

Unbounding Rationality:
Why AI is a Fundamental Issue for Strategy

Felipe A. Csaszar

Ross School of Business, University of Michigan

Ann Arbor, MI 48109

`fcsaszar@umich.edu`

September 7, 2025

*Chapter submitted for inclusion in the forthcoming book, provisionally titled
“Strategic Management at a Crossroads,” edited by Jay Barney and Todd Zenger.*

1 Introduction

The field of strategic management is accustomed to change, but the pace and profundity of recent developments in artificial intelligence (AI) represent a discontinuity of a different order. To grasp the velocity of this shift, one need only revisit Michael Wooldridge’s authoritative history of AI, published in 2020. In a figure chronicling the milestones of AI achievement, Wooldridge listed several tasks that remained “nowhere near solved”: understanding a story and answering questions about it, human-level automated translation, interpreting the narrative of a photograph, and writing interesting stories. Barely five years later, this list reads not as a set of grand challenges but as a catalog of conquered territory. Fueled by the transformer architecture (Vaswani et al. 2017) and thrust into the global consciousness with the launch of ChatGPT in late 2022, these breakthroughs have increasingly demonstrated that modern AI can master subtle, context-dependent, and generative tasks of human language (OpenAI 2023, Bubeck et al. 2023).

This development strikes at the very heart of our discipline. At its core, strategy is concerned with the application of intelligence to the firm’s most consequential problems: identifying and choosing among long-term courses of action to create and capture value. It is the domain of CEOs, boards, and consultants, a realm of analysis and judgment traditionally thought to be uniquely human. The sudden and dramatic advance in AI therefore raises fundamental questions for the theory and practice of strategic intelligence. Specifically, two questions become inescapable: First, to what extent will AI be able to perform the cognitive work of strategists, thereby altering strategy process? Second, how will AI materially affect the sources of firm heterogeneity and competitive advantage, thereby altering strategy content? Central to my approach, I treat these questions through the lens of human–AI collaboration: the practical locus of decision-making in the foreseeable future will be hybrid, with humans orchestrating and auditing AI-generated analyses rather than being replaced by them (Raisch and Fomina 2025, Choudhary et al. 2023, Iansiti and Lakhani 2020).

Given the velocity of these developments, offering definitive answers is a perilous exercise.

Nevertheless, this chapter argues that the answer to both questions is “yes.” AI is not merely another tool for incremental efficiency gains; it is a technology that may reshape the foundational assumptions of strategy theory, the core processes of strategic decision-making, and the ultimate content of strategic choices. The purpose of this chapter is to substantiate this “yes” claim and to propose a path forward for the field, ensuring that strategy scholars can meaningfully contribute to the critical academic and practical debates that AI will undoubtedly spur.

To do so, the chapter is structured as follows. Section 2 provides the affirmative case, arguing why AI is a fundamental issue for strategy by examining its potential to overcome the cognitive bounds that have long defined human decision-making. Section 3 directly confronts and addresses two powerful objections to this view: one rooted in the logic of competitive advantage and the other in the nature of strategic creativity. Section 4 then outlines a concrete research agenda, suggesting how the field can adapt its methods and focus to make tangible progress in this new era. Finally, Section 5 concludes by reflecting on the broader benefits of this intellectual pivot, for both the strategy field and for society at large.

A secondary, yet equally important, objective of this chapter is to help establish a shared language for this emerging conversation. The strategy field currently lacks a common conceptual framework to discuss AI’s implications, often causing scholars with different assumptions to talk past one another. By providing clear arguments, addressing core objections, and outlining avenues for inquiry, this chapter aims to build a lexicon and a common ground upon which a more cumulative and productive scholarly discourse can be built.

2 The Affirmative Case: Why AI is a Fundamental Issue for Strategy

The assertion that AI represents a fundamental issue for strategy rests on four interlocking arguments. The first is empirical: recent advances demonstrate AI’s capacity to master cognitive tasks of a complexity analogous to strategic decision-making. The second is theoretical: AI directly challenges bounded rationality, the foundational assumption upon which much of modern strategy theory is built. The third reflects an emerging scholarly consensus that a profound shift is underway. The final argument is pragmatic: the risks of ignoring this transformation, both for individual scholars and for the discipline as a whole, are simply too great. I elaborate on each of these points in turn.

First, AI has established a strong track record of matching or surpassing human performance in domains once considered the exclusive province of human intellect. From the intricate logic of chess and the profound intuition of Go to the creative nuances of art and literature and the abstract reasoning of the International Mathematical Olympiad, AI has increasingly matched or surpassed human performance in arenas defined by complexity, ambiguity, and creativity (Campbell et al. 2002, Silver et al. 2016, OpenAI 2023, Castelvechi 2025). This progression is accelerating, with the latest generation of models moving from human-comparable to above-expert performance in high-stakes professional domains. For instance, in medicine—a field that, like strategy, requires integrating diverse information to make consequential judgments—the most advanced AI now surpasses pre-licensed human experts in complex multimodal diagnostic reasoning (Wang et al. 2025). Given this relentless advance, it is a reasonable extrapolation to view the cognitive work of strategy as the next frontier.

What has dramatically accelerated this timeline for our field is the advent of Large Language Models (LLMs). For decades, strategy was insulated from automation precisely because its core computational challenge is language. The inputs to strategic analysis—news

reports, consultant briefings, customer feedback, competitor announcements—are textual. The outputs—strategic plans, board presentations, memos, speeches—are likewise linguistic. The transformer architecture (Vaswani et al. 2017) substantially lowered this barrier, endowing machines with the ability to process, reason about, and generate sophisticated, context-aware language. This breakthrough is not merely incremental; it unlocks the very medium of strategy for computational analysis. Early evidence already validates this potential in a strategic context. In a recent study, business plans generated by an LLM were, on average, rated more favorably by experienced investors than plans submitted by entrepreneurs to a leading startup accelerator (Csaszar et al. 2024; see also Doshi et al. 2025, Girotra et al. 2023). Taken together, these developments provide a powerful proof of concept: AI can perform the core strategic tasks of generating and evaluating novel courses of action at a level comparable to, and perhaps soon exceeding, that of trained humans.

Second, and more profoundly, AI strikes at the theoretical core of our field by offering a path toward unbounding rationality: the systematic relaxation of the traditional cognitive constraints on human decision-making. The concept of bounded rationality—the idea that human decision-makers are limited in their attention, memory, and processing capacity (Simon 1955)—is arguably the most critical assumption in the study of strategy. It explains why firms are heterogeneous, why managerial cognition matters, why heuristics and biases shape outcomes, and why organizational processes are necessary to coordinate intelligence. AI, however, operates under a fundamentally different, and generally less restrictive, set of bounds. It is not constrained by the same cognitive limits in memory, computation, or attention.

