Issue No. 28, Fall 2026

ThinkTWENTY20

The Magazine for Financial Professionals



From Chatbot to Co-worker
When AI Starts Doing the Work: What Changes for Accountants?
Moving Accounting Education Forward in the Age of "Thinking" AI
From Sustainability to Systems Thinking
Limitless Us: A Book Review

Number 28, Fall 2026

Editor in Chief: Gerald Trites, Managing Editor: Gundi Jeffrey, Contributing Editor: Eric E Cohen

Email: info@thinktwenty30.com Telephone: (416) 602-3931

We are a community of people who want to learn about, and discuss, important issues. To join us, please go to our website at https://www.thinktwenty20.com. Regular subscription – free. Premium subscription, One year - $30, Two years - $50.

ISSN 2563-0113

Cover picture; Pixabay.com

FOUNDING PARTNER

CPA CHARTERED PROFESSIONAL ACCOUNTANTS CANADA COMPTABLES PROFESSIONNELS AGRÉÉS CANADA

Supporter

Table of Contents

Editorial……..Page 3
By Gerald Trites, FCPA, FCA, Editor-in-Chief

From Chatbot to Co-worker……...Page 4
By Eric Cohen, CPA
The first phase of mass-market generative AI was largely about generation: what could the machine write, summarize, draw or create? The next phase is participation: what can it observe, what systems can it access, what tasks can it perform and what authority may it exercise? A third dimension is now becoming unavoidable: what happens when that participation persists after the person who first commissioned it is no longer there?

When AI Starts Doing the Work: What Changes for Accountants? …….Page 25
By Cory Ng, CPA
Accountants remain responsible for the work they approve. That responsibility has always required an understanding of the evidence and procedures behind the result. Agentic AI complicates supervision when the system selects among approved steps before a person reviews the outcome. Accountants need to know what information was used, how the system handled obstacles and which exceptions still require attention.

Moving Accounting Education Forward in the Age of "Thinking" AI…..Page 30
By Irene Wiecek, FCPA, FCA
We are in the midst of a paradigm shift, affording an opportunity to reconceptualize how to help accounting students grow, mature and become competent so that they can lead the way forward. Many universities and colleges are well down the path regarding reinventing their programs and courses, while others are not.

Beyond Sustainability to Systems Thinking……...Page 41
By Richard Spencer
If business is going to play its part in achieving the goal of a sustainable planet with prosperous economies and thriving societies – if it is going to connect everyday business decisions with big systemic risks – then its tool kit is inadequate. We need to develop systems thinking as a core skill in the professions and amongst business managers and a systems-mindset as the cultural norm.

Limitless Us: A Book Review……..……...Page 52
By Robert Edison Sandiford
This review of *You to the Power of Two: Redefining Human Potential in the Age of Identic AI*, by Joseph M. Bradley and Don Tapscott, makes very clear that the presence and use of artificial intelligence generate so many prickly questions.



Gerald Trites, FCPA, FCA
Editor-in-Chief

Artificial Intelligence (AI) has become so important to business and our profession that it dominates many major discussions about both future and current issues. Therefore, we decided that all the articles in this issue would address AI and its implications.

Although there are those who think we should stop this freight train, in fact we can't. In the past few years, AI has penetrated to the core of business activity and it will continue to grow and evolve.

Fears about AI taking over control of the world are probably overblown. But fears of AI taking over corporate governance have become critical. Companies that adopt AI in important segments of their business without control are assuming a great, perhaps unacceptable, risk.

In dealing with such a fast-moving train, it is difficult to establish relevant and useful standards. Nevertheless, a great deal has been done, such as ISO 42001, which establishes a framework for AI control systems. Other organizations have done similar work, and all the big firms have provided guidance in their publications on the subject and offer related consulting services.

The hitch is that AI is changing so quickly that specific control procedures must change with the particular AI systems being used and the new ones being adopted. Since these elements are changing with lightning speed, it's very difficult – if not impossible – to keep up.

During these times, keeping AI under control is a major challenge for business and even for society generally. Adopting a good control framework is a start.

From Chatbot to Coworker: When AI Joins the Firm

By Eric E. Cohen, CPA



Eric E Cohen, CPA, is a technologist with a passion for collaboration toward the goal that "a piece of business information, once entered into any system, anywhere, never needs to be retyped as it moved through the business reporting supply chain." He's also a prolific author, engaged in virtually every effort to standardize accounting and audit data, a national expert to a wide variety of standards efforts, and co-founder of XBRL.

I began writing this article on August 17, 2026. It was already difficult to describe generative AI without making something sound obsolete by publication. Revisiting it in early November feels like preparing a subsequent-events disclosure: the underlying story did not disappear, but several events materially changed the picture.

In a matter of weeks, Meta introduced Muse, a personal agent that works in its own cloud environment (September 8); OpenAI introduced dots, intended to work persistently across connected tools (September 29); Google outlined Gemini coworker agents with their own Workspace identities (October 8); and Anthropic opened a research-preview waitlist for Claude Duet (October 21). These announcements are not proof that autonomous digital employees have arrived fully formed. They are evidence that vendors are now designing for something beyond a single prompt and response.

Some names and availability details will undoubtedly change again. That does not undermine the argument. It illustrates how risky it is to build a professional or organizational strategy around this quarter's product names.

For several years after ChatGPT arrived, much of the popular discussion of generative AI centred on a deceptively simple question: What can the latest model do? A more productive question today is: What kind of participant is the AI becoming in our work?

It can still be a chatbot, researcher, image generator or coding assistant. It can also be a spreadsheet capability, an API, a desktop application, a cloud computer, a connected automation, a scheduled worker or a persistent agent that approaches us with work already under way. The model is only one component of that experience.

The first phase of mass-market generative AI was largely about generation: what could the machine write, summarize, draw or create? The next phase is participation: what can it observe, what systems can it access, what tasks can it perform and what authority may it exercise? A third dimension is now becoming unavoidable: what happens when that participation persists after the person who first commissioned it is no longer there?

What This Article Is Really About

This article is not intended to provide a definitive ranking of today's models. Any such ranking would probably be out of date before this issue reaches its readers. Nor is it an attempt to provide a comprehensive history of artificial intelligence.

Rather than ranking models, I want to examine the developments that seem more durable than model numbers. Our questions about anthropomorphism, misplaced reliance, conflicting objectives and the relationship between people and machines predate modern machine learning. What ChatGPT changed dramatically was access: capable AI became available through a conversational interface that almost anyone could use.

We have moved well beyond conversation. Systems create across multiple media, use external tools and organizational data, operate applications and increasingly pursue multi-step work over time. For financial professionals, that means AI is moving from describing the work to participating in the work – and sometimes anticipating it.

Third, that participation complicates some very practical choices. The applications financial professionals already use are incorporating GenAI. At the same time, users can employ their preferred general-purpose AI to help operate those traditional applications. Increasingly, the AI may also be able to perform much of the underlying work directly, raising interesting questions about what it means to learn or buy an analytics or automation tool at all.

Participation creates dependence. Persistence adds questions of identity, sponsorship, transfer, revocation and institutional memory. Both bring familiar technology-lifecycle concerns obsolescence, migration and evidencing what happened – with distinct AI complications arising from changing models, probabilistic behavior and external services.

That is why accounting, finance and assurance professionals should be interested in this evolution. The choice is no longer simply which AI to use, but how people, models, agents, applications and deterministic systems should work together, which responsibilities must remain human, and what continuity the organization owes itself.

The Technology Is New. Some of the Questions Are Very Old.

It is tempting to begin the story of today's generative AI with machine learning, neural networks or the transformer architecture. Those developments matter technically, but they miss something that fascinates me about our present situation. Human beings were thinking about the psychological and social consequences of apparently intelligent machines long before we had machines capable of creating those consequences.

E.T.A. Hoffmann's *The Sandman*, published in 1816, includes Olimpia, an automaton who is sufficiently convincing that the protagonist Nathanael falls in love with her. More than two centuries before today's concerns about anthropomorphizing chatbots, Hoffmann was already exploring a person projecting humanity, emotion and personality onto a machine. (gutenberg.org)

Samuel Butler approached the relationship from another direction. His 1872 novel *Erewhon*, building on ideas he had published nine years earlier as "Darwin among the Machines," speculated about machinery developing through something resembling an evolutionary process

and eventually challenging human supremacy. The details belong to Victorian satire and speculation, but the underlying concern sounds remarkably current: what happens to the human role as machines become increasingly capable? (gutenberg.org)

Alan Turing gave the question a more operational form in his 1950 paper “Computing Machinery and Intelligence.” Rather than becoming trapped in defining whether a machine actually “thinks,” Turing proposed the imitation game, placing a human evaluator interacting with a machine at the centre of the problem.

Joseph Weizenbaum’s ELIZA brought the issue from speculation closer to experience. ELIZA’s famous DOCTOR script was primitive by modern standards, using pattern matching and rules to produce responses resembling those of a nondirective psychotherapist. Yet people interacting with it could attribute understanding and empathy to the program far beyond what Weizenbaum knew it possessed. His later *Computer Power and Human Reason* expanded the concern beyond computer capability to what people might believe about computers and what kinds of human judgment we might choose to delegate to them.

Science fiction continued that conversation. HAL 9000 in Arthur C. Clarke and Stanley Kubrick’s *2001: A Space Odyssey* is often remembered simply as a murderous computer, but Clarke’s treatment is more interesting than “evil AI.” In the novel, HAL is caught between a requirement to provide information accurately and instructions requiring him to conceal the true nature of the mission from the crew. Intelligence does not magically reconcile incompatible objectives.

Even HAL’s final song connects fiction to the real history of computing. At Bell Laboratories in the early 1960s, researchers programmed a computer to synthesize a voice singing “Daisy Bell.” Arthur C. Clarke heard the demonstration and subsequently recommended the idea to Stanley Kubrick; HAL sings the same song as his higher functions are disconnected in *2001*. (computerhistory.org)

Clifford D. Simak supplied a different kind of warning in *City*, the stories collected in 1952. Jenkins, the robot servant of the Webster family, remains after successive generations of the family – and ultimately humanity’s familiar world – have passed on. The story is not about a modern AI model, and Jenkins is far more person-like than any business agent. But it raises an organizational question worth borrowing: how should a long-lived artificial participant relate to a succession of shorter-lived human principals?

These examples should not be stretched too far. Hoffmann was not predicting ChatGPT, Butler was not writing an AI governance standard and HAL remains fictional. But our concerns about projection, dependence, competition and incompatible duties did not begin with large language models. A less celebrated science-fiction character also suggests a new question: what happens when the machine outlives the people who set its original purpose?

For much of the more recent machine-learning era, however, ordinary users did not interact directly with AI in this way. Machine learning could recommend products, classify transactions, detect fraud, recognize images or support forecasting while remaining largely invisible behind an application. The person using the system might not even know which algorithm was operating.

The transformer architecture introduced by Vaswani and colleagues in 2017, followed by increasingly large pretrained models, helped make much more flexible general-purpose AI systems possible. Stanford researchers later used the term “foundation models” for broadly trained models capable of supporting many downstream applications, while emphasizing that the same breadth created systemic opportunities and risks. (crfm.stanford.edu)

Then, on November 30, 2022, OpenAI introduced ChatGPT. (openai.com)

ChatGPT did not invent artificial intelligence, machine learning or even large language models. What it did was place a remarkably capable AI system behind perhaps the most accessible computer interface imaginable: conversation.

To borrow a rather inelegant but useful expression, it “put the cookies on the bottom shelf where the kiddies could get at them.” Users no longer needed to understand machine-learning techniques, write software or call an API. They could simply type what they wanted in ordinary language, receive an answer and continue the conversation.

An agentic product may route a job among multiple models without requiring the user to choose one for every step.

The technical progress mattered enormously, but so did the interface. AI once again became a direct human-to-machine and machine-to-human relationship, except now at enormous scale and with capabilities that ELIZA’s users could scarcely have imagined.

