RBC
Thought
Leadership

AI CASE STUDY

Turning Disruption into Momentum

Manulife's AI Flywheel

Artificial intelligence is poised to reshape how value is created across Canada's economy. To understand that shift, RBC Thought Leadership interviewed more than two dozen firms that are on the front lines of building or deploying AI for our report, *Bridging the Imagination Gap: How Canadian Companies Can Become Global Leaders in AI Adoption*.

Building on that report, our series of case studies goes a level deeper. Here we follow how Manulife, a global insurer and asset manager, used generative AI as a catalyst to rethink how the organization learns, shares, and scales new ideas. The company's experience shows that successful AI adoption is not a technology challenge alone—it's a challenge of capability-building, governance, and empowering people to work differently.

Summary

Manulife, a global asset manager headquartered in Canada, saw AI as a chance to move beyond incremental efficiency gains and reimagine products and operations. Leadership judged the sector "too comfortable," set a clear ambition to become a digital-customer leader, and treated OpenAI's Large Language Model in 2022 as a tipping point. A hands-on executive session turned AI from a niche experiment into a CEO-level agenda item, signalling that real impact would require structure, governance, and integration—not one-off pilots.

- **Build absorptive capacity (infrastructure).** Manulife created a multi-tier learning stack and embedded ~200 data science and machine learning experts, and used leadership rituals to grow the "stock of prior knowledge," so new AI advances could be absorbed and embedded faster.

- Institutionalize adaptive capacity (the engine). Leaders normalized copying—if one team built something useful, others reused it. This turned isolated wins into shared playbooks and spread improvements quickly. By embedding that habit, Manulife accelerated the cycle of adopt, invent, select, scale, building adaptive and innovative capacity together.

- **Balance speed and safety (govern ance by outcomes).** Responsible AI principles, expanded model-risk frameworks, cross-functional review, and real-time telemetry treated fast iteration and strong oversight as complements, not one-off pilots

It was mid 2020. Jodie Wallis, then Manulife's Global Chief Analytics Officer, had summoned the company's top executives into a Toronto boardroom. She knew the meeting would mark a turning point: OpenAI's breakthrough latest large language model (LLM), GPT -2 had just been released, and, at nearly 100 times stronger than its previous models, GPT-2's implications stretched far beyond the technology itself. For Manulife, a 137-year-old insurer built on actuarial precision and risk discipline, the question was whether this new capability would be treated as a passing novelty, or as the spark for deeper change.

For years, AI at Manulife meant prediction and automation—underwriting models, fraud detection, lead scoring. Even as the frontier advanced with machine-learning models that could conjure hyper-realistic images, these applications still felt contained within the realm of "computer things." They were useful and very impressive but safely bounded by expectation.

To Wallis, large language models like GPT shattered those boundaries. Designed for an iterative exchange, they created value not through a single output but through an unfolding dialogue—shifting the dynamic from command-and-response to something closer to collaboration. LLMs could now reason with a human-like cadence, inviting conversation rather than instruction. The breakthrough was not a more polished "answer," but the model's ability to so fluidly augment inquiry itself—generating new directions of thought and discovery.

That shift—from bounded tasks to open-ended discovery—was as unsettling as it was exhilarating. Wallis framed the moment with unusual candor: "Our industry has been too comfortable. This technology isn't just another tool—it's a fork in the road. We either harness it, or risk being reshaped by it."

Around the table, reactions varied: curiosity, excitement, apprehension. The challenge was immediate. Should Manulife treat generative AI as an experiment at the margins, or as the new trajectory of the business itself? Wallis herself was convinced of the answer, but she also knew the technology was still raw—too raw, perhaps, for the boardroom to fully accept. The choice would force hard calls about strategy, governance, culture, and investment, all at the breakneck pace at which the frontier was advancing.

In such moments of technological upheaval,

corporate boards look to figures like Wallis to distinguish passing trends from transformative forces. Unlike the technologist-soothsayers popular at the time, her task was consequential: to foresee how generative AI might reshape an institution built on actuarial discipline, and to ensure Manulife seized the opportunity rather than being undone by it. Frame the moment correctly, and new value could be unlocked; misjudge it, and the consequences could be existential.

But foresight alone would not suffice. Wallis knew no memo or slide deck could capture the implications of generative AI; words on a page risked being dismissed as abstractions. The only way forward was direct confrontation. To overcome that gap, one had to experience it themselves. Fortunately, the technology itself offered an answer—the opportunity to turn the crystal ball around and let skeptical peers glimpse inside for themselves.

