



© Banstanks/AdobeStock

LIGHTS - Technology

AI skill herding and capital investment: When does following the crowd pay off?

ESCP Impact Paper No.2026-38-ENG

Pedro MASI, Yuanyang LIU, Chuanren LIU
University of Tennessee, Knoxville

Hao ZHON
ESCP Business School

[AI skill herding and capital investment: When does following the crowd pay off?

Pedro Masi*
Hao Zhong**
Yuanyang Liu*
Chuanren Liu*

ESCP Business School

Abstract

As firms race to invest in artificial intelligence, many are copying competitors' AI hiring and talent strategies in the hope of attracting capital and staying competitive. But imitation is not a universal success strategy, and its returns depend on a fundamental tension: whether capital markets are responding to what firms genuinely possess internally, or to what firms publicly signal they intend to become. Drawing on two complementary firm-level datasets, AI skill demand from job postings and realized AI workforce profiles from employee resumes, linked to investment outcomes, this paper shows that the capital market response to AI skill herding is driven by observable strategic intent rather than realized workforce composition. Firms that publicly signal AI hiring alignment with peers attract more capital investment, but only when those signals are backed by absorptive capacity: specifically, stronger STEM human capital and closer proximity to AI knowledge sources amplify the capital attraction benefits of AI herding. By contrast, herding measured from actual employee profiles, reflecting realized rather than intended AI workforce composition, does not generate capital returns, suggesting that investors respond to the observable signals of AI strategic positioning rather than to underlying AI workforce reality.

Keywords: artificial intelligence, investment, herding, absorptive capacity, human capital

* University of Tennessee, Knoxville

** ESCP Business School

ESCP Impact Papers are in draft form. This paper is circulated for the purposes of comment and discussion only. Hence, it does not preclude simultaneous or subsequent publication elsewhere. ESCP Impact Papers are not refereed. The form and content of papers are the responsibility of individual authors. ESCP Business School does not bear any responsibility for the views expressed in the articles. Copyright for the paper is held by the individual authors.

AI skill herding and capital investment: When does following the crowd pay off?

The managerial dilemma

Across industries, executives face mounting pressure to adopt AI. However, a stark paradox has emerged: while 88% of firms use AI in at least one business function, only 5% of AI projects generate meaningful impact on P&L statements (Ernst & Young, 2025; McKinsey & Company, 2025). The AI skills gap lies at the heart of this disconnect. Most AI projects fail not because of technical limitations, but because employees lack the skills needed to translate AI investments into measurable value (Nunes, 2025). This creates smoke and mirrors around AI adoption without delivering substantive results. Some companies have taken more aggressive approaches. For example, Klarna replaced 700 customer service positions with AI assistants built with OpenAI models (Digital Applied Team, 2026). The strategy quickly backfired when the AI agents delivered lower-quality responses to customers, forcing the company to reverse course and rehire human customer service representatives.

So, what should companies do? One natural response is to imitate rivals: hire AI talent, expand AI-related roles, and mirror the skill profiles of peer firms that appear to be ahead. This reaction is understandable. In uncertain markets, imitation can look like a rational shortcut. If leading firms are investing in AI talent, following suit may signal modernity, strategic ambition, and readiness to couple AI investments with human capabilities that can in turn translate into tangible business value. This skill-based imitation strategy is commonly known as "herding" (Koch, 2017). But AI herding raises a critical question for managers and investors alike: does following the AI talent trend improve firm performance, or does it create expensive symbolism without real capability? Our central argument is simple: **AI skill imitation pays off only when firms can absorb what they copy.** Firms that follow peer AI talent patterns are more likely to attract capital when they also possess strong internal knowledge foundations. Without those foundations, firms may imitate the appearance of AI readiness yet fail to generate the underlying value investors seek.