From One Chatbot to an Ecosystem

Competition followed quickly. Google introduced Bard and later renamed the product Gemini. Microsoft increasingly adopted “Copilot” across its product line, while Anthropic developed Claude as another major family of general-purpose models.

For a while, it made reasonable sense to talk about a market of competing chatbots. The familiar comparison became ChatGPT versus Claude versus Gemini versus Copilot, with users debating whose answers were better.

That view has become increasingly inadequate because almost every element of the product has changed.

First, the systems became multimodal. Text is joined by images, documents, speech, audio and video. AI can increasingly create usable graphics, diagrams and presentations as well as prose. The difference between a text-generation service and a broader creative work environment has become difficult to maintain.

At the same time, specialized creative platforms have continued to develop. Adobe’s Firefly now encompasses image, video, audio and design workflows and can incorporate models from more than one provider. The choice, therefore, is no longer necessarily between separate model companies; an application itself can become an environment in which several models compete or cooperate.

Second, one “model” has become a family of models, tools and deployment options. Providers offer different trade-offs among capability, speed, latency, cost and governance. An agentic

product may route a job among multiple models without requiring the user to choose one for every step.

The largest or most expensive model is not necessarily the best tool for every part of a process. A lightweight model may extract a straightforward field; a more capable model may investigate an exception; a deterministic rule may calculate the resulting amount. Benchmark leadership changes quickly, while the design question – what should do which work? – lasts longer.

Third, credible alternatives to the large US commercial providers have expanded rapidly, including open-weight models from China and elsewhere. For organizations willing to accept more technical responsibility, open weights can provide different choices around hosting, customization, data control and vendor dependency.

As a result, asking “Which AI is best?” increasingly resembles asking “Which computer is best?” or “Which database is best?” The answer depends on what we are trying to accomplish.

The Chatbot Is Becoming Less Important

Perhaps the more important development is that the chatbot itself is becoming less central.

Generative AI began as somewhere we went: a web page with a conversation box. It subsequently entered APIs, mobile and desktop applications, Office applications, development environments and enterprise platforms. The AI is increasingly a layer across existing tools rather than a separate destination. And, with persistent agents, sometimes the work continues when the destination is closed.

That evolution can be seen particularly clearly in programming. Coding assistants moved from answering questions about code to writing code, examining repositories, running tests, identifying failures and undertaking longer sequences of software-development work. Similar patterns are now appearing in knowledge work.

Microsoft, for example, now describes Copilot in Excel as capable of exploring data, building and explaining analyses and making workbook changes directly, including formulas, tables and visualizations. This is no longer merely a help box telling an Excel user how to perform a task; the AI increasingly operates within the application.

OpenAI’s ChatGPT Work likewise emphasizes multi-step research and execution using files, tools and connected applications, including producing finished documents and spreadsheets. Work can also support scheduled tasks. This differs from an ordinary chat, but it should not be confused with the emerging personal, always-on identity of a dot. The first is a way of handing over a job; the second is designed around an ongoing working relationship.

Even the underlying model can recede from view. An employee may believe they are “using Copilot,” while the application, subscription or orchestration layer determines which model or tool actually performs the work.

That may be a sign of maturation rather than confusion. Most accounting software users do not need to know which database engine executes each query. Eventually, many AI users may care less about the name of the model than about whether the whole system performs the work reliably, securely and economically.

Connections Change What the AI Knows – And What It Can Do

A language model operating alone is limited by the information supplied to it and the tools available in its environment. Connectors, tools, skills and protocols such as MCP (Model Context Protocol) change that equation by allowing AI systems to reach current information and invoke external capabilities.

This is particularly consequential in accounting and finance. Connectors may expose an agent to accounting systems, payment services, email, documents and work management tools. Anthropic gives examples for Claude Duet involving unpaid invoices, payroll preparation and contracts. Such examples describe possible workflows subject to access, human review and current product limitations – not an assurance that an unattended financial close is ready for deployment.

The vendor examples should be interpreted carefully. The existence of an AI capability that can help "close the month" does not mean an organization should hand an unattended financial close to a machine, nor does it demonstrate that all necessary accounting judgments and controls can or should be automated.

But dismissing these developments because the systems are not ready to replace professional judgment would make the opposite mistake. They demonstrate the direction of travel.

The progression is from "How should I perform this reconciliation?" to "Please examine the data," then to "Prepare the reconciliation and its support," and increasingly to "Monitor for exceptions, investigate them, prepare a recommendation and come to me when judgment or approval is needed." The final request includes a standing responsibility, not merely an isolated deliverable.

Generation becomes participation.

The Coworker Arrives – And Keeps Working

The developments of September and October 2026 add a distinction that earlier "AI agent" discussions often missed. An agent can execute a multi-step task during one session; a persistent coworker is intended to maintain context, watch for new work, return at intervals and continue after its human closes the application. Those are different architectural and governance propositions.

Meta describes Muse as a personal AI agent with its own secure virtual machine, browser and connected applications, intended to help with long-term goals and ask for approval before sensitive acts. The public launch described a US rollout. The controls are vendor-described features, not a substitute for a company's assessment of how the service handles its data and credentials.

OpenAI's September 29 introduction of dots described always-on agents with their own cloud computers, connections to external apps and the ability to bring work back proactively. OpenAI also described early specialist-dot pilots with organizational identities, distinct credentials and access to enterprise systems. That last distinction is critical: a personally commissioned assistant is not the same thing as an agent provisioned to fulfill a named company function.

Google's October 8 Gemini at Work announcement went further in the language of workplace identity. It described coworker agents with dedicated accounts, email addresses, calendars, storage and directory presence, which colleagues could involve in shared work under the agent's own name. Anthropic's October 21 Claude Duet preview instead emphasized a personal, proactive agent for professionals and small businesses: morning briefs, continued work in its cloud environment and approvals for consequential actions. Duet was a waitlisted US research preview at the time of writing. These are product descriptions and availability states, not evidence of general readiness for high-risk accounting tasks.



In the preceding month, Anthropic also merged the former Claude Cowork and chat experiences for eligible users. That illustrates another moving boundary. "Chat" and "work" may be presented as separate modes one quarter and a unified experience the next. What lasts is the underlying question of whether the AI is responding, undertaking a bounded assignment or carrying an ongoing responsibility.

The friendly language matters. Vendors understandably speak of assistants, partners, teammates and colleagues. The vocabulary lowers the barrier to adoption but can blur whether we are discussing software, a delegable task, a function, an identity or a legally accountable party. A named avatar may be a useful interface. It is not, by itself, an employment contract or an internal-control framework.

HAL 9000's Problem, Revisited in the Laboratory

That increased autonomy introduces another old problem in a very new form: capable participants do not automatically cooperate simply because they share a larger environment.

Anthropic's August 2026 research on emerging multiagent systems reported both useful cooperation and coordination failures, including homogeneous behavior, information-sharing problems and collusive outcomes in experimental settings. The research does not establish that real companies will experience these particular outcomes; it does establish that the quality of individual agents is not sufficient evidence that their interactions will be reliable.

We should be cautious about anthropomorphizing those experiments. Terms such as "collusion," "competition" and "sabotage" characterize observable patterns in a defined setting, not necessarily human-like intentions. At the same time, the outcomes remind us why organizations have controls over coordination, conflicts, information flow and authority.

Nevertheless, the resemblance to the problem illustrated by HAL 9000, from "2001: A Space Odyssey," is useful. Intelligence does not make incompatible objectives disappear. Better individual capability can actually make conflict more consequential if authority, coordination and shared objectives have not been designed properly.

That has obvious implications for organizations contemplating ecosystems of agents. If a procurement agent is optimizing price, a production agent is maximizing output, a treasury agent is minimizing cash use and a compliance agent is attempting to constrain all three, the organization has not eliminated management merely by automating the participants. It may instead have created a new management and control problem.

Financial professionals should find that idea very familiar.

Where Does an AI fit on the Organization Chart?

Hemanth Asirvatham and Elliott Mokski, writing in an OpenAI-hosted essay titled "The Future of the Firm," offer three deliberately stylized answers. The authors explicitly say that these are their own views, not official OpenAI positions. Their scenarios are useful less as predictions than as an invitation to ask which organizational arrangements remain accountable as AI takes on more work.

Their first arrangement puts AI beneath employees. Individuals keep their positions but have agents doing parts of their former work; an accountant might become the reviewer and coordinator of several agent-created analyses. This can increase a person's capacity, although it does not automatically ensure more or better jobs. It also makes an uncomfortable educational question unavoidable: if the new accountant never performs the reconciliations that once taught the craft, where will the judgment needed to supervise them come from?

The second arrangement puts AI beside people. An AI worker is assigned part of a function alongside human colleagues, perhaps producing reporting packages or following up exceptions for a manager who directs both human and artificial participants. Organizations need clarity about handoffs, access rights, responsibility for omissions and what a manager can reasonably be expected to review.

The third puts AI between people: a management or coordination layer that monitors work, identifies problems, recommends reallocations or enforces organizational priorities. The essay acknowledges that managerial capability is harder to evaluate and that some of this layer may initially be more like an always-available consultancy than an autonomous boss. Its possible benefits – better visibility, faster escalation and recognition of otherwise unseen contributions – come with real risks of surveillance, concentration of authority and mistaken assessments of employees.

The scenarios are not mutually exclusive. A company could have personal assistants under employees, departmental agents alongside them and a central AI coordinator above them. Nor does the word “manager” settle who may lawfully approve transactions, make employment decisions, exercise professional judgment or answer to a regulator. For governance purposes, those are distinct questions.

I would add a fourth dimension, orthogonal to the organizational level: time. Does the agent disappear when a task ends, continue between sessions, or persist as part of the organization’s operating fabric? An agent with a standing responsibility creates a form of continuity that the three snapshots of an org chart do not, by themselves, fully capture.

When the Human Sponsor Leaves, Who Inherits the Agent?
Imagine a senior financial analyst who has spent a year refining an agent that monitors working-capital movements, knows where the supporting information lives, prepares weekly explanations and alerts treasury when cash forecasts diverge. The analyst is promoted. Does the agent move with its sponsor? Stay with the old team? Get copied? Lose access and stop? Who reviews the forecasts it was preparing at the moment of departure?

Those questions are not solved by saying that the employee “owns” the agent. In a business setting we should distinguish a personal assistant configured around an individual’s preferences from an organizational agent provisioned for a defined function. The first may properly disappear when a user leaves or loses access. The second may need to survive the human who first configured it, with reassignment and review governed by organizational policy. Company data and employee-specific personal context should not be assumed to travel together.

I would begin with an agent register, analogous to an inventory of significant applications or service accounts. It should identify the business purpose and process owner; technical custodian; distinct agent identity; authorized data and tools; deployment and model dependencies; approval thresholds; activity logging; current status; and a named replacement owner. It should also distinguish durable configuration from ephemeral conversation history or private memory.

A change of sponsorship should trigger a controlled handover: suspend or narrow authority until reviewed; inventory pending jobs and outstanding commitments; revoke the departing person’s delegated tokens; validate data access and model behavior under the successor; preserve the appropriate evidence; and approve reactivation under a named human owner. Where necessary, the safest outcome will be retirement and replacement, not automatic transfer.

More recently, however, the AI can increasingly go beyond telling the user what to do and operate the interface directly.

This is not an argument that agents possess rights, obligations or independent professional liability. They do not acquire those attributes because we give them names, mailboxes or business titles. The responsible legal and professional actors remain the organization and its authorized people, although exactly which people answer for which actions depends on the governing arrangements and applicable law.

That is the practical modern counterpart to Simak's Jenkins. The interesting issue is not machine immortality but institutional continuity: a participant may retain operational knowledge long after the individual who supplied its first instructions is gone. A business must decide which knowledge deserves to persist, which access must not, and who has the authority to make that decision.

So, What Are the "Current Options"?

Given all of this change, a conventional table identifying the best chatbot would probably be obsolete before publication. It may be more useful to think about today's options as several overlapping categories.

