The Capability Gap: Why Both the Doomers and Boomers Are Wrong About AI

By Ashok Murthy

There is a war being fought in headlines, investor memos, and boardrooms — and it has nothing to do with actual artificial intelligence. It is a war of narratives. On one side stand the Doomers, convinced that AI will displace the workforce in a catastrophic spiral of unemployment and economic collapse. On the other stand the Boomers (not the generation — the optimists), who believe AI will unlock instant productivity gains and usher in a golden age of abundance overnight. Both camps are wrong. And the gap between what they believe and what is actually happening is precisely where the greatest economic opportunity of the next decade lives.

The Narrative of Doom: A $100 Billion Panic

To understand why fear dominates the AI conversation, consider a 2024 scenario published by research firm Catrini. In a fictional 2028 memo titled "The Global Intelligence Crisis," analysts described a world where AI-driven white-collar displacement triggers a financial contagion on par with the 2008 housing collapse. The mere publication of this hypothetical memo caused IBM's stock to crash 13% — its worst single day in 25 years.

The Catrini memo drew an uncomfortable parallel: *"In 2008, the loans were bad on day one. In 2028, the loans were good on day one… the world just changed after they were written."* The implication was chilling — that AI could retroactively invalidate the economic assumptions underlying trillions of dollars of financial products and employment contracts.

The logic of the "Doom" model follows what you might call the **Intelligence Displacement Spiral**: AI capability spikes → consumption collapses → credit contagion spreads → systemic failure → repeat. The model claims that a 2% decline in white-collar employment could trigger a 4% hit on discretionary spending — a multiplier effect that would cascade through the entire economy.

It is a compelling narrative. And that is precisely the problem.

Why We Buy the Panic: The Biology of Bad News

The Doom narrative is not viral because it is inevitable. It is viral because of evolutionary biology. Research consistently shows that a headline predicting economic catastrophe generates 10 to 50 times more engagement than one predicting gradual improvement. This is not irrationality — it is the product of millions of years of survival wiring. Our brains are built to prioritize threats over opportunities.

This negativity bias distorts our perception of technological transitions. The same cognitive hardware that kept our ancestors alive on the savanna now causes investors to price in catastrophic outcomes that may be statistically improbable, while systematically underpricing slow-moving, structural opportunities. The result: markets routinely misprice assets based on memes rather than mechanisms.

Understanding this is not just intellectually satisfying — it is strategically useful. If fear consistently drives capital away from a sector, and that fear is biologically exaggerated rather than analytically grounded, then the contrarian position may be the structurally correct one.

The Economics of Abundance: Why the Doom Model Fails

University of Chicago economist Alex Emis has argued that the Doom model relies on an "impossible confluence" of conditions: zero policy response, zero price elasticity, and fully passive capital owners. In other words, the catastrophic scenario requires governments to do nothing, consumers to stop adapting, and businesses to sit idle as their capital depreciates. None of these conditions have ever held simultaneously in any prior technological disruption.

History suggests the opposite dynamic — what economists call the **Jevons Paradox**: when technology makes goods and services cheaper, consumption historically *increases*. Steam engines made coal extraction cheaper; total coal consumption surged. More efficient cars led to more driving, not less. When AI reduces the cost of legal analysis, tax preparation, or insurance underwriting, the likely outcome is not mass unemployment in those sectors — it is an explosion of demand for those services at lower price points.

The **Service Sector "Raise"** model makes this concrete. If AI efficiency reduces household service costs by 40-70% across categories like tax preparation, insurance, travel planning, and mortgage processing, the average American household could save approximately $7,000 per year. That money does not disappear — it recycles into the real economy, funding durable goods, home improvement, and local services. The displacement in one sector becomes the stimulus in another.

This is not a guarantee. It is a countervailing force that the Doom model structurally ignores.

The Great Disconnect: The Real Story Nobody Is Telling

If the Doom narrative is wrong, does that mean the Boom narrative is right? Not at all. The Boom narrative makes its own critical error: it assumes that because AI *can* do something, AI *will* do it — quickly, universally, and at scale.

This is the **Fallacy of Speed**.

The chart that matters most in understanding the current AI moment is not a capability curve — it is the gap between two curves. AI capability (reasoning, language, analysis) has been growing at something approaching an exponential rate. Societal dissipation — the slow process of adoption, regulation, trust-building, and institutional change — grows at something closer to a linear rate. The gap between these two curves is **The Great Disconnect**, and it is where we actually live today.

The Doomers assume instant labor displacement. The Boomers assume instant technical adoption. Both are blind to the same thing: **social inertia is the most powerful economic force in the modern economy, and it is absent from almost every financial model.**

The Force of Friction: Four Types of Social Inertia

To understand why the Capability Gap persists — and why it will persist for years — you need to understand the four distinct types of social inertia that slow AI adoption from raw capability to real-world impact.

Think of it as a funnel. At the top sits everything AI *can* do. By the time you reach the bottom — what a firm or individual *actually does* with AI — most of the potential has been lost to friction.

1. Regulatory Inertia In healthcare, every AI application must navigate HIPAA compliance and FDA clearance cycles. In government contracting — a sector many of us know well — procurement cycles are measured in years, not quarters. Legal review alone can consume more calendar time than the technical development of the tool being reviewed.

2. Organizational Inertia The Cobalt Problem: 95% of ATM transactions globally still run on COBOL — a programming language first developed in 1959. Legacy infrastructure is not just a technical obstacle; it is a political one. Unions, severance laws, HR policies, and middle management incentive structures all create drag on AI deployment that no capability benchmark can measure.

