THE TK FOUNDRY PAPERS
FOUNDING PAPER · NO. 01

The Forging of *Agent Economics*

A Canadian Perspective on the Next Phase of Artificial Intelligence

STEPHEN THOMPSON
Founder · Thompson Kinkade Foundry
Toronto · 2026

This is the first of The TK Foundry Papers: a perspective series on the economics of intelligent work. It argues that the next phase of artificial intelligence will be decided less by the models we build than by how well we learn to organize what they do, and that Canada is unusually well placed to help build the discipline that measures it.

Every Revolution Creates a New Economics

The tracks mattered. The economic system they enabled mattered more.

In the summer of 1886, the first scheduled passenger train completed the journey across Canada, leaving Montréal on 28 June and reaching the Pacific terminus at Port Moody on 4 July, the year after the last spike was driven at Craigellachie.[1] Politicians spoke of sovereignty. Engineers celebrated one of the great infrastructure achievements of the nineteenth century, almost five thousand kilometres of track laid across mountains, forests, and prairie. They were all correct. But they were also, in a sense, looking at the wrong achievement.

The railway's greatest contribution was not the railway. It was the economy that grew around it. Communities that had existed in isolation became connected markets. Manufacturers reached customers a continent away. Prairie grain reached Atlantic ports. Whole industries emerged for no reason other than that transportation had finally become predictable enough for businesses to reorganize themselves around it. The tracks mattered; the system they enabled mattered more.

History has repeated this pattern with remarkable consistency. Every major technological revolution begins with infrastructure, and very few end there. Electricity transformed the factory long after the generators were built.[2] The internet reshaped commerce long after computers were first connected. Cloud computing changed enterprise software long after virtualization became routine. In each case the first generation celebrated the invention itself; the second generation discovered how to reorganize work around it.[3] The real transformation was never technological alone. It was economic.

Artificial intelligence is now approaching that same transition. For much of the past decade the defining question was whether machines could perform work once thought uniquely human, such as recognizing images, translating languages, writing software, reasoning through hard problems. The progress of recent years has answered a great many of those questions, and that achievement deserves its place in history. But it marks the end of one chapter rather than the completion of the story.

Every important technology eventually reaches the same point, where the invention stops being the central question and the operation becomes the central question. Railways became transportation systems. Electricity became industrial productivity. Cloud computing became operational infrastructure. Artificial intelligence is beginning the same journey, and organizations are discovering that building an intelligent system is only the beginning. The larger challenge is learning to operate it: not once, but every day, at scale, across thousands or millions of decisions.

So the question quietly shifts. It is no longer “Can AI do this?” but “How should we organize AI so that it creates lasting economic value?” That is a fundamentally different conversation. It is no longer about intelligence; it is about economics. And it is already visible across industry. Model capability keeps improving, yet the questions organizations actually ask have changed: how much does an autonomous workflow cost to operate, why does AI spend keep rising even as model prices fall, how should productivity be judged when a single business task may involve hundreds of model interactions, and how should governments measure the return on national investment in sovereign AI infrastructure.

Remarkably, there is still no common discipline dedicated to answering these questions. There are excellent engineering disciplines, excellent financial disciplines, excellent operational disciplines... but not yet a widely accepted economic framework for understanding intelligent work. That absence matters, because every mature technology eventually develops the language required to govern it. Accounting made modern finance possible. Industrial engineering transformed manufacturing. FinOps changed how organizations manage the cloud.[4] Artificial intelligence now stands at the threshold of its own economic discipline.

This paper argues that the next phase of artificial intelligence will be defined not only by more capable models, but by better ways of understanding how those models create economic value. It also argues that Canada has a rare opportunity to help shape that emerging discipline. Not because Canada will build every frontier model, and not because it commands unlimited compute, but because it has repeatedly shown that it can build trusted institutions around transformative technology. The railway did more than connect a country; it created the conditions for a new economy. Artificial intelligence presents a similar opportunity. The models are being built and the infrastructure is expanding. The next challenge is learning how to organize intelligent work. That is where this story begins.

