 Strong on capex, infra, and power

- He brings unusual clarity to:
 - Gigawatts of power demand,
 - Data-center build-out,
 - Chip supply and fabrication,
 - Capital expenditure cycles.
- This grounds AI in hard physical and financial infrastructure, rather than pure abstraction.

11.3 Clear separation of hype, usage, and value capture

- He distinguishes between:
 - Investment scale (capex).
 - Actual user engagement and behavior.
 - Where profits and moats might emerge.
- This triad—“build → use → capture”—is a useful lens for strategy.

11.4 Emphasis on “absorb → innovate → disrupt”

- The three-stage pattern helps organizations:
 - Avoid paralysis (“too early, do nothing”).
 - Avoid magical thinking (“disrupt everything immediately”).
- It supports a pragmatic roadmap for AI:
 - Start with obvious workflow improvements.
 - Learn from pilots.
 - Then aim at genuinely new products and structures.

11.5 Integration of Jevons paradox and “infinite interns”

- Framing AI as cognitive labor expansion rather than pure displacement is economically sophisticated.
- It helps shift the conversation from “How many jobs vanish?” to “What new things become possible?”

11.6 Focus on “normal tech deployment”

- By stressing that many AI problems are classic IT problems, he helps:
 - Demystify AI.
 - Remind decision-makers that good fundamentals still matter (data, security, change management).

12 Cons / limitations of Evans’ positions

At the same time, Evans’ framing has important gaps and limitations:

12.1 Underplaying distributional and labor impacts

- Jevons paradox addresses total output, not:
 - Who benefits.
 - Who loses.
- Evans spends little time on concrete labor displacement, regional impacts, or political backlash.

12.2 Thin treatment of alignment and long-term safety

- He treats error rates and “human in the loop” as operational issues.
- He does not deeply engage with:
 - Alignment research.
 - Strategic behavior of advanced agents.
 - The possibility of misaligned objectives at higher capability levels.

12.3 Limited discussion of education and human capital

- There is almost no exploration of how:
 - Education systems must adapt.
 - AI tutors and learning tools might transform skill formation.
- This is a major omission given AI’s likely impact on knowledge work.

12.4 Narrow framing of ethics and governance

- He touches lightly, if at all, on:
 - Privacy in the era of AI agents that “see all of it, and see us.”
 - Manipulation, nudging, and autonomy in AI-driven recommendation.
 - Democratic and civic consequences of algorithmic gatekeepers.

12.5 Limited global and geopolitical perspective

- The deck is largely US/Western-centric.
- It does not deeply consider:
 - Power imbalances between countries in AI infra and data.
 - How different political systems may deploy AI differently.

12.6 Weak engagement with meaning, culture, and religion

- There is almost no consideration of how AI intersects with:
 - Human meaning-making,
 - Cultural narratives,
 - Religious and spiritual frameworks.
- For a technology that touches identity, language, and perception, this is a non-trivial gap.

13 Notable absences: perspectives that could or should have been included

Some specific perspectives that could enrich or balance Evans' analysis:

13.1 Human alignment and institutional design

- How do we align AI with human institutions that are themselves misaligned or dysfunctional?
- What does “good governance” look like when different groups cannot agree on core values?
- How do we design oversight that is:
 - Effective,
 - Legitimate,
 - Flexible in the face of rapid technical change?

13.2 Education, upskilling, and lifelong learning

- AI tutors, personalized curricula, and micro-credentialing could reshape the education landscape.
- There is very little discussion of:
 - How to prepare the workforce for AI-augmented roles.
 - New institutional forms for learning in an AI-rich environment.

13.3 Global equity and the “AI divide”

- Which countries and regions will get early access to powerful models and infra?
- How will that affect existing global inequalities?
- Where do opportunities exist for leapfrogging, as some countries did with mobile phones?

13.4 Cultural, religious, and existential dimensions

- As AI systems increasingly mediate our experience of text, image, and sound, they intersect with:
 - Stories of meaning and purpose,
 - Religious interpretation,
 - Cultural narratives of what it means to be human.
- The deck is almost silent on this plane, even though it is crucial for understanding public response and long-term integration of AI into daily life.

13.5 Long-term institution building

- Past platform shifts reshaped not only companies but institutions:
 - Universities,
 - Media,
 - Governments,
 - Civic organizations.
- A fuller analysis would at least sketch how AI might:
 - Reshape institutional trust,
 - Require new kinds of regulators,
 - Create new forms of collective intelligence and deliberation.

14 What's Next

Taking Evans' analysis as a base layer:

In the next 3–7 years, most of the action will be “absorb” and “infinite interns”:

- AI tools woven into everyday workflows in education and health.
- Big gains in documentation, triage, translation, and scaffolding learning.

- The “capex and infra” story suggests:
 - Powerful AI will be widely available as a utility.
 - The hard work will lie in designing humane, aligned, context-sensitive applications—exactly the zone where PSA’s concerns live.

- The “engagement gap” implies:
 - Many people will not adopt AI deeply until it is embedded in familiar tools.
 - That creates an opening for projects that help communities learn how to use AI wisely, not just automatically.

- The absence of alignment, ethical, and spiritual perspectives in Evans’ deck:
 - Creates a space for your work (and our dialogue) to address the missing layer:
 - * How AI reshapes meaning, story, and consciousness.
 - * How to align AI with human flourishing when humans are not in agreement about what flourishing is.

In short, Evans provides a strong economic and infrastructural scaffold for understanding the AI transition. The document can build on that scaffold while explicitly adding:

- Human, ethical, spiritual, and educational dimensions.
- Concrete, grounded examples in healthtech and edtech for underserved communities.
- A longer-horizon conversation about AI, meaning, and the evolution of human consciousness that his deck leaves almost entirely untouched.