


From AI table stakes to AI advantage: Building competitive moats

When everyone has access to the same AI models, the winners are those who use them to build advantages that competitors can't copy.

by Dago Diedrich, Evan Williams, Tanguy Catlin, and Tim Fontaine



“If everyone is special, then no one is.” That line, adapted from the movie *The Incredibles*, captures the essence of a key issue with AI today. While AI adoption has exploded (nearly nine in ten organizations now use AI in at least one business function, most companies are deploying the same large language models (LLMs) to improve productivity. If everyone has the same advantage, it’s not really an advantage.

This is a trap we’ve seen before. During the digital-transformation wave, companies rushed to develop websites and apps, but competitive advantage—the distinct set of hard-to-replicate assets and operating models a company creates that earns superior returns over time—didn’t automatically follow. Our analysis of the banking sector, for example, showed that companies increased mobile-app adoption between 2018 and 2022, but leaders didn’t extend their advantage over laggards. What widened the gap was integrating digital and AI into the entire customer journey to reduce friction, which fueled online sales and contributed to significant TSR outperformance.¹

The lesson: Apps and tools can be copied. Value comes from building advantages that are hard for competitors to replicate—that is, competitive moats. That’s a critical lesson to bear in mind as CEOs and boards consider how to capture their fair share of the [enormous value just from generative AI](#) that’s at stake.

To better focus that thinking, we identified moats—six strategies and three capabilities—that can provide a competitive advantage. These moats shouldn’t be particularly new to business leaders, but AI has shifted the dynamics in each of them. Understanding what that means is the path to leaping from AI as table stakes to AI as an abiding competitive advantage.

Strategic moats

While a strategic moat is difficult for others to replicate, it should also represent an area where the business already has an advantage. This view can help CEOs determine which area to focus on and how to invest.

Economies of scale: Infrastructure to harness speed and scale

AI’s most fundamental economic effect in many sectors is to collapse marginal costs, starting with industries with a high proportion of costs associated with cognitive work. A customer service interaction that might have cost \$15 in labor now costs a fraction of that in compute—and is poised to continue dropping. Unit-cost advantage goes to firms that can radically lower the cost per customer served.

Even as start-up costs for some technology elements (compute, storage) are declining and the economics of AI are in flux, the costs for scaling value will require a broader set of technology capabilities. The strategic moat comes from turning cognitive work into infrastructure—data pipelines, fine-tuned models, integrated workflows, governance layers—that can scale at very low incremental cost. Benefits further accrue through feedback loops and margin gains that can be reinvested in better models, broader distribution, or acquisitions.

¹Eric Lamarre et al., “The value of digital transformation,” *Harvard Business Review*, July 31, 2023.

This effect is particularly powerful in service industries, where AI is replacing what were once variable labor costs. While new entrants or smaller competitors have lower barriers to entry, given relatively cheap access to LLMs and few legacy tech issues, the infrastructure barriers persist.

Resolution Life, a US and Australian life insurer, demonstrates what this moat looks like. Its AI platform automates actuarial, marketing, and financial tasks, enabling the company to introduce additional insurance products at a fraction of the historical cost. The platform also triages incoming claims in 15 seconds rather than weeks, allowing assessors to focus on complex cases as volumes grow.² Once models and integrations are in place, the same stack supports additional books of business acquired through M&A, pushing more volume through fixed infrastructure.

What this means for you: If cognitive work represents high costs in your business, focus on what you need in place to build the infrastructure so economies of scale from AI work in your favor. Senior leaders should consider consolidating volume across business units or geographies, building shared AI platforms, or using M&A to push additional volume through the same AI stack.

Key CEO questions: Which transaction costs and sources of friction is AI collapsing in our industry? What are the implications for marginal costs, scale economies, and market structure? Will there be a race to scale with our competitors?

Privileged data: Treating data like an asset class

Privileged data becomes a moat when AI models use it to deliver products and services that competitors can't, such as more accurate recommendations, improved risk scores, or better-performing tools. When properly architected, every AI interaction generates more labeled behavioral and outcome data that feeds back into training the model, creating a data flywheel. The most valuable data sets are cumulative and protected, such as historical transaction or telemetry data.