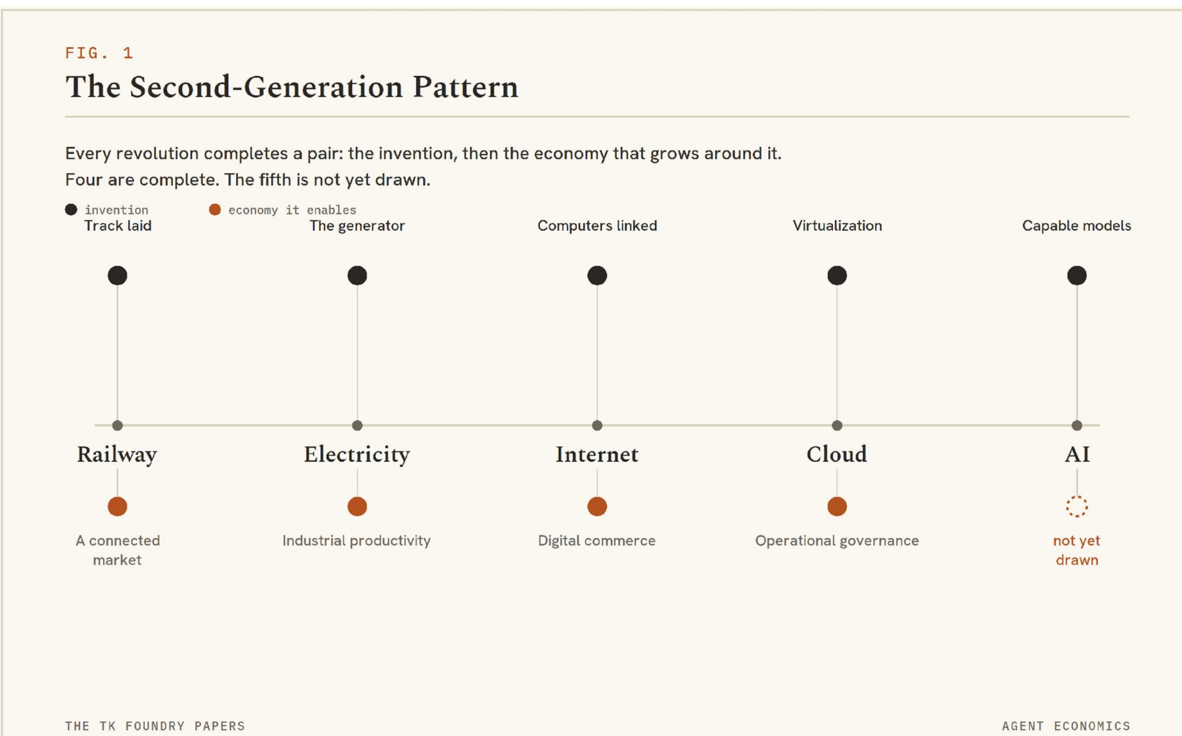

FIG. 1 *Every technological revolution completes a pair — the invention, then the economy that grows around it. AI's second marker is not yet drawn.*

CHAPTER II

THE SHIFT IN THE UNIT OF VALUE

From Intelligence to Intelligent Work

> *The modern enterprise is not asking AI to answer isolated questions. It is asking AI to perform work.*

For much of the past decade, artificial intelligence has been measured by a single question: how intelligent has it become? Every breakthrough was judged through that lens: could a model answer harder questions, write better software, reason more effectively, outperform the previous generation on ever more demanding benchmarks. Those were the right questions, and without the extraordinary gains in capability there would be no foundation on which today's AI economy could rest. But they are no longer sufficient. A remarkable invention does not automatically become a remarkable industry; history suggests something else has to happen first. The invention has to become ordinary.

There is a moment in every technological revolution when the technology itself begins to disappear, not because it matters less, but because it becomes expected.[5] Electricity stopped being an invention and became a utility. The internet stopped being a destination and became the medium through which commerce and communication simply happen. Cloud computing dissolved into the architecture of modern business; organizations no longer ask whether to use the cloud, only how to govern it. Artificial intelligence is beginning the same transition. Increasingly, organizations are not experimenting with AI. They are beginning to depend on it, and an experiment tolerates inefficiency in a way that operations never can.

The public imagination still pictures AI as a conversation: someone asks a question, a model answers, the exchange feels complete. That picture bears less and less resemblance to how AI is actually deployed inside organizations. The modern enterprise is not asking AI to answer isolated questions. It is asking AI to perform work: to review contracts, investigate suspicious transactions, prepare investment briefings, draft procurement documents, monitor competitors, generate software and test it and correct its own errors, and repeat until the job is done.

Behind what looks like a single finished task, the system has retrieved information, weighed alternatives, called external software, checked its own work, retried failed operations, delegated subtasks, and combined intermediate results before returning an answer. This is not simply artificial intelligence. It is organized intelligent work. For generations, computers automated calculation; modern software automated process; artificial intelligence increasingly automates judgement, and organizations are beginning to redesign entire workflows around it.

