



The Second Mover Advantage

Why Now is the Time to Invest in AI

July 2026

By Dr. Andrej Levin, Dr. Ivan Tretiakov and Dr. Robin Wagner



Where Strategic Clarity Meets Applied AI

We are navigating an era of unprecedented change and disruption—powered by technology, marked by complexity, where change amplifies at scale. To lead, companies need a partner that can bridge the gap between ambition and outcomes. BCG is built for this moment. We bring strategic clarity, rooted in over 60 years of deep domain knowledge, to ensure leaders make the right choices. We combine it with applied AI, shaped and wielded by our practitioners, teaming shoulder-to-shoulder with your teams to deliver transformative impact at scale. The result? Stronger returns, transferred capabilities, and change that sticks. We are BCG.

Table of Contents

1	The Race for AI Advantage Has Entered Its Decisive Phase	<u>04</u>
2	Three Curves, One Decision	<u>06</u>
	Phase 1—Visionaries	
	Phase 2—First Mover Challenge	
	Phase 3—Second Mover Advantage	
	Phase 4—Third Mover Disadvantage	
3	Three Strategic Actions to Launch Today	<u>12</u>
	About the Authors	<u>14</u>



1 The Race for AI Advantage Has Entered Its Decisive Phase

The cost of applying AI has already fallen sharply. Value realization has picked up significantly. Yet differentiation is still possible—but it will not stay that way for long. Now is the time to go all-in on AI.

The *BCG AI Radar 2026* survey found that 50 percent of CEOs believe their job stability depends on getting AI strategy right. One of their greatest worries might be **timing**: Executives who invested in AI early paid a steep premium for a technology that was not yet ready. Executives who invest last will pay a different price: they will acquire the same capabilities as their competitors at roughly the same cost, but by then AI will be table stakes without a competitive advantage to show for it.


Between these two failure modes lies a finite window—the **second mover advantage**—where costs have already fallen sharply, the technology is mature enough to deliver measurable returns, and the potential for differentiation is still real enough to justify the investment.

The **business case is already measurable**. Companies that have strategically integrated AI are now more than four times as likely to achieve significant EBIT impact as organizations still confined to early-stage pilots, as we found out in a recent study amongst European C-level executives. The question is no longer whether AI creates value—but when companies choose to capture it.

The **window of opportunity is open right now**. However, it will not stay open much longer. This report explains the mechanics behind that bold claim, shows where each major enterprise function currently sits in the window, and identifies what executives must do to capture value before it closes.

EXHIBIT 1

Now is the time to utilize the ...



Source: BCG analysis.

Methodology

This report combines evidence from multiple BCG research initiatives, proprietary analyses, and external data sources. The central framework of the "Second Mover Advantage" builds on BCG's AI Radar surveys (2024–2026), BCG's Build for the Future research, prior sector-specific AI studies, proprietary analyses on AI adoption and value creation, and a dedicated executive survey conducted with Statista Research.

The Statista Research survey was conducted between April 1 and May 8, 2026, using computer-assisted telephone interviews (CATI) with 555 CEOs, C-level executives, and senior managers across twelve European countries. Company size thresholds reflected regional market structures, ranging from organizations with 100 or more employees in Central and Eastern Europe to 500 or more employees in Germany, France, and the UK.

Unless otherwise indicated, findings presented in this report are derived from the synthesis of these research sources rather than from a single survey.



2 Three Curves, One Decision

Cost of Entry – Value Realization – Value Differentiation

Every technology cycle produces the same **timing dilemma**: invest too early, and the cost is prohibitive and the returns speculative; invest too late, and the advantage has evaporated. The resolution is not a question of judgment or risk appetite alone—it is a structural question about the interaction of three curves: **cost of entry, value realization, and value differentiation**.


The reward available from AI investment is not a simple function of technology maturity. It is the product of two distinct sub-curves—value realization and value differentiation—minus the cost of entry, and their movement over time is the analytical key to this article.

Net reward is high only when realization and differentiation are strong at the same time, so the shape it traces is an inverted-U: low in the early phase, when cost is high and realization is weak; rising steeply as realization improves while


differentiation remains; peaking when realization is near its maximum and differentiation has not yet evaporated; then declining as differentiation erodes and the investment delivers only catch-up value. **The four phases that follow trace this curve from left to right.** Exhibit 2 maps all three curves, and their product.