To grasp the implications of this shift, it is useful to adopt a cognitive perspective on the firm and consider the three primary processes through which organizations make intelligent decisions: search, representation, and aggregation (Csaszar and Steinberger 2022). AI has the potential to radically enhance all three. It can expand *search*, allowing firms to explore vast landscapes of strategic alternatives at a scale and speed that would be impossible for human

teams to consider. It can enrich *representation*, enabling the creation and manipulation of far more complex and nuanced models of competitive environments than the frameworks and simple heuristics that human cognition requires. And it can augment *aggregation*, simulating the “wisdom of the crowd” with virtual panels of experts or customers, and combining diverse viewpoints—including those of other AI agents—without the social frictions that plague human organizations. By altering these fundamental mechanisms of organizational intelligence, AI does not just offer a better tool; it changes the underlying logic of how strategic problems can be solved (Csaszar et al. 2024, Choudhary et al. 2023, Argyle et al. 2023).

Third, beyond these performance-based and theoretical arguments, there is mounting evidence that the field itself perceives a fundamental shift. The topic of AI is no longer a niche interest. At the recent 2025 Academy of Management meeting, for instance, the number of submissions mentioning “artificial intelligence” exceeded those mentioning “strategy” or “corporate social responsibility” combined. Special issues on AI and strategy at journals like *Strategy Science* and the *Strategic Management Journal*, along with new edited handbooks on the topic, have been met with a deluge of submissions.

This growing scholarly interest is also reflected in the perspectives of the field’s experts. In a poll conducted at a plenary session at the 2025 *Strategy Science* conference, the mean response from the audience of leading strategy scholars indicated an expectation that AI would perform at the level of a top-10-percent human strategist within 8 to 15 years (see Table 1). While forecasts are inherently uncertain, this convergence of opinion among experts points toward a potential for significant, near-term change—a collective assessment that warrants consideration by the field.

Finally, there are two pragmatic reasons for the strategy field to engage with AI. The first can be framed as a form of Pascal’s Wager. If one wagers that AI will be transformative and is right, the field will have prepared its theories, tools, and students for a new era. If one is wrong, time will have been spent exploring a fascinating technological frontier. But if one

	Mean	SD	Votes for each alternative						
			0-2	3-5	6-10	11-20	21-50	51-100	100+
Strategy Teaching Tasks									
Grade MBA strategy exams	8.1	22.3	56	8	6	1	2	0	4
Lead classroom discussions	13.2	21.3	17	25	15	9	6	3	2
Mentor struggling students	10.2	21.3	39	14	10	6	4	1	3
Strategy Research Tasks									
Invent new, widely accepted strategy frameworks	17.7	26.9	14	19	18	11	8	1	6
Synthesize 100+ articles into new theory	14.9	31.3	35	20	12	0	1	0	9
Write publishable strategy journal articles	12.2	22.4	22	25	16	4	5	2	3
Strategy Practice Tasks									
Diagnose Fortune 500 strategy challenges	8.9	18.6	32	26	7	5	4	1	2
Advise CEOs in major crises	9.8	16.4	23	27	13	4	8	1	1
Predict long-term industry impact of new technology	17.4	27.7	20	18	11	15	5	2	6

Notes. The number of participating individuals was 77. The question asked was: “For each of the following milestones, please estimate in how many years you think AI will be able to perform the task at the level of a top 10% professional. Please enter your answer in years.”

Table 1: Strategy scholars’ predictions on the number of years for AI to achieve expert-level strategic performance.

wagers that AI’s impact will be minimal and is wrong, the field risks intellectual obsolescence, ceding the most important strategic questions of the 21st century to other disciplines.

This leads to the second pragmatic argument: the danger of a self-fulfilling prophecy. If strategy scholars collectively decide that AI is not a “fundamental issue,” we will not invest in the research and curriculum needed to understand it. Consequently, the intellectual center of gravity for “AI Strategy” will migrate to computer science, economics, or information systems. We risk finding ourselves in the position of one of my former statistics professors, who once lamented that machine learning should have been a core part of statistics but instead became the province of computer science. To avoid a similar fate, and to ensure our field continues to be the primary locus for understanding the creation and capture of value, we must roll up our sleeves and engage directly with this technological shift.

3 Addressing Objections: Why AI’s Impact Will Be Fundamental

The affirmative case presented in Section 2, while compelling, is not without its challengers. Any claim of a fundamental shift invites skepticism, and in the case of AI and strategy, this skepticism is both healthy and intellectually rigorous. To build a robust foundation for a new research agenda, we must directly engage with the most powerful objections to the thesis that AI will fundamentally alter the nature of strategy and competitive advantage. In the spirit of Alan Turing, who in his seminal (1950) paper on computing machinery and intelligence preemptively addressed several potential “Contrary Views,” this section will confront two of the most significant objections emerging from contemporary strategy scholarship. The first, rooted in the resource-based view and the logic of competitive markets, questions whether a widely available technology like AI can ever be a source of sustainable competitive advantage. The second, grounded in cognitive science and the philosophy of knowledge, questions whether AI, as a fundamentally backward-looking and data-driven technology, can perform the forward-looking, creative, and causal reasoning tasks that lie at the heart of strategy.

3.1 The Competitive Advantage Objection: AI as a Homogenizing Force

The first major objection, articulated powerfully by Barney and his colleagues (see Barney and Reeves 2024 and Wingate et al. 2025), applies the classic logic of sustainable competitive advantage (SCA) to AI. The argument is elegant and intuitive: for a resource to be a source of SCA, it must be valuable, rare, and difficult to imitate. While AI is undeniably valuable, its digital nature makes it eminently imitable and scalable. As algorithms are open-sourced, hardware becomes commoditized, and talent diffuses, AI will become ubiquitous. Like electricity or the Internet, it will become a cost of doing business—a driver of operational

effectiveness that raises the bar for all competitors but provides a durable advantage to none. In this view, AI is not a source of differentiation but a powerful force for homogenization. The value it unlocks will be captured not by individual firms, but by the market as a whole, and true advantage will continue to reside in what Barney et al. term “residual heterogeneity”—the uniquely human domains of creativity, relationships, and passion that AI cannot replicate.

This argument is compelling, particularly when considering a hypothetical long-run equilibrium where AI technology has fully diffused. However, it overlooks the critical, path-dependent journey of “getting there.” Strategy is not played out in a frictionless steady state, but in a dynamic, Schumpeterian world of disruption and change. To assess AI’s strategic potential today, the more apt analogy is not the Internet of 2025, but the Internet of 1995. At that time, access to the Internet was also becoming widespread, yet not all firms possessed the same foresight or capability to leverage it. A select few—companies like Amazon, eBay, and eventually Google and Netflix—recognized that the Internet was not just a tool, but a new landscape upon which to build powerful moats grounded in network effects, economies of scale, and proprietary data. Their advantage was not derived from merely using the Internet, but from building novel business models that the Internet made possible.

We do not yet know with certainty how the competitive dynamics of AI will unfold, but if the Internet’s history serves as a guide, we can anticipate a similar pattern. The fundamental issue for strategists is not who has access to AI, but who can first imagine and build the novel products, services, and business models that leverage its unique capabilities. This process of discovery and implementation is neither instantaneous nor easy. It requires imagination, organizational agility, and significant investment, creating a window of opportunity for early movers to build advantages that may prove highly durable. The immense capital required to train frontier models, for instance, already creates formidable barriers to entry, challenging the notion of simple ubiquity. In the short and medium term—the timeframe in which competitive advantages are actually built—AI is poised to be a powerful engine of heterogeneity, not

homogeneity (Agrawal et al. 2023, Berg et al. 2023, McElheran et al. 2024).