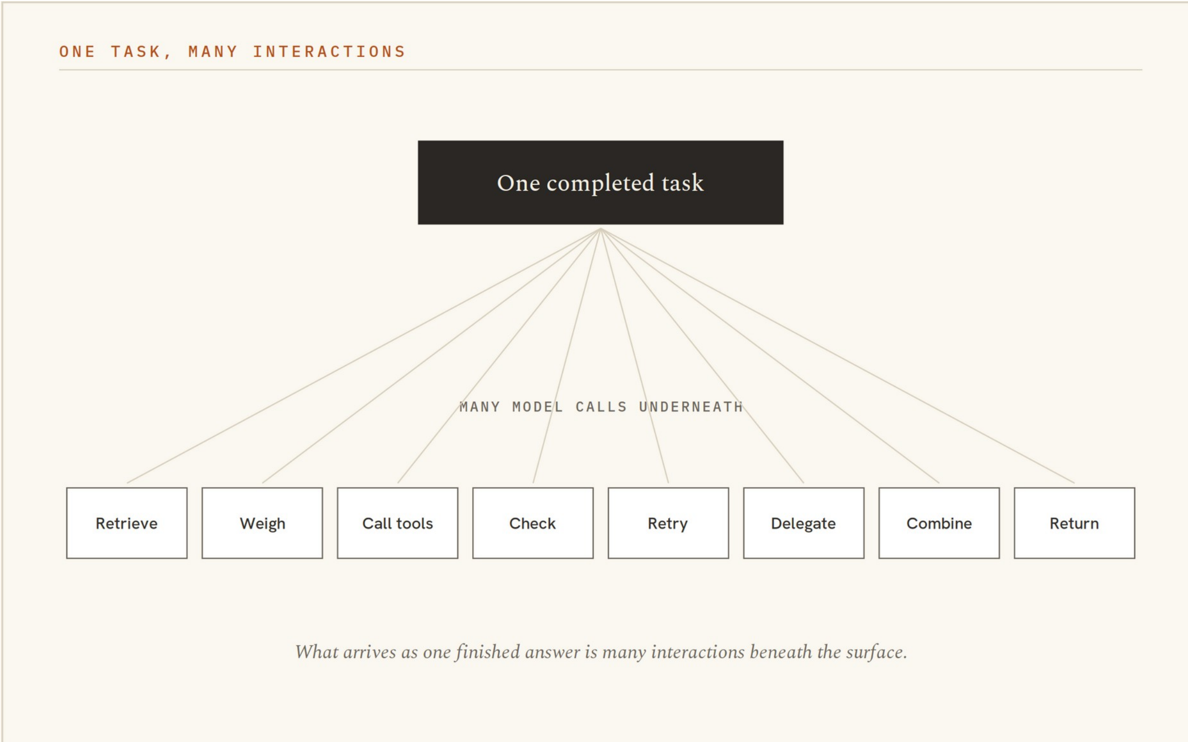

FIG. 2 *What arrives as a single completed task is many model calls underneath — the reason an outcome is not one call.*

The language of AI still revolves around the model: its performance, its size, its price, its benchmarks. Yet organizations rarely buy models for their own sake. They buy outcomes. A bank wants better lending decisions, a hospital better patient care, a government department better public services. The technology remains the mechanism; the work is the objective. Every mature industry eventually discovers that value is created at a higher level of abstraction than the technology itself, and understanding the system of intelligent work (not merely the intelligence within it), is what will determine who realizes lasting advantage.

This is why falling model prices and rising AI budgets are not a contradiction. Cheaper intelligence creates more intelligent work; more intelligent work builds larger operational systems; larger systems raise entirely new questions that are no longer primarily technical but managerial. It is the point at which artificial intelligence quietly leaves the laboratory; not because research ends, but because operations begin. The centre of gravity moves from invention to deployment, from capability to productivity, from isolated demonstrations to continuous operations. The technology remains extraordinary. Its future will increasingly be decided by the quality of the organizations that learn to operate it.

The Economics Inflection

Activity and value are not synonymous. They never have been.

Every mature industry eventually discovers that the first thing it learns to measure is rarely the thing that ultimately matters. Early manufacturers counted machines before they learned to measure output. Telephone companies once counted installed lines; today they measure customers, engagement, and lifetime value. Cloud computing began by counting servers and virtual machines, and mature cloud organizations now measure applications, resilience, and business value. The underlying technology did not become less important. It simply stopped being the most useful way to understand the business.

Artificial intelligence has reached that point. For the past few years its common language has been built around a single unit, the token. Providers price tokens, developers count them, dashboards visualize them, benchmarks compare tokens per second. The token became the natural currency of AI because it describes what a language model actually consumes and produces, and for researchers it remains indispensable. For organizations operating AI at enterprise scale, it is becoming increasingly incomplete.

Imagine the boardroom of a large Canadian bank. The chief financial officer is reviewing quarterly operating costs. No one asks how many database queries were executed, how many packets crossed the network, or how many CPU cycles the infrastructure consumed. Those numbers exist, and operations teams watch them constantly, but executives decide using different measures: revenue, operating margin, risk, productivity, return on capital. As technologies mature, management migrates toward measures that describe outcomes rather than mechanisms. Artificial intelligence has not yet completed that journey.

Consider two organizations that each spend a million dollars on artificial intelligence. On the invoice they look identical. In economic terms they may have almost nothing in common. One automates routine administrative work and frees its people to serve customers; the other runs sophisticated systems that produce impressive demonstrations and little measurable value. Their expenditure is the same. Their productivity is not. The invoice cannot tell them apart, and neither can the token. Activity and value are not synonymous. They never have been.