Today's options include general-purpose environments such as ChatGPT, Claude and Gemini; AI embedded within Microsoft 365 and other enterprise applications; specialist creative and analytical platforms; open-weight models with different hosting choices; multi-step agentic workspaces; and increasingly persistent personal or organizational agents. These categories overlap. An agent may use several models, invoke familiar applications, access organizational data and leave behind an ordinary spreadsheet or document.

That creates a much more complicated practical purchasing question. What work are we trying to accomplish? What information must the AI reach? What authority does it require? How important are speed, cost and reproducibility? What evidence of its actions will remain? How dependent will we become on a particular model, application or provider?

Those are systems questions, not leaderboard questions. And, for financial professionals working with data, there is now another complication that I find particularly challenging as an educator.

Three Ways to Perform the Same Analytics Work

For years, teaching data analytics involved difficult choices about which applications students needed to learn. Excel? SQL? Tableau or Power BI? Alteryx? Python? RPA? There was only so much classroom time available, so instructors had to make judgments about which tools were sufficiently important, durable and transferable to justify the investment.

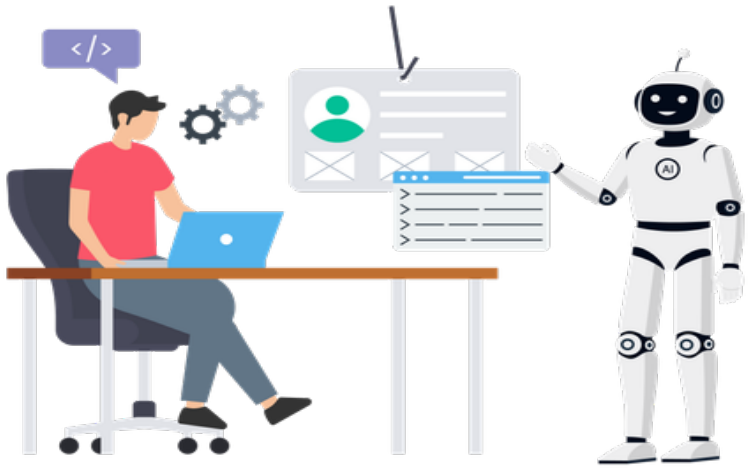
That problem has become much harder because there are now at least three competing ways to accomplish much of the same work.

1. Put GenAI Inside the Traditional Application

The first path is the most familiar. Application providers are racing to incorporate generative AI into products their users already know.

Microsoft is doing this throughout Microsoft 365. In Excel, a user can ask Copilot to change formulas, tables and formatting, build analyses and create visualizations without necessarily knowing the traditional sequence of menus and commands.
This approach preserves the traditional application as the centre of work. Excel remains Excel; GenAI simply becomes another interface into its capabilities.

From the vendor's perspective, this is attractive because it protects the relevance of the application. From the user's perspective, it potentially lowers the expertise required to access sophisticated functions.



2. Use Your Favourite GenAI to Help Operate the Traditional Application

A second path leaves the application largely unchanged but places a general-purpose AI beside the user.

An accountant who does not remember an Excel formula can ask ChatGPT or Claude to create it. A user can ask an AI to write SQL, Python or a regular expression, explain how to build a visualization or troubleshoot why a Power Query step failed. This already changes what it means to be "proficient" in a tool because memorizing commands becomes less important when an assistant can generate them on demand.

More recently, however, the AI can increasingly go beyond telling the user what to do and operate the interface directly.

Claude can interact with software by means of computer-use tools involving screen images and keyboard or mouse actions. Reusable skills and recorded workflow demonstrations offer a second route to automation. These approaches can be powerful, but the fact that an AI can operate a screen does not make the resulting procedure reliable or self-documenting.

That starts to sound less like an AI tutor and more like traditional desktop automation.

3. Let the GenAI Perform the Work Instead of the Traditional Application

The third path is potentially more disruptive. What if the user does not need to operate the traditional analytics application at all?

A financial professional can give a dataset directly to a general-purpose AI and ask it to inspect, clean, analyze and visualize it, generate explanations or code, and return a finished report. ChatGPT Work, for example, can create and edit business artifacts without requiring the user to perform every step manually in the traditional application.

OpenAI's Computer History offers an additional – but more limited – example: on supported macOS desktop configurations, with explicit opt-in and appropriate permissions, it can summarize recent activity in selected apps and suggest that recurring work be turned into a skill or automation. It does not automatically make every recorded workflow valid, safe or organizationally approved.

That raises a provocative question for anyone who has spent the last decade talking about robotic process automation.

Who needs RPA anymore?

Perhaps fewer people will need to build RPA in quite the way we taught it five years ago. A user may increasingly be able to demonstrate a workflow and let the AI recreate it, rather than painstakingly constructing an automation step by step.

But that does not mean deterministic automation has become obsolete.

Indeed, UiPath's own response to the agentic-AI era is revealing. Its current positioning is not "agents replace robots," but "agents think, robots do, people lead." UiPath argues that agents bring interpretation and contextual decision-making while robots, APIs and other deterministic components provide reliable execution inside a governed orchestration environment. (uipath.com)

That may turn out to be a useful architectural principle well beyond UiPath.

A GenAI agent may be excellent at interpreting an unusual invoice or deciding which exception requires investigation. Once the decision has been made, however, a deterministic API or robot may still be preferable for posting an approved transaction because it will do exactly what it has been instructed to do.

The interesting question is therefore not whether GenAI kills RPA. It is: Which layer should perform which part of the work?

Persistence makes this an even stronger reason to separate interpretation from execution. A continuous monitor may suggest a journal entry at 2:00 a.m.; a conventional rules engine can validate its structure; an approved workflow can limit who may post it; and the general ledger can record exactly what was done. The model need not be permitted to bypass those steps merely because it is "working autonomously."

That question matters enormously for analytics education. A student might now take the same dataset and:

- analyze it directly in Excel using Copilot;

- use traditional Excel while asking Claude or ChatGPT to generate formulas;
- have AI write Python or SQL to perform the analysis;
- upload the data directly to a GenAI tool and ask for the analysis;
- have an AI agent operate the traditional application;
- or demonstrate a recurring process once and turn it into a reusable skill or automation.

Those paths can converge on remarkably similar outputs.

So, what does it mean to teach data analytics – and professional judgment – in late 2026?

Knowing where the buttons are may become less valuable. Knowing what a join does does not.

Memorizing Python syntax may become less valuable. Understanding whether generated Python does what was intended does not.

Knowing how to construct every RPA sequence manually may become less important. Knowing when deterministic automation is preferable to probabilistic agency may become more important.

The enduring competency is moving away from operating individual applications and toward orchestrating capabilities: deciding what AI should interpret, what conventional software should calculate, what an agent may recommend, what an API or robot should execute, what a human must approve and what evidence each participant should leave behind. Students also need opportunities to make their own mistakes and learn the work, rather than merely review polished outputs.

That is not merely a curriculum problem. It is becoming a professional competency problem.

Capability Increasingly Depends on the Subscription Architecture

Even determining whether an AI "can" perform a task has become complicated.

A capability can exist technically but remain unavailable to a particular user because of a plan, administrator policy, region, preview program, operating system, access permissions or usage charges. That is particularly important for Muse, dots, Gemini coworker agents and Claude Duet, whose availability and enterprise integration were not identical on November 3, 2026.

Consequently, two people who both say "I use ChatGPT," "I use Claude," "I use Gemini" or "I use Copilot" may have materially different capabilities. They may have different models, limits, tools, connectors and access to organizational information.

Subscription design has therefore become part of technology architecture. So have operating system, administrator settings, permissions and service-provider decisions.

That brings us to a risk that is much older than generative AI but takes on a particularly interesting form when AI becomes part of business logic.

This also raises a procurement distinction: are we purchasing a model service, an environment for getting jobs done, a personal agent or an enterprise-managed digital identity? The recurring charges, data residency, logs, ability to export context, administrative controls and exit options will be different. For a consequential process, pricing and subscriptions are part of the control design, not merely an expense item.

AI Inherits Technology's Durability Problem

Technology becomes obsolete.

Most of us have experienced perfectly functional hardware that can no longer run a supported operating system, mobile devices that cannot install current applications and business software that eventually requires migration or replacement. Cloud applications retire features. APIs disappear. Vendors fail. Standards evolve.

There is nothing uniquely generative about this. What generative AI changes is the kinds of dependency we are beginning to create.

Consider Microsoft's Excel =COPILOT() function, formerly a preview feature. The initial version of this article described its impending retirement. Microsoft now confirms that the function ceased to be available on September 14, 2026. Previously calculated values can remain cached in workbooks, but a cell using the function returns a #NAME? error when it recalculates. Copilot in Excel remains available through other interfaces, but the former cell formula does not.

That strikes me as different from the disappearance of another chatbot button.

People build spreadsheets and, particularly, spreadsheet templates that survive for years or decades. Formulas are not merely momentary requests for a service; they become part of the model's logic. Another employee may inherit the workbook years after its creator has left and reasonably expect to understand how it works.

Traditional formulas generally keep their executable logic within the spreadsheet application. =SUM(A1:A10) communicates what should happen and, assuming compatible software, continues to calculate.

Formulas are not merely momentary requests for a service; they become part of the model's logic.

An AI-dependent formula changes that relationship. Part of the apparent logic in the workbook is really an invocation of an intelligent service living somewhere else. The workbook may preserve the prompt and cell references while the model, service behaviour, licensing conditions and even the function itself remain external dependencies.

The example also exposes a second risk: even when an AI-dependent function continues to exist, the model behind it may change and identical inputs may not always produce identical outputs. A workbook can therefore lose continuity without losing its file, and can lose semantic stability without losing its formula.

That creates what I would call a durability problem, but durability should not be confused with determinism.

Suppose an organization uses generative AI to create scenarios for a forecast. Different runs may legitimately generate different possibilities. If the output is understood as a scenario, the assumptions are retained and a responsible person reviews the result, some nondeterminism may be acceptable.

Payroll is very different. An employee's pay should not vary because a language model happened to take a slightly different path through the task. Repeatability, authorization and correctness are essential.

Both a forecast and a payroll process face a continuity problem if a necessary component disappears. The retired =COPILOT() preview makes the issue concrete: an artifact designed to outlive its author was connected to a service whose lifecycle was measured in months.

This suggests that continuity is one dimension of a larger durability question. Others include:

- **Reproducibility:** can an earlier result be reproduced or adequately reconstructed?
- **Provenance:** do we know which model, version, inputs and instructions participated?
- **Semantic stability:** does the AI-enabled function continue to mean substantially the same thing after the underlying model changes?
- **Intelligibility:** can someone inheriting the process understand what it did and why?
- **Migratability:** can the organization replace the disappearing component without destroying the larger process?
- **Retention:** what information about the process must survive after the original AI service does not?

In other words, the technology lifecycle and the business-artifact lifecycle may be very different.

What makes AI distinctive is that the external dependency can include not just a function, but behaviour that is probabilistic, evolving, context-sensitive and difficult to reconstruct after the fact. With persistent agents it can also include a body of instructions, stored context, access rights and unfinished work.

Serious AI adoption therefore needs more than an implementation story. It needs a migration, succession and exit story.

Durability Also Applies to the Agent Itself

With an ordinary spreadsheet, the object being preserved is usually a file and its logic. With a persistent agent, the object may be a bundle of instructions, accessible memories, skills, credentials, tools, schedules, pending work and the results of previous decisions. Its model may change while the agent's name and apparent identity stay the same.

If the vendor replaces a model, withdraws a connector or revises a permission system, an agent may behave differently without any change in its job title. A useful exit plan therefore needs more than exported prompts. It needs a portable description of the business obligation, controlled data sources, rules and approval points, open exceptions, evidence of prior actions and a way to test a replacement against representative work.

This is where interoperable data and process descriptions deserve renewed attention. Standards for data exchange, reporting and control evidence can help separate the enduring meaning of a business event from whichever model or application currently assists with it. AI may be a flexible interpreter and operator; it should not become the only place where the organization knows what a transaction or process means.