One might counter, in the spirit of the original argument, that this very foresight and capability to leverage a new technology are themselves rare, valuable, and inimitable resources, thus falling under the umbrella of “residual heterogeneity.” This is a fair and important point. However, it shifts the locus of the debate in a productive way. Rather than being immutable traits, foresight and the development of dynamic capabilities are the very levers that managers can pull and organizations can cultivate—making them a central concern of strategy scholarship and practice. Therefore, the strategic imperative is not to wait for AI to become ubiquitous, but to actively build the organizational capacity to recognize and seize the opportunities it creates. Those firms that do so early and effectively are the ones most likely to build the next generation of AI-driven SCAs.

3.2 The Creativity Objection: AI as a Backward-Looking Predictor

The second profound objection, articulated with philosophical depth by Felin and Holweg (2024), challenges AI’s capacity for genuine strategic thought. Their argument distinguishes sharply between AI’s data-driven, probabilistic prediction and human cognition’s theory-based, causal reasoning. AI models, including LLMs, are fundamentally backward-looking; they learn statistical patterns from vast historical datasets to predict the next word or classify an image. Strategy, in contrast, is fundamentally forward-looking. It is not about predicting the future based on the past, but about *creating* a future that is different from the past. This requires what Felin and Holweg call “data–belief asymmetry”—the ability to hold a belief (e.g., in the possibility of heavier-than-air flight) that is not supported by, or even runs contrary to, existing data. It is this capacity for theory, for causal intervention, and for generating new knowledge through experimentation that, in their view, separates human strategists from machines (Zellweger and Zenger 2023, Felin and Holweg 2024).

This is a crucial distinction, yet it is also important to recognize that human cognition is itself constrained by the past. We, too, learn only from the data of our own experiences

and the accumulated knowledge of others. An AI’s access to this historical data, however, is of a fundamentally different scale; it is trained on a corpus of information that vastly exceeds what any individual could ever consume. Humans have developed sophisticated cognitive techniques to transcend their own experiential limitations: we use recombination and analogical thinking to create novel ideas from existing components, we employ scientific reasoning to test hypotheses and generate new data, and we collaborate to aggregate diverse perspectives. The critical question, then, is not whether AI is backward-looking, but whether it can be equipped with analogous computational processes to generate forward-looking, novel insights. The evidence suggests it can, through at least three distinct avenues.

First, AI systems can be designed to emulate and augment the sophisticated deliberative processes that humans themselves use to foster creativity. Rather than viewing AI as a monolithic brain, a more accurate vision is to think of LLMs as foundational components—akin to the transistor—that can be assembled into complex, intelligent systems. For instance, one can construct multi-agent systems (Minsky 1986, Selfridge 1959, Wooldridge 2009) that simulate the deliberative processes of a strategic planning team, assigning different AI agents distinct roles: one to generate a baseline proposal, another to act as a devil’s advocate, a third to simulate competitor reactions, and a fourth to model customer responses—effectively creating a “McKinsey in a box” (Csaszar et al. 2024:334). This moves beyond a simple query-response model to the automation of a sophisticated strategy *process* designed to challenge assumptions and articulate novel plans (Horton 2023, Argyle et al. 2023).

Second, beyond emulating human processes, AI can achieve novel outcomes by leveraging its unique computational strengths, particularly brute-force search at a scale unattainable by human cognition. By combining generative capacity with systematic search, AI-powered systems can “boil the ocean.” This approach leverages sheer computational power to compensate for potential weaknesses in nuanced judgment, a pattern seen in other domains where AI has achieved superhuman performance. In games like Chess (Campbell et al. 2002), Backgammon (Tesauro 1995), and Go (Silver et al. 2016), AI systems combine a relatively

rough evaluation function—a heuristic for judging the strength of a given position—with a massive-scale search through billions of potential future moves. This technique, epitomized by Monte Carlo Tree Search (Newell and Simon 1956, Browne et al. 2012), allows an imperfect but fast evaluation to be amplified into profound insight (Risi and Preuss 2020, Gaessler and Piezunka 2023).

Third, even the core cognitive gap identified by the objection—the distinction between correlation and causal reasoning—is beginning to be bridged by new computational methods. Emerging research at the intersection of causal inference and deep learning is developing techniques to infer causal structures, such as Bayesian networks (Pearl 1988, Spirtes et al. 2000), from data. These methods work by reformulating the computationally intractable search for causality as a continuous optimization problem solvable by neural networks (Luo et al. 2020, Jiao et al. 2024). While still nascent, this line of inquiry suggests that the rough, probabilistic evaluations of LLMs could be complemented or grounded by models that explicitly learn causal relationships (see, e.g., Ludwig and Mullainathan 2024), directly addressing a core component of human strategic reasoning (Felin and Zenger 2017, Camuffo et al. 2024).

Ultimately, strategy performed by AI may look very different from strategy performed by humans. We do not yet know the ultimate limits of AI’s creative and causal reasoning abilities. But if the history of AI in domains like chess and Go is any guide, we should expect that AI will achieve superior performance not by mimicking human intuition, but by leveraging its own unique strengths: massive scale, computational speed, and the ability to find novel patterns in complex possibility spaces. Indeed, it remains an open question whether an artificial strategist (Csaszar 2018:615) will ultimately require explicit causal theories to guide its search or if superior predictive accuracy will suffice (a tension that echoes the Lewin-versus-Jelinek debate described in Csaszar et al. 2024:335). It will likely not “think” like a human strategist, just as a modern dishwasher does not clean plates in the same way a

human does; the goal is not to replicate the process, but to achieve a superior outcome.¹

4 Making Progress: A Research Agenda for an AI-Powered World

Having argued that AI is a fundamental issue for strategy and having addressed the primary objections to this view, the key question becomes: How can our field make tangible progress in this new era? The challenge is significant, but the strategy field is surprisingly well-prepared to tackle it. Our intellectual roots are already in place for two key reasons. First, research on strategy process—the study of heuristics, decision structures, and deliberative tools like the devil’s advocate—has long been a central concern of our discipline. Indeed, the study of strategy process has always been about designing better algorithms for decision-making, even if those algorithms were to be executed by humans. Second, the Carnegie School tradition provides a powerful theoretical legacy, framing organizations as, in essence, a form of artificial intelligence designed to make intelligent decisions despite the cognitive limits of their human components (Csaszar and Steinberger 2022). The advent of machine intelligence does not render this tradition obsolete; it makes it more relevant and testable than ever before.

Building on this foundation, this section outlines a concrete research agenda. It is not an exhaustive list but a starting point, highlighting several areas of inquiry that become critically important as we move to understand strategy in an AI-powered world. I will argue for the renewed importance of taking firms to the new technological frontier, for a renaissance in strategy process research, for a sharper focus on foresight and verifiable benchmarks, for broadening the types of papers our journals accept, and finally, for a necessary evolution in how we teach the next generation of strategists.