Why imitation alone is not enough

AI differs from many earlier technologies because it depends heavily on human capital. Buying software is one thing; building an organization that can integrate AI into workflows, decision-making, and product development is another. AI capabilities are embedded in people, teams, technical routines, and learning processes. As a result, firms cannot assume that copying rivals' hiring patterns will automatically produce similar outcomes.

This is where *absorptive capacity* matters. Absorptive capacity refers to a firm's ability to identify, acquire, understand, and apply external knowledge (Zahra & George, 2002). In the AI context, that means more than hiring a few specialists. It means having the workforce, technical base, and organizational environment needed to transform external AI signals into meaningful internal capability.

For this reason, the same imitation strategy can produce very different results across firms. For some, it becomes a credible signal of strategic readiness. For others, it remains

superficial and may even reduce confidence if investors view it as trend-following without execution capacity.

What the evidence shows

Using firm-level data that combines AI-related job postings, employee skill profiles, and capital investment outcomes, this study examines whether firms benefit from following peer AI talent strategies. The results show that AI herding on its own does not reliably improve capital investment outcomes. In fact, when considered alone, following rivals may even be associated with weaker investment performance. However, the picture changes substantially when firm-level hiring intent enters the equation. For instance, consider a firm that attracts \$1 million annually in equity investment. If that firm has strong internal STEM (Science, Technology, Engineering, and Mathematics) talent and is located near AI innovation hubs, the herding effect translates into roughly \$630,000 in additional capital inflows, a meaningful boost in investor interest driven by peer imitation on AI skills. Conversely, a firm in the same geographic proximity to AI activity but lacking internal STEM capability sees the opposite outcome: herding leads to approximately \$510,000 less in equity investment, as investors appear to penalize imitation that is not backed by genuine technical capacity. Moreover, when herding is coupled with prior AI patents, firms attract between \$50,000 and \$130,000 more in equity investment on average, suggesting that investors view a demonstrated track record in AI innovation as an additional signal of credibility that amplifies the herding effect. All effects are statistically robust, underscoring that the direction and magnitude of AI-driven herding depend critically on whether a firm has the talent and innovation record to credibly attract investor confidence.

Overall, the positive effects of AI herding emerge most clearly when firms have:

- **strong employee STEM-based capability**, which reflects deeper technical knowledge embedded in the workforce;
- **proximity to AI knowledge sources**, such as being closer to locations where AI invention and technical expertise are concentrated; and
- **strong prior innovation outputs**, which reflects measurable AI outputs generated by the firm in the form of patented products or services.

When these conditions are present together, imitation becomes more valuable. Investors appear more willing to reward firms whose AI talent moves are supported by real absorptive capacity. In other words, investors do not respond only to visible AI ambition; they respond to whether that ambition looks implementable.

What decision makers should take from this

For executives, the main lesson is that AI strategy should not be treated as a race to mimic competitors as quickly as possible. The more important question is whether the organization can absorb and deploy what it is imitating. That means asking: Do we have the technical talent base needed to integrate AI into the business? Can our teams translate new AI hires into productive routines and decisions? Are we building a workforce that complements AI, rather than simply signaling adoption? Are we investing in capability development, not just external visibility?

For investors, the message is just as critical. Not all firms that publicly align themselves with AI trends are equally positioned to benefit. A credible AI narrative depends on whether a firm has the organizational capacity to translate imitation into effective

execution. In this context, “AI washing” is a real concern. In 2025, the SEC charged the former CEO of the e-commerce app Nate for allegedly claiming the platform was “fully automated” using AI, when in reality, it relied on human workers in the Philippines and Romania to manually complete transactions (Fortune, 2025). In Europe, the EU AI Act directly addresses this risk by tightening transparency and accountability requirements (European Parliament, 2023). The regulation imposes clear obligations on firms to accurately disclose when and how AI systems are used, particularly for high-risk applications, and prohibits misleading claims about AI capabilities. By establishing enforceable standards around documentation, oversight, and truthful representation, the AI Act serves as a regulatory backstop against AI washing and helps align marketing claims with actual technological deployment.