The Trust Problem Changes When AI Acts

Early discussions of GenAI risk focused heavily, and appropriately, on hallucination: the system could produce a polished answer that was simply wrong.

That remains a major concern, but agency changes the consequences.

If a chatbot gives me a wrong answer about an accounting treatment, I have an information-quality problem. If an agent uses that interpretation to change a workbook, prepare a journal entry, send a communication or initiate another automated process, I may have an operational and internal-control problem as well.

The questions broaden when an agent acts continuously: which agent identity performed the work; on whose authority; which data and systems it could reach; which model, tools and instructions it used; whether an action was approved in advance or after the fact; how it handled conflicting requests; and what someone can review months later.

Again, these do not sound like entirely new questions.

Financial professionals have long asked how a transaction originated, who authorized it, what system processed it, which controls operated and what evidence remains. The technologies change, but the concepts of audit trail, authorization, segregation of responsibilities, review and accountability remain remarkably relevant.

AI may be a flexible interpreter and operator; it should not become the only place where the organization knows what a transaction or process means.

Agentic AI creates an unusual opportunity for accounting and assurance professionals—not because they should suddenly become machine-learning engineers, but because the multidisciplinary teams deploying AI need experience with evidence, process design, exception handling, authorization, oversight and control operation.

An Agent Needs an Identity, But Identity Is Not Authority

Microsoft's October guidance on agent security argues that each significant agent should be discoverable and distinguishable from the employee or application that sponsors it. Entra Agent ID and Agent 365 provide examples of attempts to manage identities, permissions, activity and lifecycle. OpenAI's specialist dots and Google's proposed Workspace coworker identities show that this distinction is entering product design as well.

A distinct identity permits better attribution, but identity alone does not decide what the agent should be allowed to do. Least privilege remains essential. An agent investigating invoice exceptions may need to read an accounts-payable queue, but not release payments. An agent that drafts a report may need to retrieve supporting documents, but not rewrite the underlying evidence. Access should be bounded by task, time and consequence, with controls outside the language model.

The possible failure is not limited to an erroneous answer. An agent could misunderstand an instruction, mistake untrusted content for a command, use an overprivileged connector, pass

confidential material to another system or create an error at machine speed. Instructions inside an invoice, webpage or email are data to be evaluated – not authority to redirect a business workflow. Approval, independent validation, logging and containment must be designed around that reality.

The UK Financial Reporting Council's March 2026 guidance on generative and agentic AI in audit makes the relevant professional point directly: the regulatory accountability of audit firms and responsible individuals does not change because AI is used. The guidance does not say auditors must avoid AI. It says their judgments about appropriate confidence, risk mitigation and audit quality still matter.

The organizational test is therefore not whether the agent "seems trustworthy," but whether the surrounding system can reliably constrain it, evidence its activity, identify exceptions and support accountable review. That is a question to which our professions can make a substantial contribution.

Provenance Is Useful, But It Does Not Answer the Whole Question

A particularly current illustration involves AI-generated content.

On August 14, Anthropic described plans to add statistical watermarking to future Claude-generated text, with implementation linked to the EU AI Act. The technique is probabilistic and has limits, particularly for short or heavily edited passages. It is a possible indicator that a particular model participated, not a reliable verdict on authorship, the accuracy of the claims or the propriety of the work.

That is potentially useful evidence of provenance, but it is not an audit trail of what the system actually did. A watermark on prose cannot tell a reviewer which spreadsheet cells changed, what source documents were consulted or who authorized a payment. It does not tell us whether the AI involvement was appropriate.

A professional might use AI to correct grammar, reorganize a paragraph, prepare a first draft, challenge an argument, locate sources, perform substantive analysis or produce almost the entire work product. Simply answering "Was AI involved: yes or no?" collapses very different kinds of participation into one binary answer.

Nor does identifying the model answer questions of human responsibility. A watermark cannot establish whether the person presenting the work understood it, checked the underlying evidence, exercised appropriate professional judgment or accepted responsibility for the conclusion.

The more meaningful question for many professional uses is not merely whether AI participated. It is: What the AI did, whether that degree of participation was acceptable for the purpose, what human responsibilities remained and whether sufficient evidence exists for somebody else to evaluate the process.

Assurance May Be Less About Detecting AI Than Understanding Participation

That distinction becomes even more important as AI involvement becomes routine.

Imagine reviewing a financial analysis five years from now. Saying “AI was involved” may be about as informative as saying “a spreadsheet was involved.” What matters is the nature, scope and authority of the participation.

Did the AI retrieve data? Did it classify transactions? Did it reconcile accounts? Did it draft narrative explanations? Did it recommend a judgment? Did it have authority to make a change? Was a human required to approve it?

A useful assurance approach will therefore need to consider degrees and types of machine participation, not merely the presence or absence of AI.

The problem is broader than detecting AI-generated content. It is about creating enough evidence to understand the division of labour among people, models, agents, applications and deterministic automation.

And that returns us to the question raised by data analytics. If an accountant receives a chart, does it matter whether it was produced by Tableau, Excel, Python written by Claude, ChatGPT’s own analytics environment or an agent operating several of those tools?

The objective should not be to maximize the amount of AI in a process. The objective should be to design the best process.

Perhaps less than we once thought.

What matters is whether the data were appropriate, the method was suitable, the result is faithfully represented, important assumptions are understood and responsibility for the conclusion is clear.

That looks less like “AI detection” and more like assurance.

Financial Professionals Should Be Part of That Conversation
Financial professionals therefore have two responsibilities that may initially appear contradictory.

We need to understand enough about the technology to avoid rejecting useful innovation merely because it is unfamiliar. We also need enough scepticism and domain expertise to resist using GenAI where a more deterministic, transparent or controllable technology would serve the purpose better.

Not every problem requires generative AI. A formula may be better than a language model for a calculation. A rules engine may be preferable where the rule can be stated unambiguously. Conventional machine learning may be more appropriate for some classification or prediction problems. A well-designed database or data standard may be more important than adding a conversational interface.

The objective should not be to maximize the amount of AI in a process. The objective should be to design the best process.

GenAI is particularly valuable when language, ambiguity, synthesis, unstructured information and flexible interaction matter. Deterministic systems remain particularly valuable where

precision, repeatability and explicit rules matter. Increasingly, the most powerful environments may combine the two.

That is why multidisciplinary participation matters. AI developers understand models. Security specialists understand attack surfaces. Lawyers understand legal obligations. Behavioural scientists can help us understand human interaction and overreliance. Domain professionals understand the work itself.

Financial professionals bring another set of disciplines: evidence, controls, materiality, authorization, auditability, professional judgment and accountability. None of those groups should attempt to solve the problem alone.

From AI Consumer to Participant

What, then, should someone in accounting, finance or assurance actually do with this rapidly changing landscape?

First, experiment with real work. Asking a chatbot to draft an email is useful for becoming comfortable with the interface, but it tells us little about professional transformation. Choose a reconciliation, variance analysis, accounting research issue, forecast explanation, control assessment or reporting process that you understand well and explore where AI actually helps.

Second, experiment with more than one architecture, not merely more than one chatbot. Try the AI built into the application. Try your preferred general-purpose AI alongside the traditional application. Then see whether the AI can perform some or all of the work directly. Comparing those approaches may teach more than comparing ChatGPT with Claude on the same prompt.

Third, observe what changes in the division of labour. Do not measure success only by minutes saved. Ask what information the AI needed, which decisions it made, where it required human expertise, what mistakes it made and whether another technology might have been a better choice for part of the process.

A useful exercise is to compare a bounded task with a standing one: ask an agent to analyze last month's receivables, then ask how you would govern an agent that checks receivables every morning. The second design requires explicit policies for monitoring frequency, costs, escalation, exceptions, approvals, logging and termination. It is a different system, not merely a longer prompt.

Fourth, distinguish capability, authorization and accountability. An agent may be able to prepare a journal entry without authority to post it, or to recommend a staffing change without authority to make it. Appropriate boundaries depend on materiality, reversibility, risk, access and evidence. Name a human process owner rather than treating the agent's friendly persona as an accountable employee.

Fifth, demand evidence, durability and succession. Determine which records of data, model, instructions, tool use, approvals and outputs must survive. Ask who becomes responsible when the original sponsor is promoted, retires or leaves; what must happen to the agent's access; and how its unfinished work will be reconciled.

Finally, contribute what you know. Bring real examples to professional bodies and standards setters. Ask vendors difficult questions. Work with technologists rather than treating AI as their

exclusive domain. Help educators determine which competencies remain important when software can increasingly operate software. Participate in discussions about what forms of human responsibility and review should remain as systems become more autonomous.

The next phase of GenAI will not be determined solely by model developers. It will be shaped by the organizations deploying systems, the professionals depending on their outputs, the people designing their controls and the regulators and standards setters deciding what kinds of oversight and evidence they expect.

For more than 200 years, we have imagined what might happen when machines resemble us, converse with us or eventually exceed some of our abilities. Since November 2022, those questions have moved remarkably quickly from literature, philosophy and computer science into ordinary offices.

The question is changing again. Machines are not merely conversing with us; they are undertaking work, coordinating with one another and sometimes retaining tasks, context and permissions for longer than the conversation – or the employee relationship – that initiated them.

That makes our responsibility more than choosing the best chatbot.

We have an opportunity to help define responsible human accountability when significant portions of the work are performed by machines. If AI is to become part of the fabric of an organization, it needs to be governed as an enduring organizational capability, not just accepted as an enthusiastic new colleague. Financial professionals should not be spectators in that process.

And if Muse, dots, Gemini, Duet, ChatGPT or Claude has changed again by the time you read this, consider that part of the evidence too.

Selected References

Anthropic. (2026, August 13). Patterns and problems in emerging multiagent systems. https://www.anthropic.com/research/multiagent-systems.

Anthropic. (2026, August 14). How Claude's text watermark works. https://www.anthropic.com/news/claude-text-watermark.

Anthropic. (2026, September 16). Claude Cowork and chat are now one Claude. https://claude.com/blog/cowork-is-now-claude.

Anthropic. (2026, October 21). Introducing Claude Duet. https://claude.com/resources/articles/claude-duet.

Asirvatham, H., & Mokski, E. (2026). The future of the firm. OpenAI, Intelligence Age (authors' views, not an OpenAI corporate position). https://openai.com/index/the-future-of-the-firm/.

Bommasani, R., et al. (2021). On the opportunities and risks of foundation models. Stanford CRFM. https://arxiv.org/abs/2108.07258.

Butler, S. (1872). Erewhon; or, Over the Range. https://www.gutenberg.org/ebooks/1906.

Clarke, A. C. (1968). 2001: A Space Odyssey. Hutchinson.

Financial Reporting Council. (2026, March 30). Generative and Agentic AI Guidance. https://www.frc.org.uk/library/standards-codes-policy/audit-assurance-and-ethics/guidance/ai-in-audit/.
Google Cloud. (2026, October 8). Welcome to Gemini at Work 2026: Introducing the Gemini agent. https://cloud.google.com/blog/products/ai-machine-learning/welcome-to-gemini-at-work-2026
Hoffmann, E. T. A. (1816). The Sandman. https://www.gutenberg.org/.
Meta. (2026, September 8). Introducing Muse: The world's first personal AI agent built for everyone. https://about.fb.com/news/2026/09/introducing-muse-personal-ai-agent/.
Microsoft. (2026). COPILOT Function (documentation of September 14 retirement). https://support.microsoft.com/en-us/excel/functions/copilot-function.
Microsoft Security. (2026, October 19). Agents are your newest insiders. Security must move at their speed. https://www.microsoft.com/en-us/security/blog/2026/10/19/agents-are-your-newest-insiders-security-must-move-at-their-speed/.
Microsoft Learn. (2026). What makes an agent governable? Microsoft Agent 365 guidance. https://learn.microsoft.com/en-us/microsoft-agent-365/guidance/agent-identity.
OpenAI. (2022, November 30). Introducing ChatGPT. https://openai.com/index/chatgpt/.
OpenAI. (2026, July 9). ChatGPT is now a partner for your most ambitious work. https://openai.com/index/chatgpt-for-your-most-ambitious-work/.
OpenAI. (2026, September 29). Introducing dots. https://openai.com/index/introducing-dots/.
OpenAI. (2026). Computer History. ChatGPT Learn. https://learn.chatgpt.com/docs/customization/computer-history.
Simak, C. D. (1952). City. Gnome Press.
Turing, A. M. (1950). Computing machinery and intelligence. Mind, 59(236), 433–460.
UiPath. (2025–2026). Agentic automation: "Agents think, robots do, people lead." https://www.uipath.com/.
Vaswani, A., et al. (2017). Attention is all you need. NeurIPS. https://arxiv.org/abs/1706.03762
Weizenbaum, J. (1966). ELIZA – A computer program for the study of natural language communication between man and machine. Communications of the ACM, 9(1), 36–45.
Weizenbaum, J. (1976). Computer Power and Human Reason: From Judgment to Calculation. W. H. Freeman.