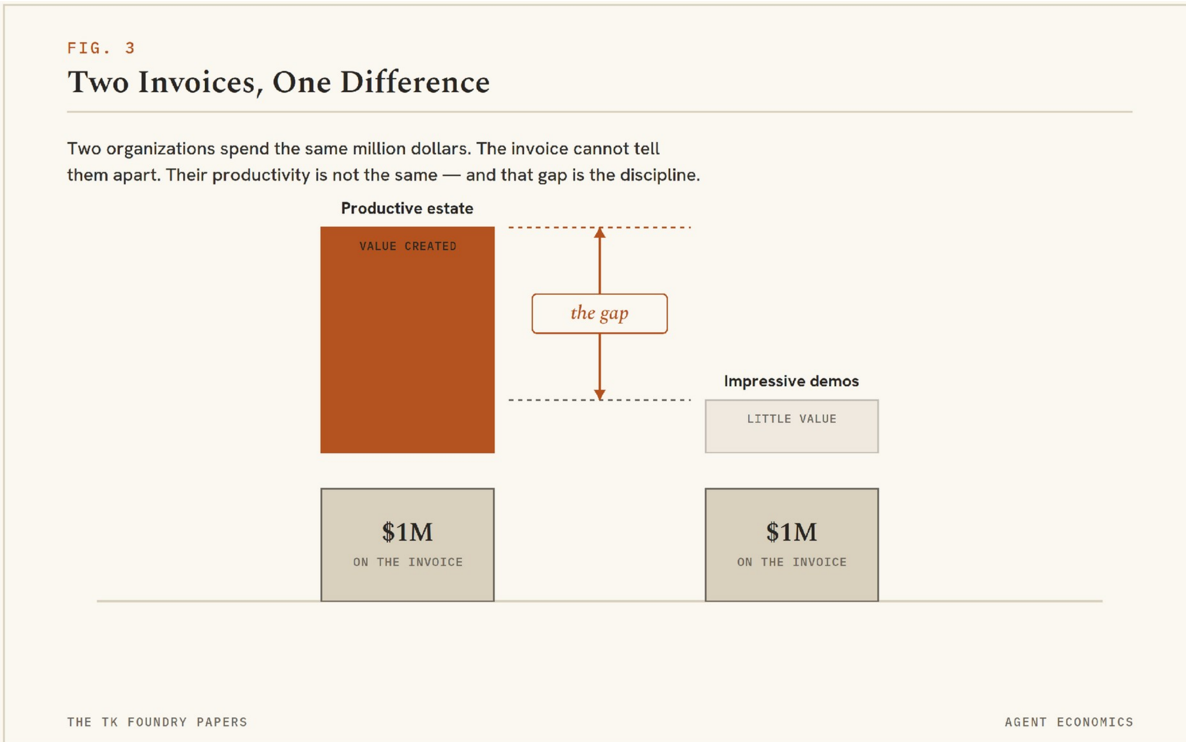

FIG. 3 *Two organizations can spend the same million dollars to entirely different effect. The invoice cannot tell them apart; the gap between outcomes is the discipline.*

The first generation of enterprise AI answered a straightforward question: can this task be automated? The next generation asks something harder: should this task be automated in this way? That question introduces economics, which begins wherever resources are scarce and choices become necessary. Should this workload use a frontier model or a smaller one? Should one intelligent workflow replace three manual processes? Should a public agency buy more compute or use the compute it already has more effectively? These are allocation questions, and allocation has always been the domain of economics. Only a short time ago organizations competed to acquire AI capability. Increasingly, they compete to operate it well.

There is a further pattern worth naming, because it is old and reliable. As the unit cost of a transformative technology falls, society rarely spends less; it simply does more.[6] Electricity grew cheaper and consumption climbed. Computing grew cheaper and organizations processed far more information. Cloud lowered the cost of software and cloud spending kept expanding. Artificial intelligence is following the same course: prices fall, capability rises, and demand grows faster still, because cheaper intelligence widens the range of work worth doing. Every reduction in cost creates new deployments; every deployment adds operational complexity; every added system raises new questions of management.

So the industry must gradually learn to speak a different language. Today it speaks largely in the vocabulary of activity: tokens, requests, latency, inference. Tomorrow it will increasingly speak in the vocabulary of age-old outcomes: completed work, productivity, public impact, competitive advantage.

The economic unit of artificial intelligence is changing: moving away from the individual model interaction and toward the completed outcome. If organizations compete through intelligent work rather than intelligent models, then understanding the economics of that work becomes strategically decisive. The next phase of AI will be defined less by building systems that are more intelligent than by building organizations that understand how intelligence creates value. Everything else follows from that.