3. Cultural Inertia Most people do not use AI daily. This is not because AI lacks utility — it is because behavioral change requires consistent reinforcement, social proof, and often institutional mandate. Shopify's approach of requiring employees to demonstrate why AI *cannot* do a task before escalating to a human is a model of forced cultural adoption. Without that kind of deliberate mandate, adoption remains voluntary and therefore slow.

4. Trust Inertia Perhaps the most underappreciated friction: verification is expensive. AI generation costs are low and getting lower. But the cost of verifying AI output — particularly in high-stakes domains like legal analysis, financial modeling, or defense contracting — remains high. Enterprises cannot safely trust "black box" outputs without audit trails, explainability frameworks, and accountability structures. Until verification costs fall, deployment will remain constrained regardless of capability.

The Inertia Penalty — and the Asymmetric Opportunity

Here is the strategic insight that follows from understanding these friction types: **large firms and small firms are not equally disadvantaged by social inertia — they are disadvantaged in opposite directions.**

Large firms have high capital, high data assets, and high distribution. But they face **Critical Inertia** — an average of 18 months between recognizing that a tool saves money and actually deploying it at scale. Legal review, security review, change management, training, and procurement cycles all compound.

Small firms and individuals have low capital and limited data. But they face **Zero Inertia**. A solo operator or small team can evaluate a new AI model the day it drops, build a personal evaluation framework within 24 hours, and begin extracting value while the enterprise competitor is still in quarterly planning.

This asymmetry — speed as the only remaining asymmetric advantage — is the defining competitive dynamic of the current AI transition. The window during which individuals and small firms can move faster than large institutions is not permanent. But it is real, and it is open right now.

The Rise of the One-Person Giant

The practical expression of zero inertia is what we might call the **One-Person Giant**: an individual or micro-team that leverages AI to operate with the functional capabilities of a much larger organization — code, design, legal, financial analysis, operations — with near-zero overhead.

The data supports this thesis. New business applications in the United States reached 532,000 in January 2026 — a level that dwarfs every prior year in the dataset. These are not all AI-native businesses, but the correlation between AI tool availability and solo/micro-business formation is not coincidental. When a single individual can generate professional-grade deliverables across multiple disciplines in hours rather than weeks, the economics of small-scale entrepreneurship fundamentally change.

As the observation goes: "One-person businesses have more leverage in the economy than ever before." This is not a motivational slogan. It is a structural description of what happens when the cost of capability drops to near zero while the cost of trust and verification — which still requires a human — remains high.

The 'Toby Lütke' Playbook: How to Position Yourself in the Gap

Shopify CEO Toby Lütke has articulated one of the most operationally useful frameworks for navigating the Capability Gap. It has three components:

1. Mandated Friction: Before asking a human colleague for help, demonstrate why AI *cannot* do the task. This inverts the default assumption from "humans first, AI as backup" to "AI first, humans for genuine exceptions." Over time, this forces teams to build real fluency with AI tools rather than treating them as optional accessories.

2. The 'Toby Eval': Maintain a personal library of prompts and task specifications, and run them against every major new model release. This is not about novelty — it is about building a systematic, empirical understanding of where AI capability has moved and what new workflows have become viable. Most organizations do not do this. The individuals who do compound their advantage with every model generation.

3. Instrument, Not Replacement: The frame that prevents both over-reliance and under-utilization. AI is the chessboard. You are the grandmaster. The question is never "will AI replace me?" — it is "how do I become a better player with a better board?" Player plus instrument consistently outperforms instrument alone.

Collapse the Integration Timeline

The final practical implication is the most actionable: **do not wait for institutional permission to validate AI tools.**

Standard corporate integration timelines span 18 months from awareness to deployment — moving through quarterly planning, legal review, pilot, and finally deployment. The agile

individual collapses this to 24 hours: new model drops, personal evaluation runs, validation confirms utility, use begins.

The competitive advantage of this approach is not just the 18-month head start. It is the compounding eval framework — the growing library of personal benchmarks, workflow integrations, and capability maps — that accumulates with every cycle. By the time an enterprise deploys the tool you started using 18 months ago, you have already iterated through three subsequent generations.

The instruction here is simple: **Be the person who validates the tool while competitors are in quarterly planning. Fail fast to build the eval framework. Do not plan your career around the Crash or the Boom. Plan it around the Gap.**

The Generational Opportunity

The Capability Gap — the space between what AI can do and what society has figured out how to do with it — is not a problem to be solved. It is an opportunity to be inhabited.

Economic value, over the next decade, will not accrue primarily to the companies that build AI capabilities. Nor will it accrue uniformly to those who are "good at AI." It will concentrate in the hands of the people and organizations who become **bridges** — who can translate raw capability into trusted, compliant, human-integrated workflows that institutions are willing to adopt.

This is a positioning argument, not a technical one. The technical capability is available to anyone. The friction — regulatory, organizational, cultural, trust-based — is the moat. Those who learn to work within the friction, to reduce it strategically, and to build the institutional trust that enables deployment are the ones who will define the economics of this transition.

The Doomers are planning for a world that is unlikely to arrive as fast as they fear. The Boomers are planning for a world that is unlikely to arrive as fast as they hope. The opportunity is in the gap between the two — and it is available right now, to anyone willing to work within the reality of how organizations and societies actually change.

Be the bridge.

This article synthesizes research from NotebookLM-sourced analysis on AI adoption dynamics, social inertia frameworks, and the economics of technological transitions.