¹This idea—that AI achieves its goals through fundamentally non-human-like processes—has a rich history. The most famous articulation is arguably from computer scientist Dijkstra (1984), who remarked that “the question of whether a computer can think is no more interesting than the question of whether a submarine can swim.” A similar, widely used analogy posits that airplanes fly without flapping their wings. In both cases, the point is that success comes from focusing on outcomes, not from mimicking a biological exemplar.

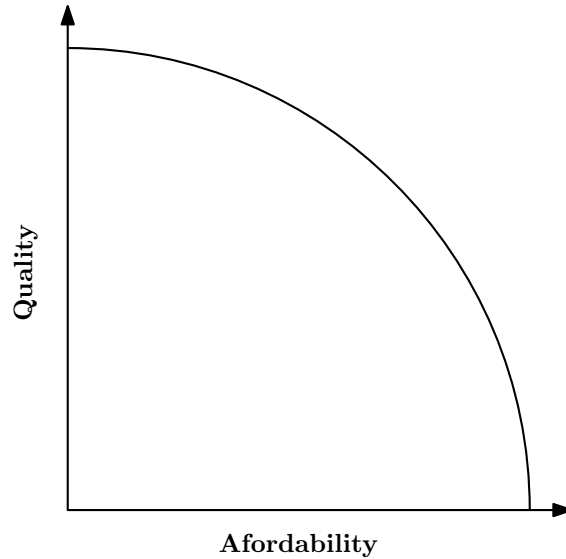


Figure 1: Porter’s (1996) distinction between operational effectiveness (moving closer to the frontier) and strategy (picking a position on the frontier).

4.1 Redefining the Frontier: Operational Effectiveness as Strategy

A well-known mantra in our field, famously articulated by Porter (1996), is that “operational effectiveness is not strategy.” We teach our students that strategy is about being different—choosing a unique competitive position—not just about running the same race faster. Porter visualized this distinction with the concept of the productivity frontier (see Figure 1). Strategy, in his view, is the act of choosing a distinct position *on* the frontier, while operational effectiveness is the process of moving *closer to* the frontier. This distinction was, in part, a reaction to the reengineering wave of the 1990s, when influential works like Hammer and Champy (1993) called for radical process redesign. Porter’s argument served as a pointed reminder that simply improving existing processes should not be confused with strategy.

While this conceptual separation is compelling, the AI revolution forces us to reconsider its practical implications. When a new general-purpose technology emerges that can dramatically shift the entire productivity frontier, the process of *getting to that new frontier* becomes a paramount strategic challenge and a key determinant of present and future performance. Firms will need to update nearly every aspect of their operations to integrate AI, trans-

forming both their strategy process (how decisions are made) and their strategy content (the sources of competitive advantage). This journey is neither simple nor instantaneous; it is a complex, path-dependent process of organizational change that requires significant investment, experimentation, and the development of new capabilities. In such a dynamic environment, the speed and efficacy with which a firm can reach the new frontier can itself become a source of temporary, and potentially durable, competitive advantage. Indeed, much of strategy research on topics like innovation, organizational learning, and dynamic capabilities is precisely about how firms can effectively move toward and redefine this frontier. Those that move faster will not only reap efficiency gains but will also be the first to discover the novel products, services, and business models that the new technology makes possible (Tushman and Anderson 1986, Lieberman and Montgomery 1988).

However, the field’s tendency to rigidly separate the journey *to* the frontier from the choice of position *on* the frontier serves as a cautionary tale. During the business process reengineering wave of the 1990s, many strategy scholars dismissed the movement as “just operations,” ceding the intellectual high ground to other disciplines. We risk making the same mistake today. The organizational transformation required to leverage AI is a fundamental strategic problem. It involves resource allocation under uncertainty, organizational design, and the cultivation of dynamic capabilities—the very heartland of our discipline. To remain relevant, we must engage with this challenge directly, recognizing that in an era of technological discontinuity, the race to the frontier is a strategic race. The competitive struggles of an incumbent like Intel in the age of AI serve as a “canary in the coal mine,” exemplifying the peril of falling behind the new technological frontier.

4.2 A Renaissance for Strategy Process Research

For decades, our theories about strategy process—the cognitive architecture that allows firms to perceive, represent, search, and aggregate information to arrive at high-quality decisions—have been rich in concepts but difficult to test rigorously. AI changes this entirely.

What was once the domain of conceptual frameworks and case studies can now be modeled, simulated, and tested computationally, opening the door to a renaissance in strategy process research.

We can begin by translating classic, human-centric process tools into explicit AI-driven systems. As recent work suggests, it is now possible to build multi-agent AI systems that automate sophisticated strategy processes (Csaszar et al. 2024). Rather than relying on a single AI to generate a complete strategy, we can orchestrate iterative cycles of proposal refinement: an initial AI agent drafts a strategic plan, which is then scrutinized by a specialized critic agent trained to identify weaknesses and blind spots. The refined proposal can be stress-tested through simulated market dynamics, with dedicated agents modeling how competitors might respond and customers might react. These reactions then feed back into the system, prompting a redesign of the original proposal—creating a continuous loop of strategic refinement that mirrors, and potentially surpasses, the iterative nature of human strategic planning and adaptation. We can systematically test whether such structured protocols outperform different types of AI systems, thereby turning classic organizational design ideas into testable computational hypotheses. Another powerful application of this approach is to create virtual panels of synthetic experts to simulate the “wisdom of the crowd,” allowing us to explore how the composition and aggregation rules of a virtual top management team affect decision quality under different environmental conditions (Breiman 1996, Schapire 1990, Ho 1995).

This research can move beyond simply mimicking human processes. We can design and test entirely new computational processes for strategy formulation. For example, we can create fully simulated economies where AI agents act as entrepreneurs, generating and refining business plans that are then evaluated by other AI agents acting as investors, allowing us to study the drivers of entrepreneurial discovery at scale (Csaszar et al. 2024). We can use evolutionary algorithms, in the spirit of Nelson and Winter (1982), to evolve populations of firm strategies and observe which ones survive under different competitive pressures (Holland

1975, Koza 1992). A central challenge, of course, will be ensuring these simulated markets are realistic enough to yield meaningful insights. But if this can be achieved, the ability to run thousands of these simulations would allow us to test specific theories about how competitive advantages emerge and erode over time.