EXHIBIT 2

Input curves: Cost of Entry, Value Realization, Value Differentiation



Net reward curve



Source: BCG analysis. Curves are directionally illustrative.

Phase 1—Visionaries

The four phases trace a single arc, and in the first, none of the curves cooperated. **Cost of entry was punishing:** running a frontier large language model cost approximately \$30 per million tokens in 2021, and enterprise-grade deployment required custom engineering that few organizations could staff. **Value realization was low**—hallucination rates were high, and the vast majority of initiatives stalled at proof of concept.

Differentiation, in principle, was enormous, because almost no one held the capability; but it could not be captured, because

the technology was not yet ready to deliver it. With cost high and realization weak, **net reward sat near the bottom** of the curve.

This is why the earliest movers—the visionaries—paid a steep **premium for a technology that was not yet ready**, often earning a negative return. Some used the head start to build advanced AI teams that still serve them today; others exhausted their organizations' patience and burned the topic internally.

Phase 2—First Mover Challenge

In the second phase, one curve began to move. **Cost of entry fell sharply:** by 2024 the price of running a capable model had dropped below \$2 per million tokens, a decline of more than 93 percent in three years, and the models available at that price were materially more capable than the ones that had cost fifteen times as much three years earlier.

Value realization improved but stayed shallow in practice—the technology worked in demonstrations, yet scale constraints and low adoption kept most deployments from maturing. Differentiation was still high, because few had truly

operationalized AI, but the hype cycle drew a crowd, generating excitement and buzz faster than results. Net reward lifted off the bottom without yet reaching its peak.

The first movers jumped on that hype and overpaid for early models. But their central problem was not ambition—94 percent of organizations maintained their AI investment throughout this period. It was **insufficient execution depth:** organizations cycled through pilots without advancing to scale, capturing the buzz of the reward curve's ascent without the returns that only depth delivers.

Phase 3—Second Mover Advantage

In the third phase, the three curves finally align. Cost of entry has fallen to a quarter of a dollar or less per million tokens, and enterprise-grade build time has compressed from roughly eighteen months to three to six. Two of the forces behind this are effectively irreversible: each hardware generation delivers more compute per dollar, and distillation lets ever-smaller models match what only a frontier model could do a year earlier—efficiency that, once achieved, does not decay. Competition among vendors and a credible tier of open-weight models pushes prices down further still.

The net reward from AI investment peaks when value realization is near its maximum and differentiation has not yet evaporated. All the evidence we see indicates that in most industries and functions, this peak is now.

Value realization has reached unprecedented scale. Agentic AI is production-ready: BCG's *Build for the Future 2025* study finds 46 percent of companies already piloting or deploying agents. Implementation playbooks are established, anchored in a proven AI-first operating model. First economic results are visible within months rather than years—future-built companies reach full deployment in 9 to 12 months, far faster than any previous digital disruption.

And differentiation, crucially, has not yet evaporated. It has narrowed, but it remains real: only 5 percent of companies capture value from AI at scale, while 60 percent see minimal returns despite heavy investment. Because a second mover reaches production in three to six months rather than the eighteen or more the pioneers required, and benefits from mature, ready-made products rather than custom builds while facing organizations finally prepared to adopt, it can still open a structural lead in efficiency, service, and speed.

With cost low, realization high, and differentiation still available, net reward is at its peak—and that is the definition of the second mover window. And because the same company can sit at opposite ends of this curve across its functions, timing must be judged function by function, not enterprise-wide.

WHY NOW: THE QUANTITATIVE CASE

The strongest evidence that the second mover window is open is not falling compute costs—it is business performance. Companies that have operationalized AI end-to-end are already **more than four times as likely to generate significant EBIT impact** than organizations still running isolated pilots.

Three measurable signals—compute cost trajectory, corporate investment acceleration, and CEO conviction—confirm that the net reward curve is currently at or near its peak across most industries. Each is grounded in data rather than projection.