Amazon illustrates how privileged data becomes a moat when generated within a closed-loop ecosystem. Across its retail and marketplace businesses, Amazon captures proprietary signals on search behavior, product views, purchases, fulfillment, and advertising response at an enormous scale. That data improves recommendations, demand forecasting, ad targeting, and marketplace optimization in ways that competitors without comparable behavioral and transaction data struggle to replicate. The economic value of that advantage is visible in Amazon's advertising business, which reached \$68 billion in revenue in 2025, showing how proprietary commerce data can be turned into a compounding, high-value asset.³

What this means for you: Manage data as a strategic asset class. That starts with prioritizing the data that underpins your differentiation and instrumenting systems to capture, enrich, and

² Kate Weber, "Resolution Life uses AI to triage claims in 15 seconds," ITnews, November 25, 2022.

³ Peter Adams, "Amazon's annual ad revenue passes \$68B, boosted by full-funnel strategy," Marketing Dive, February 6, 2026.

The strategic moat comes from turning cognitive work into infrastructure—the data pipelines, fine-tuned models, integrated workflows, governance layers—that can scale at very low incremental cost.

maintain this data at scale. Demonstrate responsible data stewardship, such as stringent data protection measures to avoid future regulatory constraints.

Key CEO questions: What unique data (including unstructured data we couldn't use before) could we capture, label, or generate (for example, through customer usage that would compound model performance to generate value faster than rivals can match?

What proprietary data could a competitor develop that would put them at an advantage over us? How would we respond?

Embeddedness: Making switching expensive

AI solutions go from convenience to necessity when embedded in core work, shaping how work gets done to deliver a better product or service at a competitive price. This is evident when AI agents orchestrate supply chains, drive sales and service flows, or generate clinical documentation inside electronic health records. Replacing them would require significant effort and cost, including rebuilding integrations, redesigning workflows, retraining employees, and potentially sacrificing accumulated performance gains.

Three “layers of embeddedness” go into creating this moat, and they compound over time:

1. Capabilities are integrated directly into core and complex systems (for example, customer relationship management [CRM], enterprise resource planning [ERP], productivity suites, and industry platforms), making workflow migration expensive.
2. These systems learn from proprietary data through feedback loops, becoming increasingly valuable over time.

3. Employees become comfortable with the solution and rely on it. Changing introduces friction, temporary productivity losses, and organizational resistance.

While increasingly sophisticated agents will shift what qualifies as embedded (for example, by abstracting away the traditional user interface and creating intuitive interfaces), those agentic capabilities are likely to create their own levels of embeddedness.

Microsoft Dragon Copilot (formerly Nuance DAX Copilot) is a good example here. The tool not only captures ambient audio from patient encounters and drafts clinical notes directly in the electronic health record (EHR) but also is deeply integrated into the EHR infrastructure and workflows at multiple healthcare providers. For example, deployments across more than 150 hospitals with Epic EHR software have shown a roughly 50 percent reduction in documentation time, as well as material reductions in clinician burnout.⁴ Some sites report doctors seeing several additional patients daily once Microsoft Dragon Copilot is in place.

What this means for you: For vendors, understand where the points of deep workflow integration are, and ensure services have feedback loops so performance improves and value increases with use. For customers, every workflow you hand over to an embedded AI system is a bet on the vendor's road map, pricing trajectory, and continued existence. Negotiate data rights and portability up front, ensure you retain access to that learning in some usable form, and maintain data protection standards.

Key CEO questions: How could we embed our product or service and make it so valuable that customers won't want to turn it off? Where are we exposed to the risk that our vendors will become so embedded that we will lose important degrees of freedom in the future?

Network effects: AI as the network architect

Networks can be powerful moats when every new user or interaction makes the platform more valuable for the next. AI reshapes these dynamics by making networks easier to build and more valuable as they scale, creating an added impetus to develop and extend them at speed.