A particularly exciting avenue involves moving beyond the concept of “local search,” which has played a central role in our theories of adaptation (Simon 1955, Levinthal 1997). AI poses the intriguing possibility that organizational search can become more global. Consider a complex strategic problem where an AI system must first identify the relevant decision variables—the “levers” that a firm could adjust. Unlike traditional approaches that assume these levers are given, an AI could discover them by analyzing vast amounts of data, industry reports, and case studies to identify N actionable dimensions of strategic choice. It could then generate scenarios representing a vast number of combinations of these levers (e.g., 3^N scenarios if each lever has low, medium, and high settings) and evaluate each combination. To do so accurately, it could devise a bespoke virtual panel of experts for each specific scenario. For instance, if a scenario involves building a power plant in a Vietnamese forest, the virtual committee could include synthetic experts on power generation, forestry, and Vietnamese political economy. Alternatively, to enhance accuracy, the system could run a virtual tournament, pitting all scenarios against each other in pairwise competitions to identify an ultimate winner. If the evaluation process is sufficiently accurate (for initial evidence, see Csaszar et al. 2024), the winning alternative would represent a high-fitness solution. Another “boil the ocean” approach could involve evaluating a vast matrix of strategic options by combining thousands of business models with thousands of potential markets (akin to the “X for Y” method of generating new ventures by analogy). Search conducted in this manner could direct a firm to jump to solutions far from its current position, fundamentally altering the incremental nature of organizational adaptation.

Furthermore, this opens up new questions about the optimal architecture of organizational intelligence. This inquiry extends beyond the design of purely computational systems—such

as whether to rely on a single, generalist AI or a distributed system of specialized AIs. Critically, it brings the human–AI interface to the forefront, making the cognitive division of labor in these hybrid teams a central research question. Understanding which tasks are best left to human judgment and which to machine analysis, and how to design the protocols that govern their interaction, becomes a fundamental challenge for the future of strategy process (Raisch and Fomina 2025, Shrestha et al. 2019).

4.3 The Primacy of Foresight and the Rise of Representational Complexity

A fundamental requirement for any machine learning system is a clear performance signal—a dependent variable it can learn to predict. For an AI to master chess, the signal is unambiguous: winning or losing the game. But for an AI to learn strategy, what is the equivalent signal? This question forces us to make explicit a concept that has always been central, yet often implicit, in our field: foresight.

At its core, strategy is about foresight, as strategists need to predict which courses of action will lead to superior performance (Csaszar and Laureiro-Martínez 2018). Our most enduring frameworks—from the Five Forces to VRIO to the strategy canvas—are fundamentally causal models designed to improve this predictive accuracy. The advent of AI does not change this goal; it makes it computationally tractable and testable. To train an “AI strategist,” we must be able to tell it when it has formulated a “good” strategy, which means we must be able to measure the quality of its foresight.

This pursuit of foresight places a new emphasis on the very models we use to generate those predictions: our strategic representations. The quality of any strategic choice depends on the quality of the representation—the model of the relevant aspects of the firm, product, and competitive environment—that informs it. Human-designed frameworks are necessarily simple, constrained by the limits of our own cognition. An AI, however, is not bound by the same constraints and can therefore operate with more complex representations. We

can imagine an AI devising a new strategic framework by extracting every relevant concept from the entire corpus of strategy research and then identifying which combination of those concepts best predicts firm performance, contingent on industry and firm characteristics. This may result in models that render the very notion of “five” forces quaintly restrictive.

The quality of any representation is ultimately measured by its predictive accuracy—how well its predicted outcome (\hat{y}) approximates the actual outcome (y). Achieving this accuracy, however, involves confronting a fundamental challenge: the classic bias-variance trade-off. This is precisely where the cognitive limits of human strategists become a binding constraint. To avoid being overwhelmed by noise (high variance), we rely on simplified models, which inherently carry strong assumptions (high bias). A framework like the Five Forces is useful because it is simple, but its simplicity is also its weakness. AI, by contrast, is uniquely equipped to manage this trade-off. Its ability to process vast datasets allows it to build and learn from highly complex, nuanced representations (low bias) without overfitting to random noise (high variance) (Csaszar and Ostler 2020).

This fundamental difference in capability suggests that the study of strategic representations will shift from analyzing the bounded mental models of managers to designing and testing the unbounded computational models of AI. Improving these computational representations is therefore on the critical path to advancing not only AI’s strategic capabilities but also the predictive power of strategy research itself.

4.4 From Art to Science: The Need for Verifiable Benchmarks

The pursuit of predictive foresight naturally leads to the question of measurement. To accelerate progress, a field needs common problems and shared metrics of success (Platt 1964). The AI community has long benefited from this, using benchmarks like the ImageNet competition for computer vision or the RoboCup for robotics to galvanize research and measure progress. Strategy, by contrast, has largely remained a non-verifiable domain—more akin to literary criticism than to physics, where the quality of a strategy is often judged

post-hoc and is subject to narrative interpretation. For AI to learn and improve, and for our field to become more cumulative, we must make the domain of strategy more verifiable. As the physicist Lord Kelvin famously noted, “if you cannot measure it, you cannot improve it.”

This requires developing two distinct but complementary types of measurement tools. The first is a suite of *descriptive scorecards* designed to characterize an AI’s strategic “personality” or style. Just as one would evaluate a human strategist on multiple dimensions, we need to develop metrics to profile an AI’s behavioral tendencies. These scorecards would capture its industry knowledge (both depth and breadth), its ability to shift between levels of analysis, its implicit risk preferences and stakeholder priorities, and its creative tendencies (e.g., its propensity to explore versus exploit). Creating these profiles is an important first step in understanding and comparing different AI strategists.

Beyond describing an AI’s style, we must also develop rigorous *performance benchmarks* to measure how well it performs. These can take several forms. One approach, inspired by the Turing Test, would be to create a “Strategy Turing Test” where human experts rate AI-generated strategic plans against human-authored ones on criteria like novelty and feasibility.² A second, more quantitative approach is to test an AI’s ability to make accurate predictions about out-of-sample outcomes—that is, to forecast firm performance or market shifts beyond the data on which it was trained. A third avenue involves designing a battery of tests to assess specific strategic capabilities. For example, one could design simulated tasks to test an AI’s depth of strategic reasoning, akin to the experimental work on iterated thinking in game theory (Nagel 1995). We could also directly measure its knowledge base, testing its grasp of real-world strategy knowledge by asking it to interpret and analyze current events from the news, as well as its mastery of conceptual strategy knowledge by quizzing it on the foundational frameworks from our field’s core texts. Collectively, such benchmarks would provide the clear performance “gradient” needed for AI systems to improve through

²Unlike the original Turing Test, which is a test of imitation, the goal of a Strategy Turing Test is not mere indistinguishability but superior performance. The aim is to generate plans that human experts judge as being of higher quality than those produced by their human counterparts.

reinforcement learning. They would also allow for direct, apples-to-apples comparisons of different AI systems and, by their public nature, create powerful incentives for developers to optimize their models for the complex, nuanced tasks central to strategy.

This effort would also force us to confront the boundaries of verifiability. Some aspects of strategy may resist quantification for the foreseeable future. Understanding this boundary—distinguishing the verifiable, analytical components of strategy from the non-verifiable, subjective ones—is itself a critical research question.

4.5 Broadening the Definition of Scholarly Contribution

The research agenda outlined above—one centered on the explicit design of strategy processes, foresight, and verifiability—implies a corresponding evolution in what the field recognizes as a valuable scholarly contribution. If strategy is to become a more cumulative and computational science, its intellectual infrastructure cannot be built solely upon the traditional hypothetico-deductive article. The very tools and concepts required for this new line of inquiry must themselves be treated as first-order scholarly achievements.