So, she placed a tablet in front of each leader, preloaded with the latest OpenAI model, and invited them to test it—to ask it the questions they might otherwise have asked her. The room fell silent as screens lit up with blinking prompts. One by one, Manulife's senior leaders began conversing with GPT-2, watching as it generated fluent answers in real time. The exercise was disarmingly simple, yet it shifted the atmosphere. Within minutes, the conversation had moved from "is this real?" to "what does this mean for us?"—the kind of pivot that months of memos and meetings could never have achieved.

It was Wallis's decision—to make her colleagues experience the frontier for themselves—that created conviction at the top. But she knew conviction alone would not be enough. To matter, it had to be built into infrastructure, and then into the agility to adapt. With that boardroom experiment, Wallis set the flywheel in motion—conviction, infrastructure, adaptation—that would carry Manulife through one of the most profound technological shifts in its history. In doing so, Manulife joined a small group of financial giants positioning Canada at the forefront of AI transformation.

To understand how this journey unfolded, RBC Thought Leadership sat down with Jason MacDonald, Chief of Staff in the Office of the CEO, and Jodie Wallis—now the company's Global Chief AI Officer—to explore how they and their colleagues steered a $72-billion insurer through one of the most profound technological shifts in its history.

Insight 1 → Show, Don't Tell: Creating Conviction at the Top

Strong buy-in from senior executives is critical at the beginning of any transformative initiative. Wallis understood that leaders had to experience AI directly for themselves. In doing so, she was putting into practice what Everett Rogers' diffusion theory had long shown: new ideas spread faster when they are trialable—safe to experiment with in low-risk conditions—and observable—when peers can see results firsthand. Together, these conditions turn abstract technology into something tangible enough to believe in.

That is exactly what unfolded in the boardroom. Once a few respected voices found the tool useful—asking follow-ups, reading fluent outputs aloud—trialability was satisfied: executives could experiment in a low-stakes, hands-on way. And because these experiments happened in public, observability took hold: colleagues could watch, compare reactions, and see the system working in real time. What could have been a solitary experiment quickly became a shared moment of discovery. Peer-to-peer reinforcement allowed skepticism to fall away and curiosity to spread, because the technology no longer seemed risky or abstract.

But conviction alone is not enough. To matter, it had to be translated into infrastructure that would let Manulife absorb and scale what leaders had seen. That is where absorptive capacity comes in.

Insight 2 → Absorptive Capacity: Building Knowledge Infrastructure

A single demo, however persuasive at the individual level, fades unless an organization as a whole can metabolize what it saw into repeatable capability. That is the job of absorptive capacity—a firm's ability to recognize the value of new information, assimilate it, and apply it to commercial ends—the infrastructure that makes later adaptation possible. Research on absorptive capacity, first developed by professors Wesley Cohen and Daniel Levinthal in the 1990s, highlights two foundations of that infrastructure:

1. **Knowledge is cumulative and path-dependent—**it builds fastest on what people already know, meaning prior knowledge is like scaffolding for future learning.
2. **Breadth of knowledge expands absorptive reach—**organizations with a wide base of prior knowledge can take in and apply new external ideas more effectively.

Absorptive capacity is about learning—building the knowledge base and routines to embed new tools. Adaptive capacity (discussed in Insight Three) is about changing—reconfiguring those routines when the frontier shifts and old paths no longer fit. Manulife needed both, but it started by deliberately building the absorptive infrastructure needed to allow the organization to learn. In doing so, Wallis's team treated culture and skills as equal pillars to technology and designed a multi-tier learning stack:

- AI 101 for anyone with an interest,
- advanced prompt-engineering and data-science for power users, and.
- tailored executive modules delivered with university partners.

They then wove AI into leadership rituals. At Manulife's Global Leadership Conference, for example, executives showcased employee-built solutions to their peers, creating a common language of use cases and governance. The goal wasn't just awareness; it was to give every layer of the company—front line to boardroom—enough context to recognize where AI was relevant and embed it in daily work.

In Cohen and Levinthal's terms, Manulife was steadily increasing its stock of prior knowledge, so each new wave of technology could be absorbed and recombined faster. Wallis's actions directly aligned with the two conditions they described: training and rituals made learning cumulative by building on what employees already knew, and broad participation across the workforce expanded the base of knowledge available to draw on. In an industry often criticized as "too comfortable," this gave Manulife a distinctive edge: the ability to build on new tools and embed them into its routines in ways that accumulated advantage over time.

But infrastructure alone is not enough. Once that foundation was in place, the challenge became keeping momentum when the frontier shifted and old paths no longer fit. That required a different capability: adaptive capacity—the engine that keeps the flywheel turning.