For policymakers and educators, the findings reinforce the importance of regional and workforce-level capability development. If access to AI talent and knowledge ecosystems shapes firm outcomes, then uneven capability distribution may widen the gap between firms that can benefit from AI and those that cannot.

For society, the findings are equally important: firms that adopt AI skills without adequate internal capabilities, risk disruptive workplace displacement, not because AI itself demands it, but because abrupt adoption strategies that fail to assess whether organizations can reliably couple AI infrastructure with workforce capabilities leave employees exposed to avoidable layoffs.

Finally, firms bear a great deal of reputational risk due to AI- commentary. As Duolingo’s “AI-first” memo showed, framing automation as a bold AI strategy while downplaying human contributors triggered a social- media firestorm: users threatened to quit, major social channels were deleted, and it took a long, very public effort by the CEO to repair the brand narrative and explain that the company was not simply replacing people with AI (Customer Experience Dive, 2025).

Three principles for evaluating AI human capital signals

Drawing on our findings, we suggest three principles that may help executives and investors assess which AI human capital signals are more likely to translate into capital outcomes.

1. Read AI job postings as strategic intent, not realized capability

When a firm announces an aggressive AI expert hiring campaign, or mirrors the AI skill profiles of leading peers, it is communicating a strategic ambition, not a demonstrated capability. Our research shows that job posting-based AI herding consistently predicts capital attraction, precisely because it is the signal that is visible to and interpretable by capital markets. But this same visibility is also its limitation: firms can post for AI talent without possessing the internal foundations to absorb, develop, and deploy it effectively. Sophisticated investors should therefore treat AI job posting signals as a starting point, not a conclusion. The critical question is whether the firm’s existing workforce can support the AI ambition being advertised. Firms with deep STEM human capital and geographic proximity to AI knowledge ecosystems, universities, AI research clusters, and established AI firms, are far better positioned to convert posting-based AI signals into realized capability. Firms without these foundations are signaling aspiration. The gap between the two is where capital is most frequently misallocated.

2. Triangulate posting signals with realized workforce data

Only 8% of organizations have reliable, bottom-up data on the AI skills their workforce currently possesses (MuchSkills, 2025). This means that when a firm signals AI strategic alignment through job postings, there is a high probability that the firm itself does not have a clear picture of its internal AI capability baseline. For investors, this structural blind spot creates a systematic risk: capital may flow toward firms that are skilled at communicating AI intent rather than toward firms that are genuinely building AI capability. The practical remedy is triangulation. Posting-based data, from sources such as Lightcast, captures what firms say they want to become. Resume and workforce profile data, from sources such as Revelio Labs, captures what firms actually are. Our research finds that these two dimensions of AI herding tell meaningfully different stories. AI herding from job postings predicts capital attraction; AI herding from actual employee profiles does not. Investors who rely exclusively on the former are pricing intent; those who cross-reference with the latter are pricing capability. The firms worth backing are those in which both signals converge.

3. Prioritize absorptive capacity as the credibility filter for all AI signals

Not all AI signals are equally credible, and the difference comes down to absorptive capacity, the organizational foundation that determines whether a firm can recognize, integrate, and exploit the AI knowledge it claims to be pursuing. Our research identifies two dimensions that investors should monitor closely. The first is potential absorptive capacity: the firm's STEM workforce depth and proximity to AI knowledge sources. These determine whether the firm can acquire and assimilate AI knowledge in the first place. A firm posting extensively for AI talent but lacking a STEM engineering base, or located far from AI talent clusters, faces structural barriers to realizing its AI ambitions regardless of how credible its job postings appear. The second is realized absorptive capacity: a demonstrated track record of converting AI knowledge into tangible outputs, most directly evidenced by prior AI patent activity. Firms with a history of AI exploitation, patents filed, products shipped, AI systems deployed, carry substantially more credibility when they signal further AI expansion. Firms whose AI narrative is built entirely on forward-looking intent, without this exploitation track record, represent a higher-risk allocation.