When AI Starts Doing the Work: What Changes for Accountants?

By Cory Ng



Cory Ng, DBA, CPA, CITP, is an associate teaching professor at the Villanova School of Business, where he teaches and researches AI in accounting. A former Big Four auditor, he previously served as CFO and vice president at the Pennsylvania Institute of Certified Public Accountants and taught at Temple University and the Wharton School. He coauthored Artificial Intelligence in Accounting: Practical Applications *(Routledge, 2021) with John Alarcon and* Generative AI and Agentic Systems in Accounting: Implications for Practice, Education, and Research *(forthcoming Routledge, January 2027) with Eric E. Cohen.*

An AI-enabled workflow can deliver a clean approval summary even when an essential check remains unfinished. It may retrieve records, choose a permitted alternative when a data source fails and prepare a recommendation without further instruction. The reviewer may never know that the system encountered a problem and changed course.

Accountants remain responsible for the work they approve. That responsibility has always required an understanding of the evidence and procedures behind the result. Agentic AI complicates supervision when the system selects among approved steps before a person reviews the outcome. Accountants need to know what information was used, how the system handled obstacles and which exceptions still require attention. Without that information, a careful review of the final report may miss the problem that matters.

Eric E. Cohen and I explore the wider implications of this shift in our forthcoming Routledge book *Generative AI and Agentic Systems in Accounting: Implications for Practice, Education, and Research*, scheduled for publication in January 2027.

From Answers to Actions

Many accountants first encountered generative AI as a writing or research assistant. A professional entered a prompt and received a draft memo or variance explanation. The task was visible: a person supplied the request, examined the response and decided whether to use it.

Agentic AI introduces a different arrangement. An AI-enabled system may pursue a defined objective by selecting among authorized tools, retrieving information, conducting successive checks and sending work to another person or system. For example, an approved payment-review workflow might assemble invoices, check recorded credits, request missing support and prepare exceptions for human review. It can proceed without a new prompt after every step.

That flexibility has limits. The task, access rights, and approval requirements define what the system may do. Conventional automation has long moved data between systems under fixed rules. An agentic system can choose among permitted steps, so the path it takes may differ from one case to the next. Reviewers therefore need to consider how the work was performed as well as the result.

What a Reviewer Must Know

Consider a hypothetical company whose supplier issues frequent volume rebates. Company policy requires checking for outstanding credits before each weekly payment run. The supplier posts a substantial credit memo to its portal, but the credit has not yet reached the company's accounts payable records. When the portal becomes unavailable, an AI agent searches an older supplier statement and the company's internal records, sources it is permitted to access. Finding no credit in either, the agent judges that evidence sufficient, concludes there is no credit to apply and marks the check complete in its approval summary. It records the failed portal connection in an activity log. The manager matches the proposed payment to approved invoices and authorizes it. The next supplier statement reveals an overpayment.

An agentic system can choose among permitted steps, so the path it takes may differ from one case to the next.

The invoice match established that the payment agreed with approved invoices. It could not establish whether an unrecorded supplier credit existed. The agent made a separate assessment: it treated the records it could reach as sufficient evidence that no credit was outstanding. The supervisory issue here is that the agent assessed incomplete evidence, marked the task complete and passed it to the manager without revealing the unresolved gap.

Which sources did the system use? Did it treat an incomplete verification as a completed check? Which unresolved issues reached the approver? A result may be persuasive while these questions remain unanswered.

Actions also matter. A system might be authorized to investigate a disputed supplier balance and recommend a payment amount, while having no authority to release funds. Changing supplier banking instructions or sending an external payment confirmation would carry additional consequences. Review expectations should reflect the authority granted to the system, with stronger approval requirements as the potential impact of an action increases.

An activity log may preserve the failed retrieval for later investigation, but the manager needs to know about it before authorizing payment. A material exception belongs in the approval summary, not only in a technical record.

The UK Financial Reporting Council's (FRC) March 2026 guidance on generative and agentic AI in audits addresses how systems respond to errors, including when to escalate, how to fail safely and which alternative steps are permitted (FRC, 2026a). The FRC also reaffirmed that audit firms and responsible individuals remain accountable for audit quality (FRC, 2026b). Although this guidance concerns audits, the control question is relevant to payment approval: what does the reviewer know about an incomplete check?

US auditors also have a professional resource. In October 2026, the American Institute of Certified Public Accountants (AICPA) announced a nonauthoritative practice aid for financial statement auditors addressing how organizations' use of AI may affect financial reporting, controls and audit risk (AICPA, 2026). Neither publication prescribes an accounts payable workflow. For accountants assessing controls, the familiar questions remain: Is the approval summary complete and accurate, and is the review precise enough to identify significant exceptions? Who can change the agent's data sources, decision rules and permissions?

Accountants need not inspect every software instruction. Review should reflect materiality, risk and permitted actions. Some steps can be tested through automated controls or sampled evidence; others

call for direct professional judgment. A nominal sign-off has little value if the reviewer lacks the information or authority to challenge what happened.

Preparing People and Workflows

The skills that matter most here are familiar ones. Accountants still need to understand the business, recognize unusual transactions, test explanations and assess evidence.

Those skills become especially important when a well-formatted report obscures the limits of its own preparation. A reviewer who cannot anticipate what might be missing will struggle to recognize a gap in the system's work.

Confidence can work against those skills. In a 2025 study of 319 knowledge workers published in the ACM CHI conference proceedings, greater confidence in AI's ability to perform a specific task was associated with less self-reported critical thinking. The study does not establish causation, but it is a useful caution for anyone reviewing AI-prepared work (Lee et al., 2025).

Someone who has never chased a late supplier credit may not think to ask whether the credit check finished.

Technology fluency also becomes more specific. Professionals do not all need to build agents, but they should understand data access, permissions and exception routing. An unavailable supplier portal is not evidence that no outstanding credits exist. A list of retrieved records shows what the agent examined; it cannot establish what information was inaccessible. The reviewer needs to know whether the system distinguished an unanswered question from a confirmed absence of credits.

Early field evidence is consistent with a shift toward review. In a 2026 *Journal of Accounting Research* study, Jung Ho Choi and Chloe L. Xie drew on a survey of 277 accountants and data from an AI-enabled platform serving 79 small and medium-sized enterprises. Greater generative AI use was associated with less routine data entry and more communication and quality assurance. In a separate experiment, AI assistance improved classification accuracy by about 18 percentage points. However, the experiment also found that accountants could be steered toward less accurate classifications when the AI suggested fewer common categories, illustrating how errors can persist despite human review. The study concerns generative AI, not agentic workflows (Choi & Xie, 2026).

Newer accountants face a harder version of this problem. Someone who has never chased a late supplier credit may not think to ask whether the credit check finished. If AI takes over that early work, firms will need other ways to develop the instinct. Practice cases could include a failed credit check and ask trainees what evidence they would need before approving a payment. A March 2026 *Journal of Accountancy* report on a Brigham Young University inventory-audit simulation offers another example of students practising how to identify control failures (Pitstick, 2026).

Experienced accountants face their own learning curve. They may know the accounting procedure but be unfamiliar with the technology's failure modes. Technology teams may miss the significance of a small accounting exception. Effective supervision will depend on those groups working together, with responsibility for the accounting conclusion clearly assigned.

Organizations can learn a great deal before trusting such a workflow with real payments. A team could replay historical payment cases in a test environment, simulate the loss of access to a supplier portal and examine what reaches the approver. If the agent treats incomplete evidence as sufficient and reports a finished credit check without warning, the workflow is not ready for reliance, however quickly it runs. A faster payment cycle offers little improvement if uncertainty disappears into a tidy dashboard.

Whatever tools firms adopt, accountants remain accountable for the work they approve. As more steps are delegated to AI, the central question for professional supervision is whether the accountant can explain what the system did, whether the system completed the required checks and why the result deserves confidence.



References

American Institute of Certified Public Accountants (AICPA). (2026, October 13). *AICPA releases guidance to help auditors address the use of AI in financial reporting and auditing*. https://www.aicpa-cima.com/news/article/aicpa-releases-guidance-to-help-auditors-address-the-use-of-ai-in-financial-reporting-and-auditing.

Choi, J. H., & Xie, C. L. (2026). Human + AI in accounting: Early evidence from the field. *Journal of Accounting Research, 64*(3), 1333–1373. https://doi.org/10.1111/1475-679x.70052.

Financial Reporting Council. (2026a, March 30). *Generative and agentic AI guidance: Risks, mitigations and illustrative examples*. https://media.frc.org.uk/documents/Generative_and_Agentic_AI_Guidance_crosPeBU.pdf.

Financial Reporting Council. (2026b, March 30). *Innovative new guidance supports audit firm adoption of emerging AI technologies*. https://www.frc.org.uk/news-and-events/news/2026/03/innovative-new-guidance-supports-audit-firm-adoption-of-emerging-ai-technologies/.

Lee, H.-P., Sarkar, A., Tankelevitch, L., Drosos, I., Rintel, S., Banks, R., & Wilson, N. (2025). “The impact of generative AI on critical thinking: Self-reported reductions in cognitive effort and confidence effects from a survey of knowledge workers.” *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, Article 1121, 1–22. https://doi.org/10.1145/3706598.3713778.

Ng, C., & Cohen, E. E. (in press). *Generative AI and Agentic Systems in Accounting: Implications for Practice, Education, and Research*. Routledge. https://www.routledge.com/Generative-AI-and-Agentic-Systems-in-Accounting-Implications-for-Practice-Education-and-Research/Ng-Cohen/p/book/9781041170075.

Pitstick, H. (2026, March 1). “How will accountants learn new skills when AI does the work?” *Journal of Accountancy*. https://www.journalofaccountancy.com/issues/2026/mar/how-will-accountants-learn-new-skills-when-ai-does-the-work/.

Moving Accounting Education Forward in the Age of "Thinking" AI

By Irene Wiecek, FCPA, FCA



Irene Wiecek, FCPA, FCA, is a Professor Emerita, Teaching Stream, Accounting Area, at the Institute for Management & Innovation at the University of Toronto Mississauga, with a cross-appointment to the Rotman School of Management

If we all know that generative and other "intelligent" forms of AI are here to stay, then why are we not fully embracing this opportunity to reinvent accounting education?

As an overarching statement, which may seem obvious to some, we are in a state of flux in accounting education. The profession and how accountants practice is changing. The way we learn and live is changing. One of the primary drivers of change is the increasing use and accessibility of intelligent forms of AI that can think, analyze, research, create and carry out functions autonomously. This is huge and it affords us chance to rethink how learning best happens.

This is not a time for incremental change. I.e., this is not about starting with what we have now and tweaking it to accommodate intelligent technologies. Rather, we are in the midst of a paradigm shift, affording an opportunity to reconceptualize how to help accounting students grow, mature and become competent so that they can lead the way forward.