For instance, the development of the verifiable benchmarks and descriptive scorecards discussed previously cannot be a mere footnote to traditional research; it must itself be a primary form of contribution. A paper that introduces, validates, and refines a “Strategy Turing Test” or a robust forecasting benchmark provides the shared dependent variables and measurement tools necessary for a community of scholars to build upon each other’s work. Similarly, the creation of novel computational models—such as the multi-agent systems that instantiate theories of strategy process—should be viewed not merely as software engineering but as a form of theory articulation. In this context, “software experience” papers, which report on the design and implementation of such systems, become instrumental for sharing the tacit knowledge gained when abstract theories confront computational reality.

This intellectual infrastructure also depends on two other pillars: high-quality data and a precise lexicon. Curated datasets designed specifically to train and test AI strategists on

tasks of foresight are essential for moving from theoretical speculation to empirical validation. Likewise, the development of a new conceptual language—terms as foundational for this era as “absorptive capacity” or the “resource-based view” were for a previous one—is a critical theoretical contribution that enables clearer reasoning and cumulative theory-building. By formally recognizing these contributions—the tools, the data, and the language—our journals can actively construct the scaffolding upon which the next generation of empirical and theoretical work will stand.

4.6 Educating the AI-Native Strategist

Finally, the intellectual pivot described in this section has profound implications for how we educate the next generation of leaders. We are entering an era of ferment, where the primary challenge for firms will be managing innovation and change. This requires both entrepreneurs who can explore the vast new landscape of AI-enabled opportunities and intrapreneurs who can lead the transformation of incumbent firms.

Our curricula must therefore continue to emphasize the timeless strategic topics of creativity, innovation, and change management. However, this is no longer sufficient. To be effective, managers will need a deep, intuitive understanding of AI’s capabilities and limitations. They must become what we might call “hybrid professionals,” fluent in the languages of both business and computation (Csaszar et al. 2025, Mollick 2024).

Business schools have not always excelled at creating this type of professional. Too often, the “technology” component of the curriculum is superficial and quickly becomes obsolete. The challenge is to provide a deep, conceptual understanding of computational thinking that will endure even as specific AI models and platforms evolve. This may require fundamental changes to our curricula, our faculty, and our partnerships with other parts of the university.

The task is daunting, but essential. The value of a human strategist in the 21st century will not lie in performing analyses that an AI can do faster and better, but in the wisdom to ask the right questions, interpret the outputs, and lead the human organization through the

complex process of change. Our teaching must evolve to cultivate that wisdom.

5 Looking Forward

This chapter began by asking two fundamental questions: to what extent will AI alter the *process* of strategy, and how will it reshape the *content* of strategy? The preceding sections have argued that the answer to both is “profoundly,” and that engaging with this transformation is now a central task for our field. This endeavor is more than an intellectual exercise; it is an engagement with one of the grand challenges of AI. Strategy operates in what has been termed a “wicked environment”: a domain characterized by ambiguity, novelty, high stakes, and noisy, delayed feedback (Churchman 1967). Developing an AI that can navigate such an environment would represent a milestone in machine intelligence.

This question is not new. In a prescient 1943 conversation at Bell Labs, Alan Turing and Claude Shannon debated the ultimate frontiers of computation: which would be harder to create, an artificial scientist or an artificial CEO?³ In the decades since, the vision of an “AI scientist” has inspired a vibrant research program, with hundreds of tools developed aiming to automate the totality or parts of the scientific process (Chen et al. 2025). The “AI strategist,” by contrast, has received far less attention. Perhaps Shannon was right, and the wicked nature of strategy makes it a fundamentally harder problem for computation than science. The alternative, however, is more provocative: is it simply that the community best equipped to formalize the problem—the strategy field itself—has not yet fully dedicated itself to the task?

This latter possibility suggests not an inherent limit, but an intellectual goldmine waiting to be explored. Indeed, the field is uniquely prepared to take up this challenge. In a sense, strategy scholarship has always been in the business of designing algorithms. Our most enduring frameworks are, in essence, structured protocols for analysis, designed to be

³As Turing quipped to Shannon: “No, I’m not interested in developing a powerful brain. All I’m after is just a mediocre brain, something like the President of the American Telephone and Telegraph Company.” (Hodges 1983:251; see also Gleick 2011:205).

executed by human minds to improve decision-making. The arrival of AI presents a historic opportunity: for the first time, these algorithms can be instantiated in silicon, tested at scale, and used as building blocks for entirely new computational processes that transcend the limits of human cognition. This promises to accelerate the development of theory and sharpen its predictive power.

The ultimate impact of AI on the practice and study of strategy, of course, hinges on a great unknown: the future trajectory of AI development. It is useful, therefore, to consider two divergent scenarios. In the first, AI capabilities plateau near their current levels. Even in this world, AI's full strategic impact would take decades to unfold. The diffusion of a general-purpose technology is not instantaneous; it is a long, path-dependent process of co-invention and adaptation. The history of the Internet is instructive. In the early 1990s, visionaries like Nicholas Negroponte in his book *Being Digital* accurately foresaw the shift to on-demand media, the rise of e-commerce, and personalized news feeds (Negroponte 1995).⁴ Yet, it took decades for the necessary business models, organizational processes, infrastructure, and consumer habits to co-evolve and make these ideas an important part of the business environment. Similarly, integrating even today's AI requires immense innovation in finding viable applications, redesigning workflows, training personnel, and navigating the complex landscape of consumer, regulatory, and stakeholder acceptance. This multi-decade journey of adaptation and co-invention represents the central challenge for practitioners and a critical research frontier for scholars for the foreseeable future.

In the second, more intriguing scenario, AI achieves superhuman strategic ability. Would the field of strategy then become a sub-discipline of computer science? The answer, I believe, depends on the actions the field takes today. If strategy scholars cede this intellectual territory, we risk the self-fulfilling prophecy discussed earlier. If, however, the field dedicates itself to this task, a different outcome seems more likely. Consider the analogy of chess. Although

⁴See also AT&T's 1993–1995 “You Will” ads, which previewed technologies like in-car navigation systems, electronic toll collection, and streaming movies to your home (<https://www.youtube.com/watch?v=RvZ-667CEdo>).

the best software has been much stronger than any human for decades, the study of chess theory continues to thrive. So too, I suspect, will it be with strategy. Human oversight will remain essential, not only to ensure value alignment and mitigate risks like amplified bias or strategic miscalculation, but because humans, not algorithms, will ultimately experience and be held accountable for the consequences of strategic decisions. Moreover, unlike chess, strategy is not a closed game with fixed rules; the very objective is often to change the rules (Helfat 2021). This open-ended, dynamic quality will likely ensure a durable role for human judgment, creativity, and leadership.