Insight 3 → Adaptive Capacity: The Innovation Flywheel

When then-CEO Roy Gori warned that the industry had grown "too comfortable," Wallis knew this complacency was dangerous in a domain where new AI models and applications were appearing at a breakneck pace, driven by massive new capital flows. Absorptive capacity had already given Manulife the infrastructure to learn and embed AI tools across the enterprise. The next challenge was agility: ensuring the company's

Numbers

- **$1.6T**
  Assets under management
- **35M**
  Customers worldwide
- **$53B**
  Market **Capitalization**
- **$5.1B**
  Net **Income**
- **38K**
  Number of **employees**
- **200**
  Data scientists and engineers embedded Across teams
- **$600M**
  Benefits attributed to digital transformation (with AI as a core part) in 2024.
- **75+**
  AI use cases deployed by end of 2025
- **75%**
  Share of Manulife's global workforce engaged with GenAI

response to advancing technology was equally swift and dynamic. Adoption couldn't be a one-off event; it had to become iterative. That insight set the stage for adaptive capacity—the engine that converts adoption into continuous reinvention.

Research underscores why this engine is critical. Prior adoption experience is the single strongest predictor of inventive capacity: organizations learn to invent by first copying. Yet when firms switch paths—moving to new models or methods —performance often dips before it recovers, as old mental models stop fitting the new approach. Adaptive capacity is therefore the discipline of riding out that trough and recovering faster, turning temporary disruption into cumulative learning. Manulife operationalized this discipline through a set of deliberate routines.

- **Adoption →** taking in new tools, practices, or patterns developed elsewhere, and embedding them into the organization's routines.
- **Selection and Scale →** filtering what works, embedding it into routines, and scaling proven solutions across the enterprise.
- **Invention →** creating original solutions internally, without relying on external patterns.

Manulife built this discipline deliberately. With a strong foundation of AI literacy embedded across the company, leadership worked to smooth adoption pathways by normalizing copying as a precursor to invention. Wallis instituted prompt-a-thons and leadership conferences where employee-built tools were showcased, creating a common language of value and risk. These rituals made it legitimate to borrow, refine, and scale what worked—ensuring adoption wasn't confined to early enthusiasts but cascaded across the enterprise. In Cohen and Levinthal's terms, this was about continuously increasing the firm's stock of prior knowledge so that when a path switch came—whether a new model, platform, or application—the organization could absorb and apply it faster.

Secondly, Wallis deliberately designed for safe path-switching. A vendor-agnostic, cloud-ready stack allowed models to be swapped 'even daily,' making technology change a managed routine rather than a disruptive reset. Scaling decisions were tied to clear business outcomes—revenue lift, cost savings, risk reduction, or productivity—so that pivots created value rather than noise.

Finally, it embedded selection capacity—the discipline to prune weak ideas quickly and scale winners. Cross-functional forums and outcome-based funding kept the portfolio focused, so absorptive capacity compounded rather than leaked.

Together, these routines formed Manulife's innovation flywheel: adoption experience generated invention; selection routines filtered the noise; flexible architecture enabled safe path-switching; and the loop restarted with each cycle stronger than the last.

> **"This wasn't a tool to give answers; it was a tool to improve the questions we asked."**
>
> —Jodie Wallis, Global Chief AI Officer, Manulife

Adaptive Capacity: The Innovation Flywheel

The spiral diagram shows Manulife's **adaptive capacity.** Each loop of **Adoption → Invention → Selection & Scale** expands embedded knowledge. When a **path-switch** happens (new model or method), performance dips briefly, then slingshots into a stronger cycle. What matters is riding out the trough and carrying momentum into the next adoption loop. Progress isn't linear—it's a **flywheel that turns disruption into momentum**.