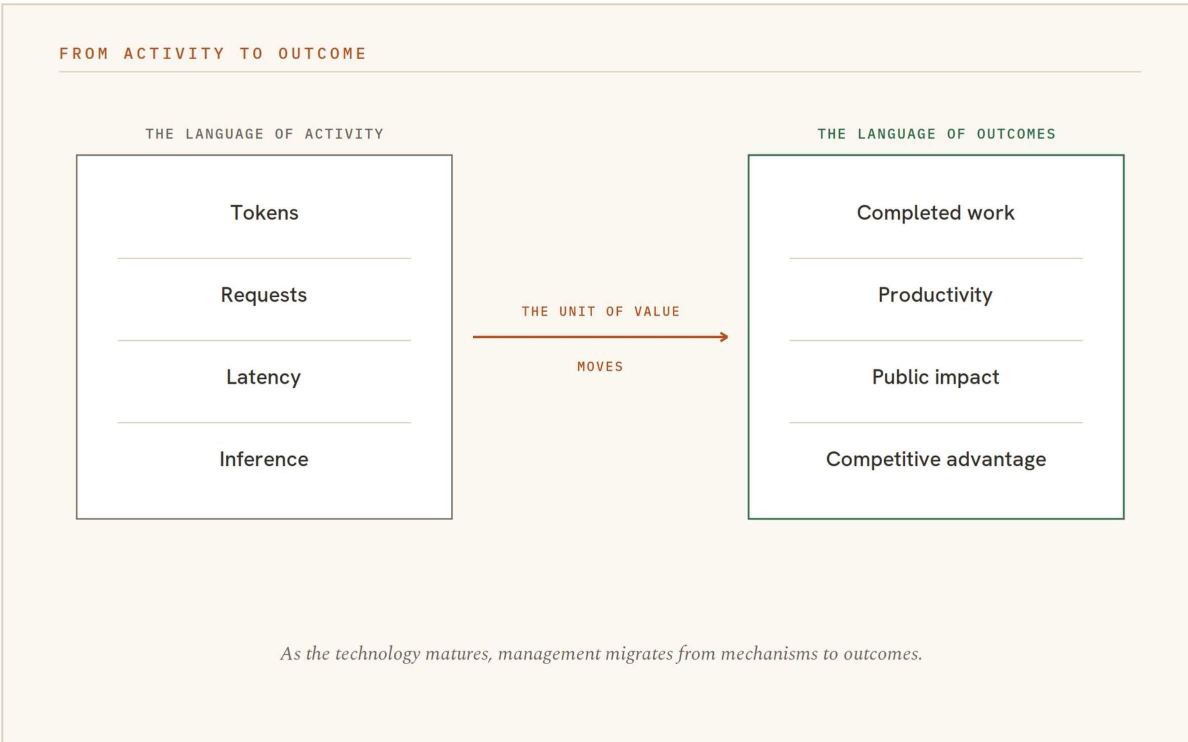

FIG. 4 *As the technology matures, the economic unit moves from the language of activity to the language of outcomes.*

The Forging of Agent Economics

> *The purpose of naming Agent Economics is not to claim ownership. It is to recognize that these questions now belong to a coherent field.*

There is a moment in the evolution of every industry when the existing language begins to lose its usefulness. The words remain accurate; they simply stop being sufficient. Engineers speak one language, finance another, operations a third, and eventually someone notices they are all describing different parts of the same system. A new discipline is born.

Accounting did not emerge because commerce suddenly became complicated; commerce had always been complicated. It emerged because organizations needed a consistent way to understand financial reality. Industrial engineering arrived when factories grew too complex to optimize by craftsmanship alone. Operations research developed when military planners found intuition inadequate to increasingly sophisticated logistics. Cloud computing eventually produced FinOps for the same reason: providers could already issue invoices, engineers already understood utilization, finance already understood budgets. What was missing was a shared framework that let all three discuss the same system in the same economic language.

Artificial intelligence has reached an equivalent moment. Modern AI has produced an abundance of technical measurement, such as latency, inference throughput, context windows, token consumption, model utilization, response quality, benchmark performance. These measures are valuable, but they answer only one class of question: how is the technology performing? Organizations increasingly need the answer to a different question: how is the organization performing because of the technology? Those are not the same question, and the gap between them is where a discipline is needed.

This paper proposes Agent Economics as the discipline that brings them together. The term is not another branch of AI research, nor a synonym for cost optimization. Agent Economics is the study of how autonomous and semi-autonomous AI systems create, consume, and transform economic value — not simply what intelligent systems can accomplish, but how organizations should deploy, govern, and improve them over time. Like every mature management discipline, its purpose is practical rather than theoretical: to improve judgement, to replace intuition with evidence, and to hold apart three pairs that are easily confused: activity from productivity, expenditure from investment, technological capability from economic performance.