Regardless of which scenario unfolds, the act of engaging with AI should benefit the discipline. It will force a greater precision in theory, more rigor in testing, and a more cumulative approach to knowledge-building. The challenge of teaching a computer to “do” strategy compels a formalization of concepts to the point where they are computationally tractable, a process that inevitably exposes hidden assumptions and clarifies causal logic. This aligns with the principle of “strong inference,” whereby a field accelerates its progress by systematically articulating and testing competing hypotheses (Platt 1964). In this way, AI can serve as the ultimate proving ground for our theories. As the computer scientist Donald Knuth famously remarked, “Science is knowledge which we understand so well that we can teach it to a computer” (Knuth 1974:668).

The history of AI is punctuated by profound philosophical debates over the limits of machine intelligence, most famously in the critique by Hubert Dreyfus in *What Computers Can't Do* (Dreyfus 1972). The response from AI pioneers like Herbert Simon and John McCarthy was not to engage in a protracted philosophical counter-argument. Rather, their approach was to treat such objections as empirical questions to be answered through construction. Their implicit stance was that the best way to discover what computers can do is to try to build them, reframing abstract qualities like “belief” into computationally tractable problems (McCorduck 2004:211–239). Adopting a similar spirit of constructive engagement—one that treats theoretical puzzles as engineering challenges—offers a powerful

way for the strategy field to advance. As this chapter has argued, such an approach would enable us to translate our theories into working, testable decision protocols; to establish shared benchmarks of foresight; and to build a cumulative base of evidence about which strategy processes more reliably create and capture value.

The stakes of this shift toward a more computational and verifiable science of strategy extend far beyond the academy. As the marginal cost of intelligence approaches zero, we are entering a new economic era, moving beyond the information age into an age of ubiquitous cognition. The project of unbounding rationality through AI unlocks what might be called *Radical Intelligence*: the capacity to embed and leverage sophisticated, near-zero-cost intelligence into every product, process, and customer interaction. This is the engine for what could be an industrial-scale discovery of value, with profound welfare implications for humanity (Amodei 2024). Just as AI has discovered novel, superhuman moves in the game of Go (Silver et al. 2016), it holds the promise of discovering novel business models and societal solutions that were previously beyond human imagination. The physicist David Deutsch (2011) offers a profound source of optimism for this pursuit: “Everything that is not forbidden by the laws of nature is achievable, given the right knowledge.” In this light, the task of the field of strategy can be seen as focusing the lens of artificial intelligence upon the frontiers of value creation, illuminating the path to a new era of human progress.

References

- Agrawal, A., J. S. Gans, A. Goldfarb. 2023. Artificial intelligence adoption and system-wide change. *Journal of Economics & Management Strategy* **33**(2) 327–337.
- Amodei, D. 2024. Machines of loving grace: How AI could transform the world for the better. Available at: <https://darioamodei.com/machines-of-loving-grace> (accessed November 29, 2024).
- Argyle, L. P., E. C. Busby, N. Fulda, J. R. Gubler, C. Rytting, D. Wingate. 2023. Out of one, many: Using language models to simulate human samples. *Political Analysis* **31**(3) 337–351.
- Barney, J., M. Reeves. 2024. AI won't give you a new sustainable advantage. Harvard Business Review, available at: <https://hbr.org/2024/09/ai-wont-give-you-a-new-sustainable-advantage>.
- Berg, J. M., M. Raj, R. Seamans. 2023. Capturing value from artificial intelligence. *Academy of Management Discoveries* **9**(4) 424–428.
- Breiman, L. 1996. Bagging predictors. *Machine Learning* **24**(2) 123–140.
- Browne, C. B., E. Powley, D. Whitehouse, S. M. Lucas, P. I. Cowling, P. Rohlfshagen, S. Tavener, D. Perez, S. Samothrakis, S. Colton. 2012. A survey of Monte Carlo tree search methods. *IEEE Transactions on Computational Intelligence and AI in Games* **4**(1) 1–43.
- Bubeck, S., V. Chandrasekaran, R. Eldan, J. Gehrke, E. Horvitz, E. Kamar, P. Lee, Y. T. Lee, Y. Li, S. Lundberg, H. Nori, H. Palangi, M. T. Ribeiro, Y. Zhang. 2023. Sparks of artificial general intelligence: early experiments with GPT-4. arXiv:2303.12712.
- Campbell, M., A. J. Hoane Jr., F. Hsu. 2002. Deep Blue. *Artificial Intelligence* **134**(1–2) 57–83.
- Camuffo, A., A. Gambardella, A. Pignataro. 2024. Theory-driven strategic management decisions. *Strategy Science* **9**(4) 382–396.
- Castelvecchi, D. 2025. DeepMind and OpenAI models solve maths problems at level of top students. *Nature* **644**(8075) 20.
- Chen, Q., M. Yang, L. Qin, J. Liu, Z. Yan, J. Guan, D. Peng, Y. Ji, H. Li, M. Hu, Y. Zhang, Y. Liang, Y. Zhou, J. Wang, Z. Chen, W. Che. 2025. AI4Research: A survey of artificial intelligence for scientific research. arXiv:2507.01903.
- Choudhary, V., A. Marchetti, Y. R. Shrestha, P. Puranam. 2023. Human–AI ensembles: When can they work? *Journal of Management* **51**(2) 536–569.
- Churchman, C. W. 1967. Wicked problems. *Management Science* **14**(4) B141–B142.
- Csaszar, F. A. 2018. What makes a decision strategic? Strategic representations. *Strategy Science* **3**(4) 606–619.
- Csaszar, F. A., M. G. Jacobides, P. Zemsky. 2025. The effects of artificial intelligence on management education. Available at SSRN: <https://ssrn.com/abstract=5325783>.
- Csaszar, F. A., H. Ketkar, H. Kim. 2024. Artificial intelligence and strategic decision-making: Evidence from entrepreneurs and investors. *Strategy Science* **9**(4) 322–345.
- Csaszar, F. A., D. Laureiro-Martínez. 2018. Individual and organizational antecedents of strategic foresight: A representational approach. *Strategy Science* **3**(3) 513–532.
- Csaszar, F. A., J. Ostler. 2020. A contingency theory of representational complexity in organizations. *Organization Science* **31**(5) 1198–1219.
- Csaszar, F. A., T. Steinberger. 2022. Organizations as artificial intelligences: The use of artificial intelligence analogies in organization theory. *Academy of Management Annals* **16**(1) 1–37.
- Deutsch, D. 2011. *The Beginning of Infinity: Explanations That Transform the World*. Viking, New York.