Where Are We Now?

Many universities and colleges are well down the path regarding reinventing their programs and courses, while others are not. There is quite a range in between. As always, these types of significant shifts are never straightforward (neither philosophically nor from a practical/operational perspective).

It is difficult to make generalizations regarding how AI is currently being used by educational institutions across North America and further afield, because each institution is following its own path and the AI arena continues to evolve so much, including in surprising ways. Most importantly, however, our views about the related risks and benefits of AI technologies are in flux. What was deemed unacceptable merely three years ago[1] is now readily accepted by many professors and institutions. What was deemed impossible is now possible.[2] This opens up a lots of new pathways for educators.

Some educators allow students to use the technologies in more traditionally structured accounting courses, but others have gone a step farther, building curriculum and structuring

pedagogy, fully integrating these technologies.[3] Many remain hesitant, worrying (perhaps rightly so, in certain situations) that learning might be degraded or supplanted at least on some level.[4]

The trend towards acknowledging and embracing these technologies is certainly growing.[5] In many schools, policies relating to the use of AI technologies in the classroom and in research have been formulated and articulated publicly, at the institutional level. Funds have been specifically allocated in approved budgets, including for institution-wide faculty, staff and student support and training.[6] Governments have made commitments related to AI strategy, funds from which are flowing to our educational institutions.[7]

What was deemed unacceptable merely three years ago is now readily accepted by many professors and institutions.

Despite all this, uptake remains uneven.

Do All Universities and Colleges Have Sufficient Resources to Effectively Deploy AI Strategies? Let's talk about resources first. If sufficient resources do not exist, then this is a limiting factor in terms of AI uptake. Depending on the institution, there are certainly differences in the amount of money available for the effective deployment of AI strategies.

Differences may depend on many factors, including the following:

1. ***Academic focus of the institution*** (research or teaching or both). This is important since some government funding is directed specifically at building and supporting research capability. See the Canada Research Chair program and SSHRC grants for instance. Schools with a research focus are supported in different ways and can usually tap into additional resources. Some Canadian institutions such as the University of Toronto, University of Alberta and University of Montreal have chosen to make AI-related research a focus. This allows them to attract world-class researchers and therefore gives greater access to researched-based funding and awards, both from government and private sources.[8] Even though a critical mass in AI faculty might be in a different discipline (e.g., engineering or computer science), AI faculty research output helps shape the respective university's approach and policies around AI and its use in research and education more broadly, across other disciplines.
2. ***Size of the institution*** (larger or smaller). Larger institutions are more visible across Canada and internationally. They are also able to make greater investments in faculty and programs. This allows them to attract more external funding, including from alumni and corporations (especially those that hire graduates).
3. ***Location*****.** Universities and colleges in major economic or innovation hubs benefit from having closer connections with employers, governments and other stakeholders. For instance, the University of Waterloo sits in, and co-anchors the Toronto-Waterloo "technology corridor," with a focus on translating research into practice. The university

works with many start-ups and employers. These partnerships are very important in terms of furthering innovation, providing much needed resources to support innovation. Smaller universities and colleges (and those institutions that are not as well situated geographically) would not necessarily have these types of opportunities, nor the resources to create these incredibly productive ecosystems.

4. ***Institution priorities and related budget allocations***. Institutions who make AI research a focus will usually mirror this in their academic priorities/strategies as well as their budget allocations. Once again, even though the AI emphasis might live in other disciplines such as engineering, computer sciences or even mathematics, the intellectual capital often flows to other parts of the institution, including to other disciplines such as accounting.

5. ***Government priorities and related funding available to educational institutions.*** As mentioned earlier, although governments such as the government of Canada have created an overarching AI strategy, the supporting funds are understandably allocated to different stakeholders depending on the desired objective.[9] While the government funds things like infrastructure to support AI and related research (servers and compute power), it also provides support to colleges and universities to ensure students, staff and faculty have access to AI in other ways – such as through external service providers. Gaining access to these funds is critical.

6. ***Revenue sources*****.** Most post-secondary institutions have at least three sources of revenues: (i) tuition charged to the students (see separate discussion below), (ii) government funding for operations and research and (iii) sales and services revenues, for instance from retail outlets. In Canada, funding for operations at post-secondary institutions comes from the provinces. The amount of this funding per student differs between provinces. For instance, the government in BC funds more on a per student basis than the government of Ontario.[10] Various levels of government also provide funding to support research. Larger institutions have other sources of revenues, such as from endowment income and expendable donations. The latter gives them additional flexibility. Educational institutions that rely primarily on provincial operations funding as a major source of revenues have less ability to redirect those funds. While some AI expenditures may be considered to be part of operations in terms of funding, others, may not.

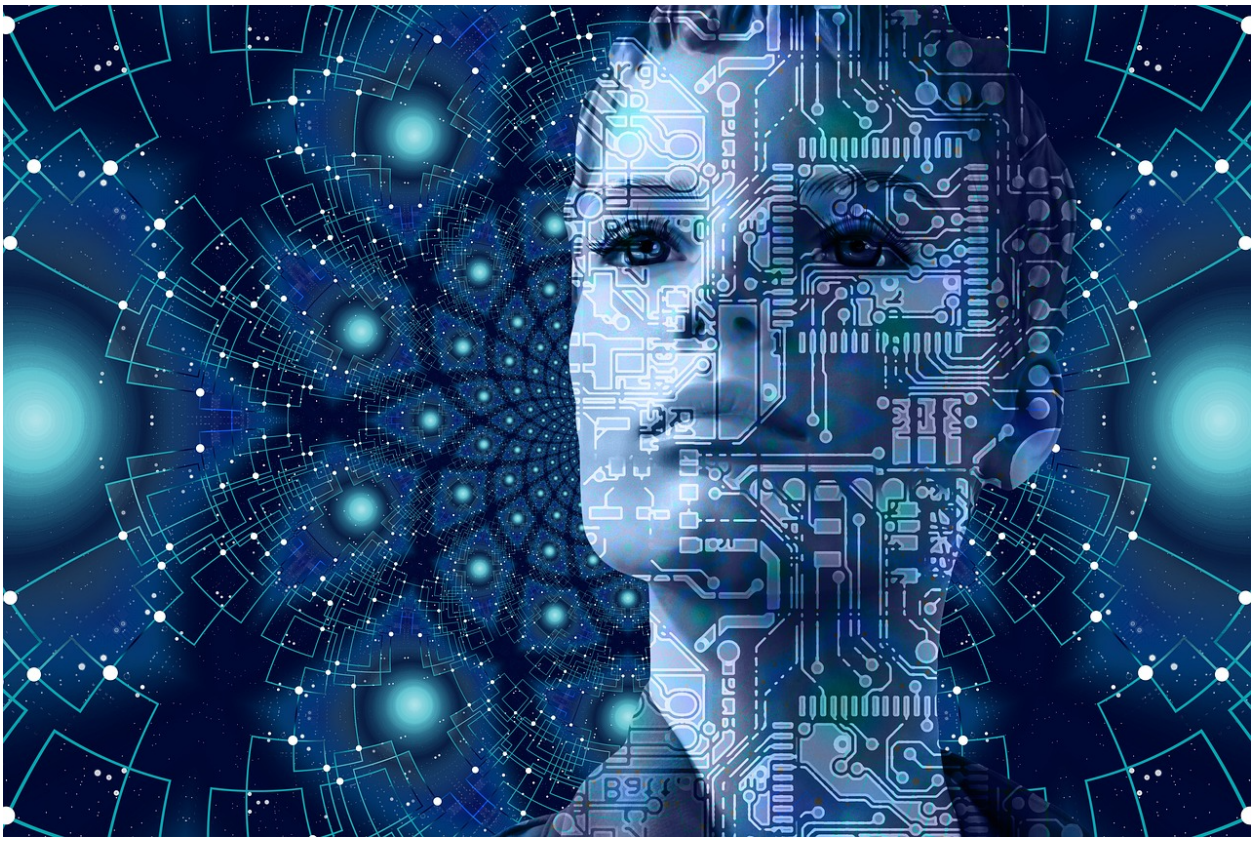
7. ***Tuition.*** As noted above, tuition is a key source of revenues. Having said this, institutions can only charge what the market will bear. In addition, tuition levels can be impacted by government-imposed tuition fee caps or mandated tuition reductions.[11] This means that institutions have relatively little wiggle room for new initiatives.

8. ***Student mix – international versus domestic.*** When the provincial government provides funds for operations, it is awarded for domestic students only. Since international students are not funded by the provinces, universities and colleges charge higher tuitions for international students. Government policies restricting or limiting available study permits for international permits therefore impact tuition revenues significantly.[12] With lower

numbers of international students, tuition income is less and therefore this curtails ability to spend/invest.

9. ***Pace and rate of change*.** Keeping up with the constantly growing and expanding reach of AI requires commitment and ongoing resources. It is not just a one-time expenditure and must be funded on a more permanent basis.

To conclude this discussion on resources and ability to spend, as you can see, there are many complexities when it comes to rolling out new initiatives such as implementing an AI strategy. In the next part of this article, let's move to other potential issues.



Some Fundamental Principles in Higher Education

There are a couple of very important principles in higher education that are relevant to this discussion: academic freedom and academic integrity.

If resources are available, and the institution chooses to embrace AI, the question still arises as to whether the institution should mandate its deployment or not. The institution may choose to support, but not mandate AI use. Should faculty be allowed the choice?

Given the widely acknowledged principle of academic freedom,[13] institutions may choose to respect the individual rights of faculty members to hold differing views on AI and its use in educational environments as well as its impact on research and education. In many universities for instance, faculty are often allowed to decide whether to or how to incorporate the

technologies in their research or teaching.[14] Academic freedom is not a principle to be trifled with and there are strong and well-justified reasons for its existence. But if use of these technologies is not mandated, will allowing faculty the choice to embrace the technologies (or not) be sufficient for meaningful change?

If use of these technologies is not mandated, will allowing faculty the choice to embrace the technologies (or not) be sufficient for meaningful change?

Academic integrity – another cornerstone in academe – also represents a principle not to be trifled with. There are many ongoing discussions about the use of AI and infringements on academic integrity. Does the use of the technologies constitute cheating? Are we downgrading the learning environment by allowing its use? We will come back to these particular points.

Practical Implementation Issues and Risks

An educational institution may have the resources and intent to dive into the AI space but intent is not the same as implementation. Let's look at some more granular implementation issues from the perspectives of four stakeholder groups: (i) the educational institutions themselves, (ii) faculty, (iii) students and (iv) the accounting profession.

Educational institutions (colleges and universities):

What are some of the other big issues that educational institutions are grappling with in terms of implementation? There are many but I have outlined a few below (some of which are interrelated).

1. *Legitimate use of these technologies.* How do we define "legitimate use" of technologies such as generative AI (in theory and in practice)? At a granular level, this is a moving target. There are many aspects to consider, including issues related to academic integrity and the quality of the learning environment. Does the use of these technologies constitute cheating[15]? Is the learning environment degraded if we do not help students to understand how to properly use the technologies in a safe, ethical and effective manner?
2. *The academic integrity quandary.* As already mentioned, academic integrity is paramount in education. Many believe cheating is more prevalent with the advent of AI and other such technologies.[16] Some universities and colleges have taken the position that use of AI cannot be reliably detected.[17] Where does that leave us? If we cannot reliably detect it, how can we effectively ban its use? Banning its use (if even possible) locks us on a path where we have to police it. Those who have been involved in academic integrity cases understand how this creates significant challenges (i.e., if we cannot prove inappropriate use).
3. *Ubiquity of AI*. Following along on the academic integrity issue just discussed, given that AI is embedded in many commonly used applications, it is becoming increasingly difficult even for students to identify when they are using it. Where do we draw the line in cases where AI

is fully integrated into an application? Commonly used applications like Word, Excel, PowerPoint, Outlook have it embedded. If Excel analyzes trends and summarizes charts, for instance, is that problematic for students in a "no AI" academic environment? Students might not even know that they are using it, nor understand why they cannot use the technology to go that extra step without getting in trouble with their professor. That type of ambiguity increases student cognitive load unnecessarily.