Importantly, it begins from a different starting point than most conversations about AI. Most begin with the model; this discipline begins with the organization. Instead of asking which model performs best, it asks which organizational outcome matters most. Instead of asking how many tokens were consumed, it asks how much useful work was completed. Instead of asking how quickly the model responded, it asks how effectively the organization met its objective. Technology remains essential, but it is no longer the centre of the conversation. The object of management is no longer the model. It is the work.

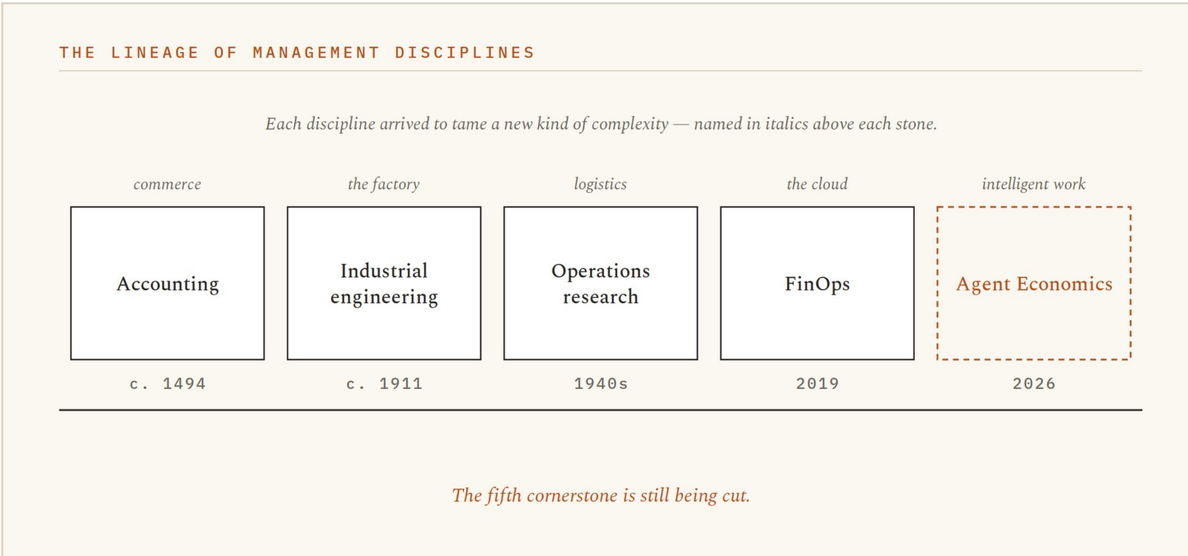

FIG. 5 *Each management discipline arrived to tame a new kind of complexity. Agent Economics is the fifth cornerstone — still being cut.*

Governments, hospitals, universities, manufacturers, and financial institutions all face a version of the same challenge: how should intelligent work be governed — not merely technically, but economically? The answer differs by sector; the objective does not. Deploy intelligence in ways that create more value than they consume. Agent Economics is therefore best understood not as a proprietary framework but as an emerging field of practice, like accounting, operations management, or FinOps. Its methods will evolve, its vocabulary will mature, its measurements will improve, and its standards will emerge through experience. No single organization will define it, nor should one.

Nor does the phrase arrive from nowhere. Researchers have already begun exploring the economics of AI agents in the academic literature — notably a 2025 University of Illinois framework by researchers Ke Yang and ChengXiang Zhai, *Ten Principles of AI Agent Economics*, which examines how autonomous agents make decisions and participate as actors within the broader economy.[7] That work asks an important question: how will intelligent agents behave as economic participants? This paper asks the complementary one: how should organizations measure, govern, and improve the intelligent work those agents perform? The first is a question for economists studying a new kind of actor. The second is a question for managers operating a new kind of workforce. Both belong to the field taking shape, and the fact that scholars and practitioners are converging on it independently is itself the strongest evidence that a discipline is forming.

The purpose of naming Agent Economics is not to claim ownership. It is to recognize that the questions now confronting organizations belong to a coherent field that deserves its own language. History suggests that once a discipline acquires a shared vocabulary, progress accelerates: researchers build on common ideas, organizations compare results, governments set standards, markets form, and knowledge compounds. Artificial intelligence has reached the point where that process can begin. The naming does not end the first chapter of AI. It begins the second. The first chapter asked whether machines could become intelligent; the second asks how societies should organize that intelligence. The first transformed computer science. The second will transform management.

Canada's Opportunity

Sovereign compute creates national capability. Only effective deployment creates national advantage.

Every country participates in a technological revolution. Very few help define the institutions that let it mature. History tends to remember the inventors, but economies are shaped more profoundly by the builders of systems. The United Kingdom gave the world the Industrial Revolution, yet its enduring influence came as much from modern finance, insurance, and commercial law as from the steam engine. The United States transformed the digital age not only through semiconductors but through the capital markets and entrepreneurial institutions that let those technologies scale. Switzerland became indispensable to global finance without being a large economy; Singapore became one of the world's great trading nations despite its geography, because it built trusted institutions around commerce. Their advantage was not merely technological. It was organizational.