Investment: doubling in a single year

According to BCG's *AI Radar 2026* survey, corporate AI investment as a share of annual revenue reached approximately 0.6 percent in 2024, rose to 0.8 percent in 2025, and is on track to reach 1.7 percent in 2026—a near-doubling in twelve months. In technology and financial services, the survey puts the figure above 2 percent. It also finds that 94 percent of organizations will keep investing even if AI does not pay off financially in 2026; only 6 percent plan to pull back.

CEO conviction: from delegation to personal ownership

The same survey shows how far ownership has shifted to the top: 72 percent of CEOs now describe themselves as the main AI decision maker in their organization—double the share a year ago. 50 percent say their own job stability depends on getting AI right. Roughly 90 percent believe AI agents will enable measurable ROI in 2026, and these executives have already committed more than 30 percent of their 2026 AI budgets specifically to agentic AI.


EXHIBIT 3

Three Quantitative Signals

Three Signals That the Optimal Moment Is Now


LLM cost¹

Cost per 1 million tokens (in USD)




Corporate AI Investment²

AI investment as % of annual revenue



CEO conviction²



¹OpenAI pricing history ²BCG AI Radar 2026 Survey (n=2,360; n=640 for CEOs)
Source: BCG analysis

Another study we conducted with Statista in 2026 confirms the conviction and exposes a paradox: for instance, this survey of 555 European C-level executives found 86 percent of German executives expect companies that fail to scale AI within three years to fall structurally behind. Additionally, 78 percent say investing today yields more mature, scalable solutions than the earliest movers got. Yet AI budgets have fallen since 2025. Thus, the conviction is nearly universal, but the action to capture it is not following.

Not All Functions Are in the Same Moment

The second mover window does not open and close simultaneously across all parts of the enterprise. Functions vary in where they currently sit on the net reward curve: how mature AI deployment is within their context, how much


differentiation remains, and how much value they put at stake for the business. These differences make prioritization both possible and essential.

The right framework for prioritization combines two dimensions. The horizontal axis is the function's current position on the net reward curve—ranging from "too early, value not yet realizable" through the peak of the window to "commoditizing, differentiation already eroded." The vertical axis is the value at stake—the size of the economic lever the function controls, measured by the share of cost base or revenue it influences. The most urgent investments are those that combine high current net reward with high value at stake. Exhibit 4 applies this framework to the major enterprise functions.

EXHIBIT 4

Function Prioritization: Where Is Each Function in the Second Mover Window?

Positions are directionally illustrative



Source: BCG analysis



Software and R&D: the clearest case for immediate action.

Software engineering and research-and-development functions sit at the current peak of the second mover window: high value realization, meaningful differentiation, and a short window before commoditization. The productivity multipliers are already verified at scale—three to five times for new software development, more than 50 percent cost reduction for legacy migration—and the operating model that delivers them, what

we call the "Agentic Software Factory", is deployable now. The differentiation window in this function will last two to three years at most. Organizations that have not transformed their software delivery model in that time will find themselves competing at commodity productivity rates against competitors who moved faster. (for a deep-dive, see the sidebar on Agentic Software Factory below)

The Agentic Software Factory: The Best Available Example

Software development is the clearest current example of the second mover dynamic at work—a function where the technology is production-ready, the productivity multiplier is large, the differentiation is still real, and the commoditization timeline is short and visible. It is the function where acting now produces the largest verifiable advantage, and waiting the longest produces the largest verifiable penalty.

We call this operating model the "Agentic Software Factory"—a concept inspired by the "dark factories" of manufacturing that run with the lights off because no human is on the floor. AI coding agents autonomously build, test, and deploy software around the clock, while humans define business intent, set quality guardrails, and review outcomes—but do not write the code. The economics are striking. Where copilot-style assistance delivered incremental gains, organizations operating at this level report three to five times productivity on average, and more than 50 percent cost reduction at scale in legacy migration. In our own work, a five-day AI task force converted two business-critical applications initially estimated at hundreds of person-days, reaching over 50 percent productivity gains at scale.

The commoditization trajectory is equally clear: AI-orchestrated delivery will be industry standard within two to three years. The organizations that build the infrastructure, operating model, and institutional knowledge to operate this way by then will have accumulated a development-velocity, technical-debt, and time-to-market advantage that latecomers cannot replicate. The Agentic Software Factory is not a future concept. It is a deployable capability today—with differentiated returns that are finite in duration.