First, AI reduces the "cold start" problem that has historically limited the growth of new networks. It can generate initial supply (such as listings, content, or product descriptions) and compress time to value for new users by personalizing recommendations from the outset. Automating content creation and interaction tools can further help users contribute and consume network content.

Second, AI strengthens network effects by improving how participants are matched and how value is surfaced. As networks grow, AI ranks, recommends, verifies, and filters content, ensuring that the right participants find each other and that high-quality interactions rise above the noise. Credit card networks illustrate this dynamic: AI matches offers to the customers most likely to redeem them and measures performance precisely, increasing engagement and attracting more partners. The network deepens not just because it is larger, but because it becomes more intelligent.

⁴ Andrea Fox, "Nuance AI copilot now fully embedded in Epic EHR," Healthcare IT News, January 19, 2024.

AI solutions go from convenience to necessity when they are embedded in core work and shape how work gets done to deliver a better product or service at a competitive price.

TikTok demonstrates how AI can accelerate engagement. Its recommendation engine ensures users see relevant content immediately.

As AI agents increasingly mediate network interactions (for instance, by searching, comparing, and transacting on behalf of users), discovery shifts from human browsing to algorithmic decision-making. Agentic commerce, for example, could orchestrate up to \$1 trillion in US retail by 2030, according to estimates.⁵ Advantage will shift to platforms that learn to play effectively in this agentic world, either by owning the agentic layer or by being positioned to be selected by agents.

What this means for you: Look for opportunities to create network effects that you previously dismissed as too costly or risky. If you have an existing network, ensure your models improve matching quality, reduce noise, and increase trust with every transaction. As transactions increasingly flow through AI agents, be clear about who owns the agent and who captures value when agents transact.

Key CEO questions: Where are the transformational growth opportunities for AI by building new networks or improving the quality of existing networks? How will we compete in markets and ecosystems increasingly mediated by agents?

Business model disruption: Shifting who owns the customer and how value gets priced
Business model disruption becomes a competitive moat typically when incumbents face channel conflict, economic constraints, or organizational barriers. AI accelerates this kind of disruption along two primary vectors:

⁵Katharina Schumacher, Roger Roberts, and Katharina Giebel, "[The agentic commerce opportunity: How AI agents are ushering in a new era for consumers and merchants](#)," McKinsey, October 17, 2025.

- The customer relationship. With the “agentification” of business, we are entering a new phase of disintermediation as AI agents capture customer relationships. In brokered markets, such as commercial and life insurance (where third-party distributors account for more than 85 percent and 50 percent of sales, respectively⁶) or mortgages (where about 75 percent are broker-originated in Australia), AI agents can replicate the functions of brokers and can make direct sales.

NEXT Insurance, for example, built an AI-powered platform that lets a café owner or electrician get a quote, buy coverage, and download a certificate in about ten minutes, all without a human broker. More than 600,000 entrepreneurs now use the service, and a survey of 1,500 small-business owners found that about 60 percent already buy insurance fully online without an agent.⁷

This disintermediation is also happening as ChatGPT and Claude integrate third-party apps. When a user books a service through the agentic interface, the AI captures the relationship and the data, while the supplier becomes interchangeable. The AI agent also accumulates cross-app behavioral data, building new habits and creating switching costs for users that can shape demand across the ecosystem.

- *Outcome-based pricing.* While outcome-based pricing is not a new concept, AI makes it more feasible. AI can better measure outcomes in real time, scaling the ability to make continuous optimizations and improving the predictive pricing accuracy.

Michelin's EFFIFUEL, for example, guarantees fuel savings over multiyear commitments. It uses real-world data analytics and telematics to track fuel use against contractual savings targets.⁸ This creates long-term contractual lock-in, reinforcing the data and embeddedness moats, making it hard for competitors to dislodge.

Service industries whose economics and incentives are built on billable hours or day rates are vulnerable to this pricing disruption because they face internal resistance to cannibalizing their legacy model. AI-first competitors that have offers designed around outcomes can provide more attractive terms.

What this means for you: If you're in a brokered market, identify where AI allows you to connect with the customer and assess channel conflict issues. If you're at risk of being disintermediated, identify what your economic leverage points are to protect and enhance customer relationships. If you bill by time or throughput, identify how to switch to outcome-based pricing.