Conclusion

In fast-moving technology markets, imitation often looks like the safe bet. But when it comes to AI, copying rivals' skill sets is not enough, and not all AI human capital signals are equal. Two fundamentally different signals are at play: a realized capability signal, derived from the actual AI skills present in a firm's current workforce, and a strategic intent signal, derived from AI skill requirements in job postings. The first reflects what the firm has already built; the second reflects what it is publicly committed to become. Confusing the two is among the most common and costly mistakes in due diligence.

Our results show that capital markets respond to the intent signal, not the capability signal, the act of publicly aligning AI hiring plans with peer strategies is what attracts capital, while quietly possessing that capability does not. Yet intent only converts to value under the right conditions. Firms with strong employee-level STEM capability and proximity to AI knowledge sources can turn imitation into a credible signal of future potential, one that attracts meaningful capital. Similarly, firms with measurable, patented

AI outputs command greater investor interest when that innovation track record is coupled with peer-driven talent acquisition. For firms without those foundations, however, imitation risks becoming costly symbolism rather than a value-creating strategy.

This makes absorptive capacity the critical intervening variable. An intent signal backed by STEM workforce depth, proximity to AI knowledge clusters, and a demonstrated track record of exploiting prior innovations is a credible signal; it tells investors that the firm has the organizational infrastructure to identify, integrate, and exploit AI capabilities as they become available. An intent signal without these foundations is entirely different: at best, a genuine aspiration that may take years to materialize into productive capability; at worst, strategic mimicry designed to attract capital the firm is not positioned to deploy effectively. Before crediting any AI signal, whether forward-looking or capability-based, the first question should be whether the firm's absorptive foundations make that signal real.

References

Customer Experience Dive (2025), 'Duolingo went "AI-first" and then came the consumer backlash', *Customer Experience Dive*, 7 August. Available at:

<https://www.customerexperiencedive.com/news/duolingo-ai-first-consumer-backlash-lessons/757133/>

Digital Applied Team (2026), 'Klarna reverses AI layoffs: why replacing 700 workers backfired', *Digital Applied Blog*, 9 March. Available at:

<https://www.digitalapplied.com/blog/klarna-reverses-ai-layoffs-replacing-700-workers-backfired>

Ernst & Young (2025), *Work reimagined survey 2025: AI, workforce and productivity trends*, EY Global Insights Report. Available at:

https://www.ey.com/en_gl/insights/workforce/work-reimagined-survey

European Parliament (2023), *EU AI Act: first regulation on artificial intelligence*, European Parliament. Available at:

<https://www.europarl.europa.eu/topics/en/article/20230601STO93804/eu-ai-act-first-regulation-on-artificial-intelligence>

Fortune (2025), 'Albert Saniger and Nate shopping app fraud case highlights risks of AI misuse', *Fortune Magazine*, 11 April. Available at:

<https://fortune.com/2025/04/11/albert-saniger-nate-shopping-app-fraud-ai-justice-department/>

Koch, A. (2017), 'Herd behavior and mutual fund performance', *Management Science*, 63(11), pp. 3849–3873.

McKinsey & Company (2025), *The state of AI: agents, innovation, and transformation*, McKinsey Global Survey. Available at:

<https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>

MuchSkills (2025), 'AI Workforce Planning: Why Most Companies Are Flying Blind',

MuchSkills Blog. Available at: <https://www.muchskills.com/blog/ai-workforce-planning>

Nunes, R. (2025), 'The AI training gap: why organizations struggle to build AI-ready talent', *RuiNunes Insights*. Available at: <https://ruinunes.com/ai-training-gap/>

Zahra, S.A. and George, G. (2002) 'Absorptive capacity: a review, reconceptualization, and extension', *Academy of Management Review*, 27(2), pp. 185–203.