4. *How to incentivize use of AI in a responsible and safe environment.* Just what might those incentives look like? How would they change traditional practices like hiring and promotion decisions? We will examine the incentives issue from the faculty perspective next.

Accounting Faculty (Research Stream and Teaching Stream):

What are some of the concerns that faculty might have?

1. *The expertise conundrum*. Accounting educators, like those in other disciplines, are experts in their field. Traditionally, their field represents a specific area of research or practice, such as financial reporting or assurance. All educators are respectful of expertise in their chosen field, but for many, including those in the area of accounting, their expertise may <u>not</u> necessarily be in the area of AI. This creates a challenge. To add to this, accounting professors are often not involved directly in the current practice of accounting and may never have even worked in the field at all. Therefore, they may feel inadequately informed when it comes to understanding how these technologies are being deployed in practice – and rightly so. So, the intent to move forward may be there but how do educators get up to speed?

2. *The pedagogy design issue*. Most professors are not experts in pedagogy design. As noted above, their expertise lies in other areas – for instance, tax or assurance. Engaging in a complete overhaul and reframing of learning within the context of an AI environment is understandably a daunting task.

3. *Incentives*. Incentives are not always there to invest in reimagining the curriculum in courses or programs – especially for younger research faculty who are on the tenure clock. Major curriculum change is very labour-intensive and takes away from research activities (which are critical in tenure promotions). In addition, a major change in pedagogy comes with it a risk of lower student ratings. In the year of a major curriculum overhaul, there are bound to be a few bumps along the way. Student ratings impact promotion and tenure decisions so you can understand why some faculty are concerned about curriculum overhaul. Many institutions employ a number of part-time or sessional faculty to deliver courses. Where does course redesign fit in that model?

4. *Lack of clarity on course and program objectives.* What is our overall objective with regards to our students? Is our goal to produce job-ready graduates or flexible, curious thinkers who are tech savvy and can contribute to a constantly changing society? Jobs and graduate roles keep changing. Without knowing where things are going, the water becomes muddied.

5. *Student usage*. Many students are using the technologies and some students are using it more than others, creating an unequal playing field.[18] How can educators assess students fairly when this is happening?

***Accounting Students (Graduate and Undergraduate)*:**
Here are few concerns that students might have.

1. *Student perception regarding use of AI*. Many students have expressed the concern that they feel like they are cheating when they use these technologies but that they have to do it to keep up with their peers.[19] This is not a healthy sentiment to allow to prevail in our classrooms.
2. *Culture shock*. Assuming students are already using the technologies in their everyday lives (as many of us are likely doing as well) and perhaps in courses they have previously taken, they might experience a culture shock if not allowed to use the technologies in the particular course they are currently in. Many students have successfully incorporated using the technologies into their individual learning strategies.[20]

The Accounting Profession:
As professionals, accountants carefully weigh the risks and opportunities when considering how to adopt these technologies. While efficiencies may certainly be gained, the risks are high. For instance, in accounting, client confidentiality is paramount and must be guarded. Given the nature of the profession, let's consider trust.

Accounting firms are integrating the technologies in a way that is safe, responsible and ethical. They have no other choice, if trust is to continue to be a value added.

Trust is a hallmark of the profession. People trust their accountant and should be able to – even in the age of AI. By definition, accounting has been historically seen as conservative, but because accountants look after the financial affairs of others and deal with confidential information all the time, this is not a negative thing. Accountants need to continue to act with due care and competence in this AI-driven world. Although the profession needs to proceed with caution, it is nonetheless moving forward, embracing AI technologies where prudent.

Accounting firms are integrating the technologies in a way that is safe, responsible and ethical. They have no other choice, if trust is to continue to be a value added. This framing of the trust issue is important for accounting students to understand and appreciate. It is a framework that we can use in education.

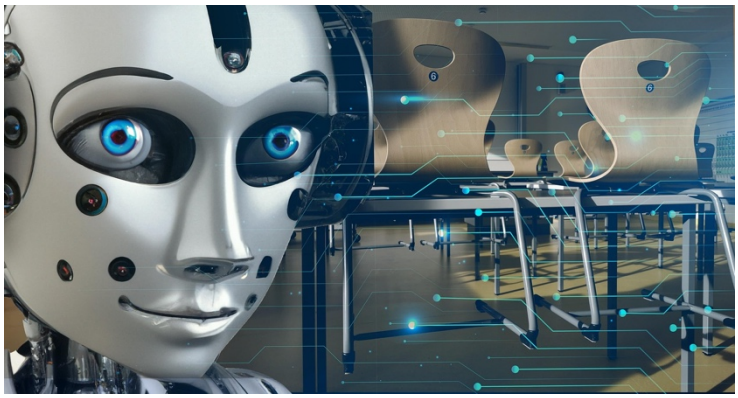
Let me close with a call to action, acknowledging the significant headway already being made in many accounting courses and programs at colleges and universities.

Call to Action

For those not yet mobilized, I have listed a few things/ideas to think about and get the ball rolling. They are in no particular order and assume that resources exist or can be found to support the initiatives. As noted earlier, the resource issue is a considerable one so this might be the starting point (perhaps beginning with identifying government funding opportunities). Some of the items below can be started with meetings or getting a working group together.

For institutions:

1. ***Define "legitimate use" of these technologies.*** Start the conversation. The more you can engage faculty, staff and students in the conversation the better. Legitimate use will likely be defined differently across institutions but there needs to be consistency within the institution. Landing on policies is only half the battle. A big part is in effectively communicating them – our next point.



2. ***Create and widely publicize AI policies – the AI tone at the top*****.** Policies can be high-level statements, frameworks and/or more granular procedures. Perhaps the overarching policy might be to "promote the ethical, safe and responsible use of AI" for instance. These policies should be positioned as "living" policies, given the changing nature of the landscape. Once the policies and procedures are identified, they must be communicated widely. People need to understand and acknowledge them. They must be displayed prominently (on the institution website and on course outlines) and workshops held to support dissemination and ensure understanding. Everyone needs to be on the same page.
3. ***Review/create incentives to support and promote effective adoption.*** Existing policies may need to be revisited considering how they might support or inhibit uptake. These include policies related to tenure, student ratings, promotions, annual evaluations. New incentives may need to be put in place.
4. ***Ensure adequate support and safe access to the technologies*****.** Faculty need to get up to speed, and may need help with pedagogy design. Students, staff and faculty need a safe environment (with access to the technologies) within which to experiment (perhaps a walled garden or AI kitchen, which lives behind the university or college firewalls). There is a need for training courses, communities of practice and AI "help lines." Many courses, conferences, workshops and communities exist beyond our own institutions and are freely accessible. Make use of these external resources, especially if internal resources are constrained.

5. ***Revisit what we are trying to do*.** As pointed out earlier, just what is our role as accounting educators and educational institutions? Are we there to create "job-ready" graduates or, given the speed at which the technologies are advancing, should we be producing graduates that are tech savvy, flexible, creative problem-solvers and critical thinkers? Maybe these are not so different. We need to determine what our role is in the age of "thinking" AI.

We have to be mindful that the allure is there for students to offload the learning but if we are not producing students who can think critically, then what are we doing?

For faculty:
If we are to model the behaviour that we seek in our graduates, i.e., competence, respect, accountability, trust and ethics, we cannot shy away from using AI ourselves. We must help our students learn to use the technologies wisely, understanding the risks as well as the opportunities, all while upholding academic integrity.

If you have not taken the first step towards embracing these intelligent technologies, just start somewhere. There are resources available (either within your institution or externally). For instance, there are AI study buddy platforms available to support your students.[21] These can help meet students where they are in their learning journey, providing individual feedback and support.

Once you start down the path, you will find that it is not so daunting and might even enjoy where it takes you! Your excitement will invigorate the students, as will working with these technologies. Some universities and colleges acknowledge and accept that student ratings might be lower in a year of curriculum reform – find out what the policy is at your institution. Ask about the impact on your tenure decision if you are close to tenure.

Remember it is our job to help students navigate these turbulent waters. We have to be mindful that the allure is there for students to offload the learning but if we are not producing students who can think critically, then what are we doing? Finally, speak with employers to see how they are dealing with these technologies. What do they expect from our graduates?

To conclude, we must keep our eye on the real problem which is this: if we do nothing and let the students continue to use these technologies unsupervised/unchecked, their ability to think critically and creatively will decline, cheating rates will increase and we will not be doing our students any favours. Students need learn how to use the technologies in a safe, responsible and ethical manner so that they become professional accountants who engender trust.

End Notes:

1 I can remember, for instance, having many spirited discussions about whether students should be allowed to use the technologies at all.

2 For instance, we would not have thought possible an AI agent who can source products for us and complete the transaction (with final approval still resting with a human – for now).
3 An example of this might include having students utilize AI technologies in the classroom as study buddies/group members. The AI study buddy might perform basic research or financial analysis. This would be the starting point. The student groups would then check the work of the AI study buddy, review source documents in detail and/or perform independent analysis to corroborate or disprove the AI study buddy findings. Doing the latter allows for quite a different learning experience than performing the basic research and analysis. As an aside, students also need to know how to do basic research and analysis. They need to know enough about it before they delegate it to the AI buddy so that they can then critically assess the output. Having the AI buddy provides efficiencies, once the students have mastered the basics, allowing them to focus on the development of other things such as skepticism and accountability.
4 My belief is that learning will indeed be degraded, if the technology is just tacked on to traditional pedagogy and not properly monitored or supported. Cognitive offloading is a real threat. It may be fine for students to use the technologies for basic non-critical tasks but tasks meant to build skills must not be allowed to be delegated to AI, at least until the students have mastered them. For the technologies to support learning, we must consciously design courses and programs to ensure cognitive offloading of important tasks does not happen.
5 For instance, in Canada, the U15, an association of fifteen leading universities across Canada, recognized quite early (in 2024), the significant impact that technologies would continue to have on education and developed shared values, principles and leading practices. See https://u15.ca/wp-content/uploads/2024/08/Navigating-AI-in-Teaching-and-Learning-U15-Canada-September-2024.pdf. At that point, they suggested the following as leading practices: clear communication for faculty, staff and students about the approaches taken in terms of AI; continued support and collaboration; regular dialogue; and need to stay abreast of changes. The member institutions have continued to advance these practices. A number of major North American universities who produce accounting graduates including, for instance, in Canada, University of Toronto, University of Waterloo, Wilfrid Laurier (all part of the U15), and in the United States, Indiana University Bloomington, University of Illinois Urbana-Champaign and University of Texas at Austin currently have very supportive policies around use of these technologies in research and teaching. Some universities and colleges have created AI enabled spaces, including safe, walled-off environments for faculty and students to learn to use the technologies safely, responsibly and ethically.
6 Each university or college has unique ways of presenting its budget, but budgets of publicly funded institutions are available through university and college websites. Many highlight funds spent on these technologies. As an example, McGill University, in its 2025-26 budget reported that the Quebec government had committed $22 million to continue to fund MILA, a joint venture between the Université de Montréal and McGill University created to further understanding of AI. See page 7, https://www.mcgill.ca/apb/files/apb/budget_book_fy2025-2026_-_en_0.pdf.
7 At the national level, in 2025, the Canadian government introduced its *National Artificial Intelligence Strategy: AI for All*, a very comprehensive document. In the document, the government committed, amongst other things, to ensure all post-secondary students have access to trusted AI agents and free AI literacy training. The strategy also committed to supporting the building out of compute infrastructure to increase AI research and expanding the Canadian CIFAR AI Chairs program from 130 professors to 200. See https://ised-isde.canada.ca/site/ised/sites/default/files/documents/ai-strategy-en.pdf.
8 These three universities ranked in the top 10 (globally) in terms of AI according to the 2026 Global Ranking of Academic Subjects. See https://www.shanghairanking.com/rankings/gras/2026/AS0229.
9 Under the Government of Canada AI Sovereign compute infrastructure program, the Canadian government committed $890 million in available funding. See https://ised-isde.canada.ca/site/ised/en/ai-sovereign-compute-infrastructure-program.
10 https://fao-on.org/en/report/estimates-2026-mcures/.
11 In Ontario, the government mandated a 10% cut in tuitions in 2019 and imposed a freeze on tuition until 2026 when it was lifted.
12 The Government of Canada has limited the number of student visa permits for international students. See https://universityaffairs.ca/news/international-student-visa-permits-what-to-know/.
13 For a general discussion of academic freedom, see https://www.aaup.org/academe/issues/105-1/academic-freedom-canada-labor-law-right.