Canada now faces a similar moment. Much of the global conversation about AI has understandably centred on the race to build ever more capable frontier models. That race matters, and it is becoming one of the most capital-intensive industrial competitions in modern history. For most countries, attempting to replicate it is neither practical nor strategically necessary because technological leadership has rarely depended on inventing the foundational technology. More often it has belonged to those who understood how that technology reshapes industries.

Canada has repeatedly shown that kind of leadership. Its universities have produced globally respected researchers, and institutions such as the Vector Institute, Mila, and Amii have helped make the country one of the world's centres of AI scholarship. Just as important are the strengths that attract less attention: stable public institutions, trusted financial systems, strong regulatory traditions, deep expertise in highly regulated industries, and a long habit of balancing innovation with public stewardship. Those are economic assets. The rise of intelligent work places extraordinary demands on trust and organizations need confidence that intelligent systems can be measured, governed, and improved through evidence rather than optimism. Those are institutional requirements as much as technological ones.

Canada's recent investment in sovereign AI infrastructure illustrates the opportunity precisely. The national conversation has understandably focused on computational capacity, and that capacity matters. But infrastructure creates possibility; it does not guarantee productivity. Every railway needed businesses able to use it. Every electrical grid needed industries willing to reorganize around electric power. Every cloud platform needed organizations that learned to govern computing differently. Artificial intelligence follows the same pattern. Sovereign compute creates national capability. Only effective deployment creates national advantage.

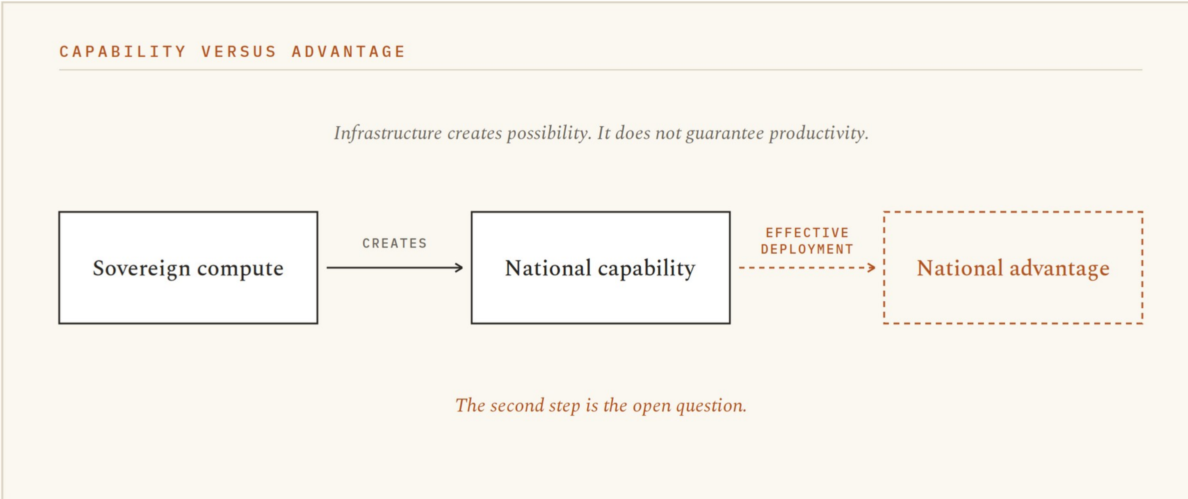

FIG. 6 *Sovereign compute creates national capability. Advantage is the open question that effective deployment must answer.*

So rather than competing on computational scale alone, Canada can compete by becoming a trusted leader in the economics of intelligent systems. The most influential frameworks emerge through openness, evidence, and practical experience, and Canada has long excelled at convening researchers, governments, and industry around shared challenges. Artificial intelligence offers another occasion to do exactly that; not because Canada holds every answer, but because it is well placed to ask the right questions. How should intelligent work be measured? How should autonomous systems be governed? How should organizations tell experimentation from productivity? How should governments evaluate the return on public investment in AI infrastructure?

These questions concern the future organization of work itself, and so the opportunity before Canada is larger than building another successful technology company. It is helping to establish the intellectual infrastructure on which an emerging economy will depend. Capital markets required accounting. International trade required commercial law. Cloud computing required operational governance. The age of intelligent work will require its own foundations, and Canada is unusually well positioned to help build them; not because of geography or scale, but because it has repeatedly shown that thoughtful institutions can become enduring competitive advantages. Countries will keep competing to build more capable models. A quieter competition has already begun: the competition to understand how intelligence creates value once it becomes part of everyday economic life. Canada has earned the opportunity to take part in that conversation. The question is whether it chooses to lead it.

CHAPTER VI

AN INVITATION, NOT A CLAIM

Building the Discipline

Disciplines are not declared into existence. They are built.