- Dijkstra, E. W. 1984. The threats to computing science. Manuscript EWD898, Department of Computer Science, University of Texas at Austin. Available at: <https://www.cs.utexas.edu/~EWD/transcriptions/EWD08xx/EWD898.html> (accessed August 31, 2025).
- Doshi, A. R., J. J. Bell, E. Mirzayev, B. S. Vanneste. 2025. Generative artificial intelligence and evaluating strategic decisions. *Strategic Management Journal* **46**(3) 583–610.
- Dreyfus, H. L. 1972. *What Computers Can't Do: A Critique of Artificial Reason*. Harper & Row, New York.
- Felin, T., M. Holweg. 2024. Theory is all you need: AI, human cognition, and causal reasoning. *Strategy Science* **9**(4) 346–371.
- Felin, T., T. R. Zenger. 2017. The theory-based view: Economic actors as theorists. *Strategy Science* **2**(4) 258–271.
- Gaessler, F., H. Piezunka. 2023. Training with AI: Evidence from chess computers. *Strategic Management Journal* **44**(11) 2724–2750.
- Girotra, K., L. Meincke, C. Terwiesch, K. T. Ulrich. 2023. Ideas are dimes a dozen: Large language models for idea generation in innovation. Available at SSRN: <https://ssrn.com/abstract=4526071>.
- Gleick, J. 2011. *The Information*. Pantheon Books, New York.
- Hammer, M., J. Champy. 1993. *Reengineering the Corporation: A Manifesto for Business Revolution*. HarperBusiness, New York.
- Helfat, C. E. 2021. What does firm shaping of markets really mean? *Strategy Science* **6**(4) 360–370.
- Ho, T. K. 1995. Random decision forests. *Proceedings of the Third International Conference on Document Analysis and Recognition*. ICDAR '95, IEEE Computer Society, 278–282.
- Hodges, A. 1983. *Alan Turing: The Enigma*. Simon and Schuster, New York.
- Holland, J. H. 1975. *Adaptation in Natural and Artificial Systems*. University of Michigan Press, Ann Arbor, MI.
- Horton, J. J. 2023. Large language models as simulated economic agents: What can we learn from Homo Silicus? arXiv:2301.07543.
- Iansiti, M., K. R. Lakhani. 2020. *Competing in the Age of AI: Strategy and Leadership When Algorithms and Networks Run the World*. Harvard Business Press, Cambridge, MA.
- Jiao, L., Y. Wang, X. Liu, L. Li, F. Liu, W. Ma, Y. Guo, P. Chen, S. Yang, B. Hou. 2024. Causal inference meets deep learning: A comprehensive survey. *Research* **7** Article 0467.
- Knuth, D. E. 1974. Computer programming as an art. *Communications of the ACM* **17**(12) 667–673.
- Koza, J. R. 1992. *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. MIT Press, Cambridge, MA.
- Levinthal, D. A. 1997. Adaptation on rugged landscapes. *Management Science* **43**(7) 934–950.
- Lieberman, M. B., D. B. Montgomery. 1988. First-mover advantages. *Strategic Management Journal* **9** 41–58.
- Ludwig, J., S. Mullainathan. 2024. Machine learning as a tool for hypothesis generation. *The Quarterly Journal of Economics* **139**(2).
- Luo, Y., J. Peng, J. Ma. 2020. When causal inference meets deep learning. *Nature Machine Intelligence* **2**(8) 426–427.
- McCorduck, P. 2004. *Machines Who Think: A Personal Inquiry into the History and Prospects of Artificial Intelligence*. 25th ed. A. K. Peters, Natick, MA.

- McElheran, K., J. F. Li, E. Brynjolfsson, Z. Kroff, E. Dinlersoz, L. Foster, N. Zolas. 2024. AI adoption in America: Who, what, and where. *Journal of Economics & Management Strategy* **33**(2) 375–415.
- Minsky, M. L. 1986. *The Society of Mind*. Simon and Schuster, New York.
- Mollick, E. 2024. Reinventing the organization for GenAI and LLMs. MIT Sloan Management Review, available at: <https://sloanreview.mit.edu/article/reinventing-the-organization-for-genai-and-llms/> (accessed August 31, 2025).
- Nagel, R. 1995. Unraveling in guessing games: an experimental study. *The American Economic Review* **85**(5) 1313–1326.
- Negroponte, N. 1995. *Being Digital*. Knopf, New York.
- Nelson, R. R., S. G. Winter. 1982. *An Evolutionary Theory of Economic Change*. Harvard University Press, Cambridge, MA.
- Newell, A., H. A. Simon. 1956. The logic theory machine: A complex information processing system. *IRE Transactions on Information Theory* **2**(3) 61–79.
- OpenAI. 2023. GPT-4 technical report. arXiv:2303.08774.
- Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, San Mateo, CA.
- Platt, J. R. 1964. Strong inference. *Science* **146**(3642) 347–353.
- Porter, M. E. 1996. What is strategy? *Harvard Business Review* **74**(6) 61–78.
- Raisch, S., K. Fomina. 2025. Combining human and artificial intelligence: Hybrid problem-solving in organizations. *Academy of Management Review* **50**(2) 441–464.
- Risi, S., M. Preuss. 2020. From chess and Atari to StarCraft and beyond: How game AI is driving the world of AI. *KI - Künstliche Intelligenz* **34**(1) 7–17.
- Schapire, R. E. 1990. The strength of weak learnability. *Machine Learning* **5**(2) 197–227.
- Selfridge, O. 1959. Pandemonium: A paradigm for learning. D. V. Blake, A. M. Uttley, eds., *Proceedings of the Symposium on Mechanisation of Thought Processes*. Her Majesty’s Stationary Office, London, UK, 511–531.
- Shrestha, Y. R., S. M. Ben-Menahem, G. von Krogh. 2019. Organizational decision-making structures in the age of artificial intelligence. *California Management Review* **61**(4) 66–83.
- Silver, D., A. Huang, C. J. Maddison, A. Guez, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, S. Dieleman. 2016. Mastering the game of Go with deep neural networks and tree search. *Nature* **529**(7587) 484–489.
- Simon, H. A. 1955. A behavioral model of rational choice. *Quarterly Journal of Economics* **69**(1) 99–118.
- Spirtes, P., C. Glymour, R. Scheines. 2000. *Causation, Prediction, and Search*. MIT Press, Cambridge, MA.
- Tesauro, G. 1995. Temporal difference learning and TD-Gammon. *Communications of the ACM* **38**(3) 58–68.
- Turing, A. M. 1950. Computing machinery and intelligence. *Mind* **LIX**(236) 433–460.
- Tushman, M. L., P. Anderson. 1986. Technological discontinuities and organizational environments. *Administrative Science Quarterly* **31**(3) 439–465.
- Vaswani, A., N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, I. Polosukhin. 2017. Attention is all you need. *Advances in Neural Information Processing Systems*, vol. 30. Curran Associates, Inc., 6000–6010.

- Wang, S., M. Hu, Q. Li, M. Safari, X. Yang. 2025. Capabilities of GPT-5 on multimodal medical reasoning. arXiv:2508.08224.
- Wingate, D., B. L. Burns, J. B. Barney. 2025. Why AI will not provide sustainable competitive advantage. *MIT Sloan Management Review* **66**(4) 9–11.
- Wooldridge, M. J. 2009. *An Introduction to Multiagent Systems*. 2nd ed. John Wiley & Sons, Chichester, UK.
- Wooldridge, M. J. 2020. *A Brief History of Artificial Intelligence*. Flatiron Books, New York.
- Zellweger, T., T. Zenger. 2023. Entrepreneurs as scientists: A pragmatist approach to producing value out of uncertainty. *Academy of Management Review* **48**(3) 379–408.