Marketing and sales: differentiation still real, moving fast.

In customer-facing commercial functions, the window is open but narrowing. In automotive, BCG and OpenAI asked in *Will AI Become the Best Car Sales Advisor?* what GenAI does to an OEM's top line depending on whether it acts. Their answer is unambiguous: industry players must embrace change and act decisively now. By 2030, more than 40 million car-buying journeys a year, up to 20 million in China alone, will be meaningfully shaped by GenAI. OEMs that move decisively stand to gain up to 20 percent in top line globally; those that wait risk losing up to 15 percent.

This gap is not created by the technology but by timing, and its width tracks local conditions: in China, where adoption is fast and brand loyalty low, the spread between leaders and laggards reaches 51 percentage points; in Europe, where adoption is slower and loyalty higher, it is narrower—roughly +12% against -7%—but the direction is the same. The principle generalizes: in the second mover window, the reward for acting and the penalty for waiting rise together. Fast movers take volume from laggards and price better through sharper customer matching; laggards face attrition and a downward price spiral.

The automotive analysis applies broadly: Firms deploying GenAI across the buying and servicing journey are already generating measurable conversion and price realization gains. The key differentiation today is between organizations that have moved to AI-native customer journeys and those still operating on legacy acquisition and engagement models. Within 18 months, that gap will define competitive position in most consumer sectors.



Operations and customer service: act now, before the window narrows further.

Operations and supply chain carry among the highest value at stake of any function, and AI-First customer service shows why the window is open today: where AI handles tickets end-to-end, with humans managing edge cases, it delivers a structural cost and capacity advantage—1,500 tickets per agent in an AI-First model versus 150 to 250 in a traditional one, with response times collapsing from hours to seconds. This represents meaningful differentiation now, but the adoption rate in customer service is higher than in software or marketing. The window is open, but closing faster here than in R&D.



Finance, HR, and procurement: the window is opening, not yet at peak.

Finance, HR, and procurement functions are at an earlier stage of AI maturity in most organizations. Point solutions are established—AP automation, resume screening, spend analytics—but end-to-end AI-First transformation of these functions remains rare, so their net reward has not yet peaked. Procurement is the standout case: because it governs external spend, its value at stake is far higher than that of finance or

HR, which mainly raise internal productivity. While all three sit early on the reward curve, procurement has the largest lever waiting to be pulled. This creates a different investment logic: building foundational data infrastructure and beginning structured deployment now, in anticipation of the peak that arrives in 18 to 36 months. Waiting until the peak to start building is equivalent to starting a factory after the first shipping deadline.

Phase 4—Third Mover Disadvantage

In the fourth phase, the curves that made the window valuable close it. Cost of entry keeps falling—but that decline now reaches every competitor equally and **confers no advantage on anyone**. Value realization is high for everyone: the capability has become table stakes. And differentiation has evaporated, because once a capability is universal, further investment delivers parity, not advantage. Net reward falls back toward the bottom, offering only catch-up value.

This is the price the last movers pay. They acquire the same capabilities as their competitors at roughly the same cost, but **with no competitive advantage to show for it**. Waiting shaved a little more off the price of entry; it did nothing about the differentiation lost in the meantime.

The real cost of waiting is the competitive gap opened by rivals already operating at the efficient frontier—and that gap does not come back down. Those who move last are left behind as competitors scale, the cost gaps between them widen, and what they face is not a temporary lag but a structural loss of competitiveness that is hard to reverse.



3 Three Strategic Actions to Launch Today

The strategic case for acting now is compounded by a financial forcing function that most leadership teams have not yet fully internalized.

Capital markets are beginning to price AI execution capability into valuation models, and analysts will soon demand evidence of measurable AI impact to justify the investment levels being committed.

As mentioned before, the *BCG AI Radar 2026* survey found that 50 percent of CEOs believe their job stability depends on getting AI strategy right. This is an anticipation of shareholder pressure that is forming. The logic is straightforward: if corporate AI investment doubles to 1.7 percent of revenue in 2026, and 94 percent of organizations say they will continue investing even without short-term returns, the earnings cycles of 2027 and 2028 will be the moment when analysts require a coherent account of what that investment produced.