Every generation inherits technologies that redefine what is possible. Far fewer inherit the chance to define how those technologies are understood. Artificial intelligence has reached that moment. The science has advanced at extraordinary speed; systems that seemed remarkable a few years ago are now widely accessible, and organizations in every sector are experimenting with intelligent systems that assist professionals, automate routine work, and increasingly take part in decisions once reserved for human judgement. That progress is undeniable. It is also incomplete. History suggests that every technological revolution succeeds or fails not on the quality of its inventions but on the quality of the institutions that grow around them. Technologies create possibility. Disciplines create progress.

This paper has argued that artificial intelligence is entering its second phase. The first phase asked whether machines could become intelligent; the second asks how organizations should use that intelligence: responsibly, productively, and sustainably. Those questions cannot be answered by engineers alone, nor by economists, governments, or software companies on their own. They require a common language capable of connecting all four. That is the role Agent Economics seeks to play: not as a proprietary framework and not as a commercial category, but as a management discipline for an economy in which intelligent systems increasingly become participants in everyday work.

Like every discipline before it, Agent Economics will not arrive fully formed. It will evolve. Its vocabulary will change, its measurements will improve, some assumptions will prove incomplete, and others will settle into accepted practice, because disciplines mature through evidence. If this paper has one central argument, it is a simple one: the next decade of artificial intelligence will be defined less by the intelligence we create than by the wisdom with which we organize it. History rarely rewards societies that merely invent powerful technologies. It more often rewards those that learn to govern, measure, and improve them. Industrial engineering transformed manufacturing by teaching factories to be productive. Modern accounting transformed commerce by letting organizations see financial reality. FinOps transformed cloud computing by giving enterprises a shared language for operational complexity. Artificial intelligence stands at a similar threshold.

Canada enters this period with unusual advantages. It has helped shape modern artificial intelligence through research, education, and public investment, and it has earned international credibility through institutions that balance innovation with stewardship. Those strengths position it to contribute something the world increasingly needs. That is not simply another technology, but a way of thinking: a discipline, a shared understanding of how intelligent work creates economic value. Disciplines are not declared into existence. They are built.

For Thompson Kinkade Foundry, this paper marks the beginning of that work. The papers that follow in this series will take up the practical questions introduced here, namely, how organizations should measure intelligent work, how operational efficiency should be understood, where AI waste originates, how governments should benchmark national capability, and how enterprises should compare performance across increasingly autonomous systems. These are questions worth pursuing seriously, and our contribution will be to pursue them with intellectual honesty, operational evidence, and an unwavering commitment to separating observation from assumption. Ideas become valuable when they are useful to others.

There is a tendency, in periods of rapid change, to believe that the defining achievements are the inventions themselves. History offers a quieter lesson: the inventions capture our attention, but the disciplines shape our future. The first generation of artificial intelligence asked whether machines could think. The second asked whether they could reason. The next will ask a different question: whether societies can organize intelligence well enough to create enduring economic value. That question belongs to no single company, university, or government. It belongs to all of us. This paper has argued that it deserves a name, and has offered one: Agent Economics. The work of answering it has only just begun.

REFERENCES

Footnotes

1 Historical dates regarding the Canadian Pacific Railway are a matter of public record. The last spike was driven at Craigellachie on 7 November 1885, and the first scheduled passenger train arrived at Port Moody on 4 July 1886.

2 The historical lag between the invention of generators and the reorganization of factories was definitively explored by economic historian Paul A. David in "The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox," *The American Economic Review* 80, no. 2 (1990).

3 This framework of technological revolutions unfolding in two phases — an installation period centred on the invention, followed by a deployment period in which the economy reorganizes around it — draws on Carlota Perez, *Technological Revolutions and Financial Capital* (2002).

4 "FinOps" is a trademark of the FinOps Foundation, a project of the Linux Foundation. Its work in establishing a shared operational language for cloud computing serves as a historical model for the discipline of Agent Economics.

5 The observation that mature technologies recede from attention as they become expected reflects Mark Weiser's foundational writing on ubiquitous computing. See Mark Weiser, "The Computer for the 21st Century," *Scientific American* (1991).

6 This pattern — where falling unit costs expand rather than shrink total consumption — was first described by William Stanley Jevons in *The Coal Question* (1865), and is widely known as the Jevons paradox.

7 Ke Yang and ChengXiang Zhai, "Ten Principles of AI Agent Economics" (University of Illinois Urbana-Champaign, arXiv:2505.20273, May 2025). This paper regards their academic treatment of AI agents as economic actors as a complementary contribution to the same emerging field.

Any errors of interpretation regarding these foundational frameworks are the author's own.

THE TK FOUNDRY PAPERS
Founding Paper No. 01 · The Forging of Agent Economics · Thompson Kinkade Foundry Inc., Toronto · 2026. A perspective paper published to open a field of practice; it advances no proprietary claim over the discipline it names. · tkfoundry.ai