Key CEO questions: Where could we use AI to disintermediate brokers, own customer relationships directly, or charge for outcomes rather than inputs? What business model could a well-funded attacker launch against us to take advantage of friction and cost arbitrage in our product and service offerings?

⁶Insurance Forums Staff, “Independent distribution drives individual life premium to sales record in 2024,” Insurance Forums, March 27, 2025.

⁷Tim McDermott, “Survey: 60% of small business owners buy business insurance without an agent or broker,” ERGO NEXT, March 20, 2025.

⁸“Michelin Solutions launches EFFIFUEL—A commitment to helping truck fleets save fuel,” Cision, July 11, 2013.

Constrained assets: AI meets the physical world

As AI drives down the marginal cost of digital intelligence, competitive dynamics will shift more toward the physical world. Companies that control and optimize the real-world bottlenecks will hold the strongest moats. Large-scale physical assets—mines, ports, airports, power generation, data center sites, distribution networks—already have effective moats.

For companies whose physical assets are less dominant, however, the strategic moat emerges when they combine control of new or existing physical assets with AI-driven optimization. Firms that integrate AI into a physical asset that is already a competitive advantage (such as supply logistics, medical equipment, or field operations) can enhance that advantage, including by creating new proprietary data sets that competitors cannot quickly replicate. Reproducing it requires not just software but also capital investment, operational expertise, workflow redesign, and, often, regulatory approval.

Amazon's fulfillment network illustrates this dynamic. Its moat is not just software, but a vast, strategically located warehousing and logistics infrastructure enhanced by AI for inventory placement, routing, and robotics.

A similar logic applies to human-AI production systems in physical workflows. Many activities, from wealth management to manufacturing, still require human judgment, physical presence, or regulatory accountability.

John Deere shows how AI can create a moat when it is tied to hard-to-replicate physical assets. Its See & Spray technology uses computer vision and machine learning on field equipment to identify weeds and apply herbicide selectively in real time. The advantage comes not only from the model but also from John Deere's installed equipment base, agronomic data, dealer network, and integration into real farm operations. Competitors would need to replicate not just software but also machines in the field, service capabilities, customer relationships, and years of operational learning.⁹

What this means for you: As AI commoditizes knowledge, focus on where your company controls the physical systems that competitors cannot easily replicate. The goal is not simply to apply AI to existing assets, but to build physical networks whose value compounds when combined with intelligence.

Key CEO questions: Which scarce physical, operational, intangible, or regulated assets in our business become more valuable as AI makes digital intelligence cheap and abundant? Where could combining AI with hard-to-replicate assets (infrastructure, field force, supply chain, installed base, or licenses) create an advantage that software-only competitors cannot match?

⁹ Katy Hill, "More farmers are adopting John Deere's See & Spray. Here's why," CropLife, May 1, 2025.

As AI drives the marginal cost of digital intelligence down, competitive dynamics will shift more to the physical world.

Capability moats

A capability moat is an organizational strength that is difficult to build but enables a company to repeatedly translate AI into sustainable advantage.

Rewiring for velocity: Increasing speed of learning and deployment

Because AI performance improves with experimentation and data, organizations with a learning and development velocity that is consistently greater than their peers can create competitive moats.

Our research has demonstrated that companies in the top quartile of software development velocity achieve four to five times faster revenue growth and 60 percent higher total shareholder returns than bottom-quartile peers.¹⁰ Agentic AI is already showing signs of significantly accelerating every stage of the development and deployment cycle, moving from two-week to 24-hour development cycles (or sprints).

Building up a high-velocity learning and execution advantage requires businesses to [rewire how they work](#). This is much more than training developers how to use agentic tools. It's about building up a coherent set of capabilities: a strategy that targets transformational value and reimagines end-to-end workflows; top-tier tech talent; an operating model built around small, rapidly moving cross-functional teams; flexible technology platforms; data embedded in the business; and systems that reuse to drive adoption and scaling.