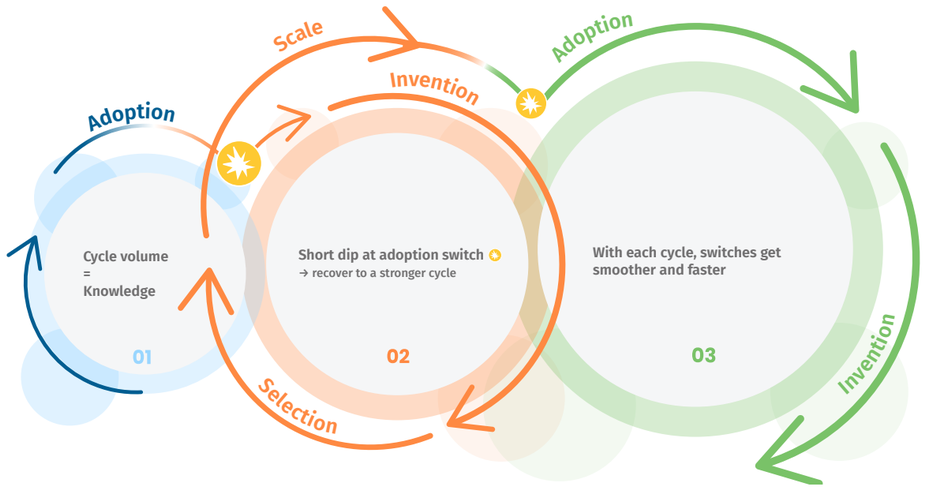

Insight 4 → Responsible Governance: Balancing Speed and Safety

From the outset, the company made responsible AI governance a design choice. In the absence of clear national rules, it created its own responsible AI principles and operating rules to ensure experimentation and deployment stayed aligned with ethical, privacy, and compliance obligations.

Manulife expanded its existing model risk frameworks to address GenAI's unique challenges—vetting third-party vendors, monitoring outputs for bias or hallucinations, and requiring ongoing performance assessments for every model in production. A cross-functional governance committee reviewed use cases for ethical and privacy risks, aligning policies with evolving global guidelines. Governance was embedded as a living process, not a static policy.

Critically, Manulife treated fast iteration and strong oversight as complements, not trade-offs. Continuous model monitoring—tracking accuracy, drift, and usage—was used to tighten controls in real time. This outcome-based approach allowed models to stay in production as long as they met error and bias thresholds, and to be adjusted or pulled the moment they didn't. Iteration was welcome, but never at the expense of trust.

This proactive stance enabled Manulife to scale GenAI quickly and responsibly, building confidence with compliance teams, customers, and policymakers, even in the absence of clear regulation. The broader lesson is that firms in sensitive sectors should not treat regulation as a brake. By self-imposing principles, operationalizing oversight, and demonstrating to regulators that innovation can be pursued responsibly, companies can get ahead of uncertainty. For policymakers, the takeaway is equally important: enabling real-time oversight and outcome-based guardrails may achieve safety faster than prescriptive, one-off compliance checks.

> **"Instead of posing a problem for the machine to solve, humans could now use technology to embark on a process of discovery together, where the real value lay in the insights that unfolded along the way."**
>
> —Jodie Wallis, Global Chief AI Officer, Manulife

Results: From Experiment to Enterprise Transformation

Within just a year of embracing generative AI, Manulife achieved broad-based adoption at a speed few incumbents match. Its proprietary assistant, ChatMFC, went from pilot to near ubiquity: within months, 40% of employees were using it monthly, and by early 2025, more than 75% of the global workforce was actively engaged with GenAI tools, training, or use cases. Adoption was not siloed to tech teams; it touched nearly every function, from sales and service to back-office operations.

The impact on productivity was equally striking. In call centers, AI tools shaved 30 – 40 seconds off average call times without lowering customer satisfaction. Across the enterprise, generative AI was no longer a side project—it had become embedded in the daily flow of work.

Customer-facing gains were even more visible. Newer advisors ramped up faster, using AI coaching to practice and refine interactions. Meanwhile, advisors reported that AI freed them to focus on client relationships, creating the unusual outcome of a technology initiative that delivered both efficiency and deeper human engagement.

At the strategic level, the flywheel was spinning. By mid-2025, Manulife had 35+ GenAI use cases in production and 70 more in queue. Early deployments alone contributed an estimated $4.7 million in benefits, while the broader digital transformation program (with AI at its core) yielded over $600 million in 2024 benefits—savings, new sales, and better risk outcomes. Looking ahead, the company projects a threefold return on AI investments over five years. These results affirm that Manulife's design choices — hands-on executive engagement, outcome-gated scaling, perpetual-beta governance—transformed AI from novelty to institutional capability.

Endnotes

[1]Generative Pre-training Transformer

Project Lead

Reid Mckay, Director, Technology Policy, RBC Thought Leadership

Contributors

Caprice Biasoni, Design Lead, RBC Thought Leadership
Nora Bieberstein, Director, Strategic Programs, RBC Thought Leadership
Yadullah Hussain, Managing Editor, RBC Thought Leadership
John Intini, Senior Director, Editorial, RBC Thought Leadership

Special thanks to **Janice Stein**, Founding Director, Munk School of Global Affairs & Public Policy and **Nicole Harris**, Graduate Student, Munk School of Global Affairs & Public Policy.

Published by

RBC Thought Leadership

December 2025