This creates an additional dimension of urgency for prioritization specifically. A broad AI investment program spread across many functions, all at proof-of-concept stage, generates costs without generating a narrative of impact. A focused program—two or three functions, end-to-end transformation, measurable before-and-after outcomes—generates both strategic advantage and a credible account of ROI. BCG's AI Radar 2026 data on the three CEO archetypes makes the contrast concrete:

- Trailblazers, who represent approximately 15 percent of CEOs and commit 60 percent or more of their AI budget to agentic AI while upskilling more than 70 percent of their workforce, are already generating track records that analysts will reward.

- Pragmatists—the 70 percent who invest selectively without deep commitment, accumulating costs without building comparable evidence of return.
- Followers are doing neither, which is the least defensible position of all.

The obstacle is rarely analytical. The Statista survey found executives across maturity stages can assess AI's business value; far fewer feel confident making the investment decision, and that gap closes only once transformation is underway.

Across Europe, decision-confidence ranges from 84 percent in the Nordics to 37 percent in Eastern Europe, and the barriers cited are structural (data quality, legacy IT, fragmented architecture), not budget, which ranks last. This is why the investor clock is dangerous: the companies least able to answer for their AI spend are not those short of capital, but those that never built the confidence and data foundations to integrate deeply. Understanding is not the bottleneck. The capacity to commit is.

Executives who recognize the need for action in their own corporation should launch the following strategic actions today:

1 Map every major function in your company against the net reward curve.

Using Exhibit 4 as a starting point, assess each function on three questions: how far AI adoption has spread among competitors, whether differentiation is still available, and how much value the function puts at stake. The output is a ranking of the two or three functions where realization, differentiation, and value at stake align most favorably—not an attempt to transform everything at once. For most enterprises, Software and R&D tops that list, followed by the primary customer-facing function.

2 Commit to end-to-end transformation in the highest-priority functions.

As the First Mover Challenge showed, the failure of 2022 to 2024 was not insufficient ambition but insufficient execution depth. The distinction between Phase 2 and Phase 3 is the willingness to pursue end-to-end process redesign rather than

AI layered on top of unreformed processes—and the payoff of that depth is measurable: the share of companies achieving significant EBIT impact (3 percent or more) rises from 10 percent among early-stage organizations to 42 percent among operational adopters and 65 percent among the strategically integrated, making organizations that operationalize AI end-to-end more than four times as likely to achieve significant EBIT impact as those still running isolated use cases. This is the structural change that generates the measurable outcomes investors will demand. For the prioritized functions, the commitment must be full transformation, not enhanced pilots.

3 Build the measurement infrastructure to make impact visible.

Every function-level AI transformation must be accompanied by a before-and-after measurement framework: baseline metrics established before deployment, outcome metrics tracked after. Revenue, cost per unit of output, cycle time, error rate, customer satisfaction—the specific metrics vary by function, but the principle is constant. Organizations that build this infrastructure now will have credible ROI narratives ready when analysts demand them. Those that do not will be unable to distinguish their transformation investments from their competitors' costs.

The second mover window is open.
Prioritize the highest-reward functions.
Commit to end-to-end transformation.
Build the measurement infrastructure to prove the impact. The time for each of these is now.

The second mover advantage in AI is a real and currently accessible opportunity—but it is a timed one. The cost curve has already fallen significantly. Value realization potential has reached high maturity. Differentiation is still meaningful in the highest-value functions, but it is eroding at different rates in different parts of the enterprise. The organizations that move now—with focus, commitment, and measurement—will build structural advantages in the functions that matter most. Those that continue to run pilots while waiting for perfect conditions will find that the window closed while they were watching it.

About the Authors



Dr. Andrej Levin

Managing Director and Senior Partner
Hamburg

Levin.Andrej@bcg.com



Dr. Ivan Tretiakov

Principal
Düsseldorf

Tretiakov.Ivan@bcg.com



Dr. Robin Wagner

Principal
Stuttgart

Wagner.Robin@bcg.com



For information or permission to reprint, please contact BCG at permissions@bcg.com. To find the latest BCG content and register to receive e-alerts on this topic or others, please visit [bcg.com](https://www.bcg.com). Follow Boston Consulting Group on [LinkedIn](#), [Instagram](#), [Facebook](#), and [X](#).

BCG