Rewired companies typically improve their EBITDA by 10 to 30 percent, with an average of 20 percent, according to recent McKinsey analysis.¹¹ These capabilities create a flywheel effect, leading to the gap between leaders and laggards widening by roughly 60 percent in recent years.¹²

DBS Bank illustrates how rewiring the business increases velocity. Over the past decade, DBS shifted to a “managing through journeys” operating model of cross-functional squads, cleaned its data and made it accessible, and built an AI platform to standardize models for easy reuse.

¹⁰Developer velocity is the software development capability made up of empowering developers, creating the right environment for them to innovate, and removing points of friction; Shivam Srivastava, Kartik Trehan, Dilip Wagle, and Jane Wang, “[Developer Velocity: How software excellence fuels business performance](#),” McKinsey, April 20, 2020.

¹¹Eric Lamarre, Kate Smaje, Robert Levin, Alex Singla, and Alexander Sukharevsky, [Rewired: How Leading Companies Win with Technology and AI](#), Wiley, April 2026.

¹²“[Rewired and running ahead: Digital and AI leaders are leaving the rest behind](#),” McKinsey, January 12, 2024.

These efforts have been part of a broader rewiring effort that has helped them reduce the time it takes to develop and deploy AI solutions to two to three months, from 12 to 18 months.¹³

What this means for you: Treat organizational velocity as a strategic differentiator, not an operational metric. Measure your clock speed from idea to proven value to scaled deployment—and remove the bottlenecks that slow it down. Companies looking to truly rewire their organization have to start by building conviction across the entire C-suite, target domains where they have economic leverage, and commit both real resources and their top people to lead the transformation.

Key CEO questions: What kinds of decision bottlenecks and deployment friction could you eliminate to run orders of magnitude more experiments per year? What actions will have the biggest impact on reducing learning cycle times? What would be the impact of a rival that learned five times faster than we did?

Regulation and compliance: Integration into AI solutions

As AI expands the use of sensitive data and automated decision-making, regulatory scrutiny is intensifying across markets. The EU AI Act has provisions on copyright protection, security, and transparency regarding the use of data and content created by AI models.

The strategic moat develops when regulatory compliance is embedded in the solution development process and the technology stack: built-in audit trails, explainability, data lineage tracking, bias monitoring, and human-in-the-loop controls. If a company has regulatory permission (for example, as Waymo has for self-driving cars in some places) or patents (for instance, as some pharmaceutical companies have with their glucagon-like peptide-1s [GLP-1s]), they have a window to profit while competitors are finding ways over the hurdle. Building these capabilities requires legal expertise, risk infrastructure, governance processes, and capital buffers—assets that incumbents in regulated industries often already possess.

While attackers could use LLMs to navigate regulations and compliance requirements, or take advantage of regulatory “gray zones,” over time, enforcement tends to catch up. When it does, the advantage shifts toward firms that have already built scalable compliance infrastructure.

What this means for you: Treat compliance as strategic infrastructure by embedding auditability, transparency, and governance directly into your AI systems. This helps regulatory compliance scale with your growth. Clarify which parts of your value proposition depend on regulated activities and where you may face gray-zone competition.

Key CEO questions: Where in our market will regulatory gray zones create temporary openings for attackers? How could we turn compliance, auditability, and governance into product features that make us easier to adopt or more trusted than less mature competitors?

¹³ Eric Lamarre, Kate Smaje, Robert Levin, Alex Singla, and Alexander Sukharevsky, [Rewired: How Leading Companies Win with Technology and AI](#), Wiley, April 2026.

Because AI performance improves with experimentation and data, organizations with a learning and development velocity that is consistently greater than their peers can create competitive moats.

Trust: Your anchor for the customer relationship

In high-stakes domains such as finance, healthcare, and identity, trust is a strategic moat because it functions as a gatekeeper to adoption. In banking, for example, more than half of consumers now use gen AI tools and want banks to offer them as well. Nearly all say they would eventually switch to another provider if their current bank didn't keep up with this technological shift.¹⁴

Customers are more willing to share data and accept automated decision-making when they believe a firm is safe, fair, and accountable. That access, in turn, improves model performance and deepens switching costs.

AI, however, has unique trust challenges. Global surveys indicate that only about 30 percent of people globally embrace AI, while 35 percent reject it.¹⁵ Building that trust comes from having explainable AI models, explicit consent, auditability, and clear guardrails. Organizations that embed responsible AI into how they design, deploy, and monitor systems gain permission to scale faster. High AI performers are more likely to have formal validation processes, and sufficient investment in responsible AI correlate with capturing material EBIT impact.¹⁶

JPMorgan Chase has been ranked first on the Evident AI Banking Index for four consecutive years.¹⁷ It's one of only a few banks publicly reporting realized AI returns, approaching \$2 billion.¹⁸ In tightly regulated markets, that combination of performance and openness underpins both regulatory trust and investor confidence.

What this means for you: Winning companies treat trust as a speed enabler. Identify the core sources of trust in your businesses (safety, fairness, reliability, transparency) and embed them directly into your AI systems through automated governance and policy-as-code controls. Integrate risk and compliance guardrails early in AI solution development.

¹⁴ ["Global Banking Annual Review 2025: Why precision, not heft, defines the future of banking."](#) McKinsey, October 23, 2025.

¹⁵ "Technology's tipping point: Why now is the time to earn trust in AI," World Economic Forum, March 21, 2024.

¹⁶ ["Moving past gen AI's honeymoon phase: Seven hard truths for CIOs to get from pilot to scale,"](#) McKinsey, May 13, 2024; Gabriel Morgan Asaftei, Roger Roberts, Abby Sticha, and Cécile Prinsen, ["State of AI trust in 2026: Shifting to the agentic era,"](#) McKinsey, March 25, 2026.

¹⁷ Evident AI Index: 2025 edition, Evident Insights, accessed April 2026.

¹⁸ Lisa Monica and Ramdhani Pratama, "AI brings huge efficiency gains, saving JPMorgan US\$2 billion per year," IDN Financials, October 29, 2025.

Find more content like this on the
McKinsey Insights App



Scan • Download • Personalize



In high-stakes domains such as finance, healthcare, and identity, trust is a strategic moat because it functions as a gatekeeper to adoption.

Key CEO questions: How could demonstrable AI safety, fairness, and privacy become an asset that unlocks integration into our customers' and/or partners' most high-value processes? In which decisions would customers, employees, regulators, or partners trust us with AI before they would trust anyone else—and why?

Turning models into moats: Early actions

The forces unleashed by AI have shifted the locus of competitive advantage. For boards and CEOs, this suggests a clear agenda:

- *Align on your moats—and make trade-offs explicit.* The firms pulling ahead aren't necessarily building all nine moats, but they rarely rely on just one. Many of these moats reinforce one another. Privileged data strengthens business model innovation. Trust unlocks deeper data access. Identify the one to three where you have a structural advantage, then align and commit to them explicitly. This requires committed investment and making hard trade-offs.
- *Build the enabler systems that support your strategic moat.* Define the enabler system that will reinforce your strategic moat over time: how data is captured, how models improve, how solutions scale, and how quickly you can iterate. Prioritize which workflows and set of capabilities are needed to power the relevant end-to-end workflows.
- *Govern the moat like a core business, not a set of experiments.* Moat building is a multiyear program. Boards and executive teams should track a small number of leading indicators tied directly to the chosen moat—such as cost per transaction, network liquidity, or share of customer interactions mediated by AI (embeddedness). This ensures sustained focus, disciplined resource allocation, and accountability for outcomes.

In the age of AI, competitive advantage won't come from having the cleverest model. It will come from being the organization that turns common models into uncommon moats faster than anyone else.

Dago Diedrich is a senior partner in McKinsey's Düsseldorf office; **Evan Williams** is a partner in the Sydney office, where **Tim Fountaine** is a senior partner; and **Tanguy Catlin** is a senior partner in the Boston office.

This article was edited by Barr Seitz, an editorial director in the New York office.

Copyright © 2026 McKinsey & Company. All rights reserved.