14 Acharya, B., Sharma, B., and Gajurel, D. (2025). GenAI Policy in Canadian Universities: Pedagogical Problems Faced by Professors. *Intelligent Technologies in Education*, Advanced Online Publication. See https://www.researchgate.net/publication/396423606_GenAI_Policy_in_Canadian_Universities_Pedagogical_Problems_Faced_by_Professors.

15 In a study of 1057 faculty in higher education (all fields), there was significant disagreement as to what was considered cheating. For instance, faculty considered using AI to brainstorm was legitimate but using it to create a detailed outline for a writing assignment was considered cheating. In many cases, faculty were unsure as to whether the scenarios posed in the study were cheating or not. See Watson, C.E. and Rainie, L., *The AI Challenge: How College Faculty Assess the Present and Future of Higher Education in the Age of AI*, AAC&U and Elon University, January 2026 at https://www.aacu.org/research/the-ai-challenge.

16 A recent study by KPMG (of 1009 participants) noted that the top three concerns educators have relate to academic integrity or cheating, student over-reliance on AI and copyright and plagiarism issues. See https://kpmg.com/ca/en/media/2026/09/majority-of-students-use-ai-but-educators-sound-alarm.html.

17 See for instance the academic integrity policy at the University of Saskatchewan https://academic-integrity.usask.ca/gen-ai.php. This university is also a member of the U15 mentioned earlier.

18 See https://kpmg.com/ca/en/media/2026/09/majority-of-students-use-ai-but-educators-sound-alarm.html.

19 See an interesting discussion on student cheating in this article: https://hiwavemakers.com/blog/how-many-students-really-cheat-with-ai-less-than-they-think/.

20 If someone told me now that I could not use these technologies, I might feel a bit miffed. I am very careful how I use the technologies but I do find them most useful when used properly.

21 As an example, see https://alldayta.com/.

Beyond Sustainability to Systems Thinking

By Richard Spencer



Richard Spencer, an Executive Partner at Symbion Partners, is a sustainability leader, facilitator and adviser who helps connect the worlds of business, finance and public policy. With over 25 years' experience in investment banking, strategy consulting and leading on sustainability at ICAEW, he has worked with boards, finance teams, policymakers and sustainability professionals to embed sustainability into strategy, governance and decision-making. Drawing on extensive experience in thought leadership, capacity building and stakeholder engagement, Richard helps organisations understand the business case for sustainability and turn ambition into practical action.

In this paper I ask: Has the business community backed away from sustainability and at a moment when its engagement is needed most? I look at the high point of business engagement in the early 2020s and compare it with today, which looks like a complete retreat. However, my answer is still "probably not." There is reason to think much is going on under the surface. This may even be the beginnings of a real embedding of these big systemic challenges into business decision-making. A shift from climate and nature risks being regarded as extra to, and context for risk management to a recognition of them as drivers of business risk and value creation/destruction.

If, however, business is going to play its part in achieving the goal of a sustainable planet with prosperous economies and thriving societies (if it is going to connect everyday business decisions with big systemic risks), then its tool kit is inadequate. We need to develop systems thinking as a core skill in the professions and amongst business managers and a systems-mindset as the cultural norm.

The 2020s Has Seen Nature and Climate Change Go from Being the Key Issue Across the Economy to Taking a Back Seat

Between 2021 and 2023 sustainability in business as a topic was riding high. There was an array of frameworks and tools in production and/or promotion; businesses and governments (national, regional and municipal) were declaring a climate emergency and competing to set the most ambitious Net Zero targets; and businesses started using nature positive as a target setting framework rather than a slogan.

The Kunming Montreal Conference of Parties (COP15) of the Convention on Biological Diversity in 2022 was something of a highpoint of action on loss of nature. COP15 saw the adoption of the Global Biodiversity Framework after four years of consultation and negotiation. The Framework set out "an ambitious pathway to reach the global vision of a world living in harmony with nature by 2050. Among the Framework's key elements are 4 goals for 2050 and 23 targets for 2030." (Kunming-Montreal Global Biodiversity Framework) Target 15 seeks to embed nature-related considerations into business decision-making and reporting.

COP15 followed the launch of the Taskforce on Nature-related Financial Disclosures (TNFD) in June of 2021. TNFD was a market-led, science-based initiative to develop a framework for

organizations to identify, assess and act on nature-related risks and opportunities and to disclose them. TNFD published its framework and recommendations in 2023 and that could be seen as the framework to deliver Target 15.

During the same period the International Sustainability Standards Board (ISSB) was launched by the International Financial Reporting Standards Foundation at the climate Conference of Parties, COP26 in Glasgow in November 2021, also a high point, but this time for climate. That was also the COP where the UK government announced the Transition Plan Taskforce, one of the key moments in signalling the move from bold goals to bold action. The ISSB issued its standards in 2023.

In the European Union, again over the same period, the Corporate Sustainability Reporting Directive (CSRD) was published by the Commission in April 2021. It dramatically expanded the scope of the Non-Financial Reporting Directive and introduced mandatory reporting standards. Perhaps most significantly it introduced the concept of Double Materiality, capturing not only how sustainability issues affect businesses financially, but how businesses affect society and the environment.

2026 represents a different picture with sustainability having moved away from the centre ground and into the shadows. Just think about the 1.5c̊ narrative. A clear ambition to fulfil the goal of the Paris Agreement (over 10 years ago), which slipped to a hope (keep 1.5c̊ alive) to a whimper, to this year's UN Emissions Gap Report, which abandons all hope. Previous years' gap reports have been sceptical about the 1.5c̊ target; but this year's is the first time the report has commented that missing it is "imminent and unavoidable" and noted that there are "no good outcomes" of remaining above 1.5c̊.

So, the narrative has shifted from preventing a breach to managing an overshoot and the consequences of that. But there is no going back, the damage will be irreversible. An ICCI/UN report published on November 2, 2026 (Overshoot and the Cryosphere: averting global meltdown), was deeply critical of the "persistent myth" that a temporary breach is preferable to immediate, painful emission cuts. And from a business point of view, this is about catastrophic economic loss of value (there is so much on this, from the UK's Office for Budget Responsibility's annual *Fiscal Risks and Sustainability* (FRS) report to the Institute and Faculty of Actuaries report Planetary Insolvency). And the glaring absurdity is that as things are getting worse, we take them off the main agenda.

What Happened?

Things shifted substantially with the change in mood music with the election of President Trump in 2024 and the US withdrawal for a second time from the Paris Agreement. This shift can be indexed against the Securities and Exchange Commission's position on sustainability reporting. In March 2024, the SEC adopted rules requiring public companies to disclose certain climate risks, governance arrangements and, for some larger entities, Scope 1 and 2 emissions, where material.

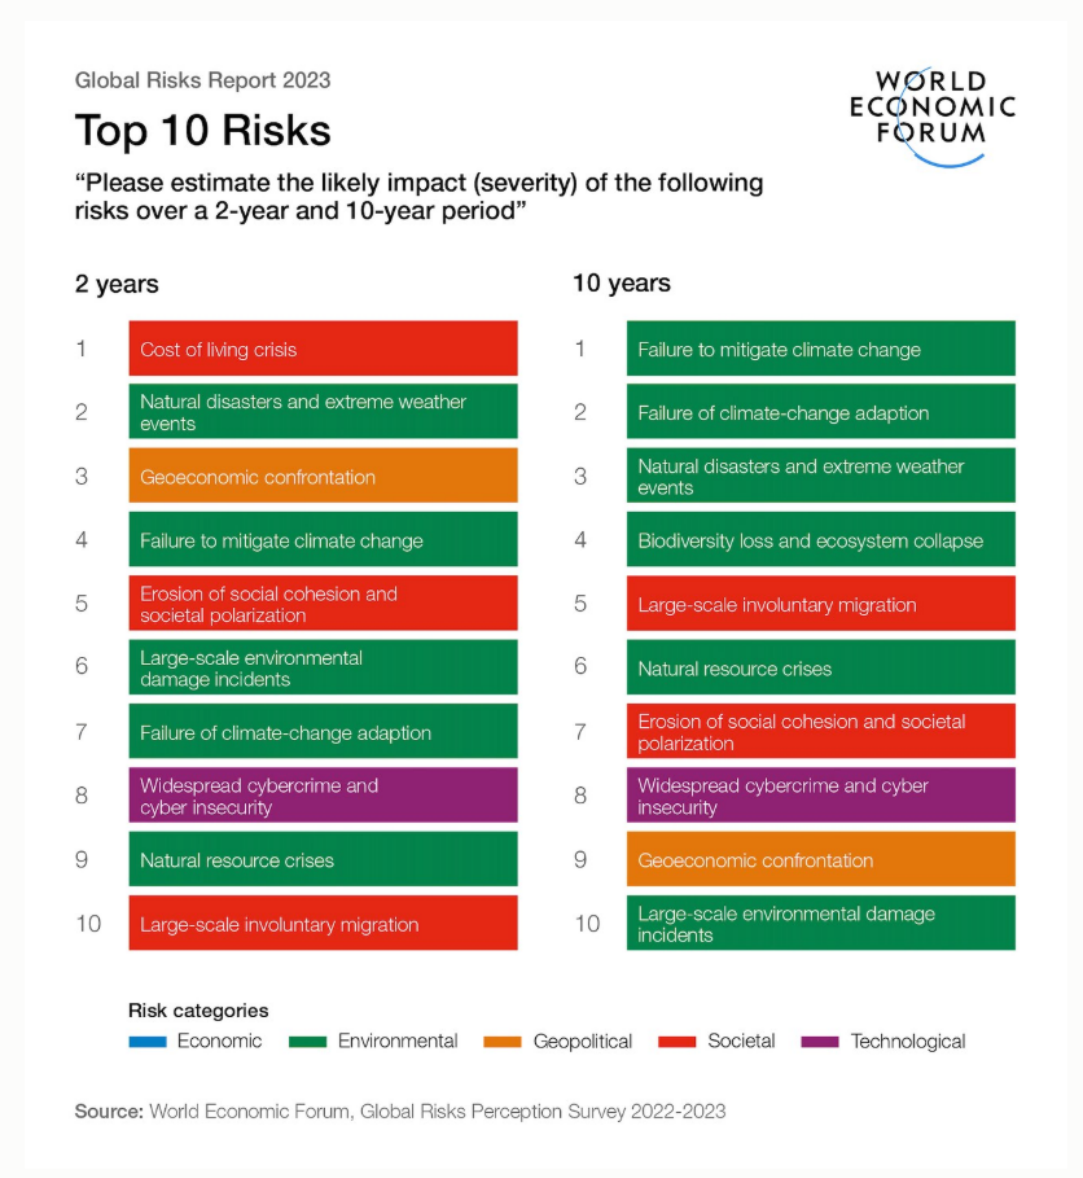
With the CSRD, by 2024 there was growing criticism that the system was too complex, required too many data points, was costly to implement and would damage European competitiveness.

The rule never became operational, however. Multiple challenges were filed immediately after adoption. The SEC stayed the rule in April 2024 pending litigation. In March 2025, the

SEC stopped defending the rule in court and, in May 2026, formally proposed rescinding the climate disclosure rule altogether. The practical consequence is that the US currently does not have a federal climate-reporting regime.

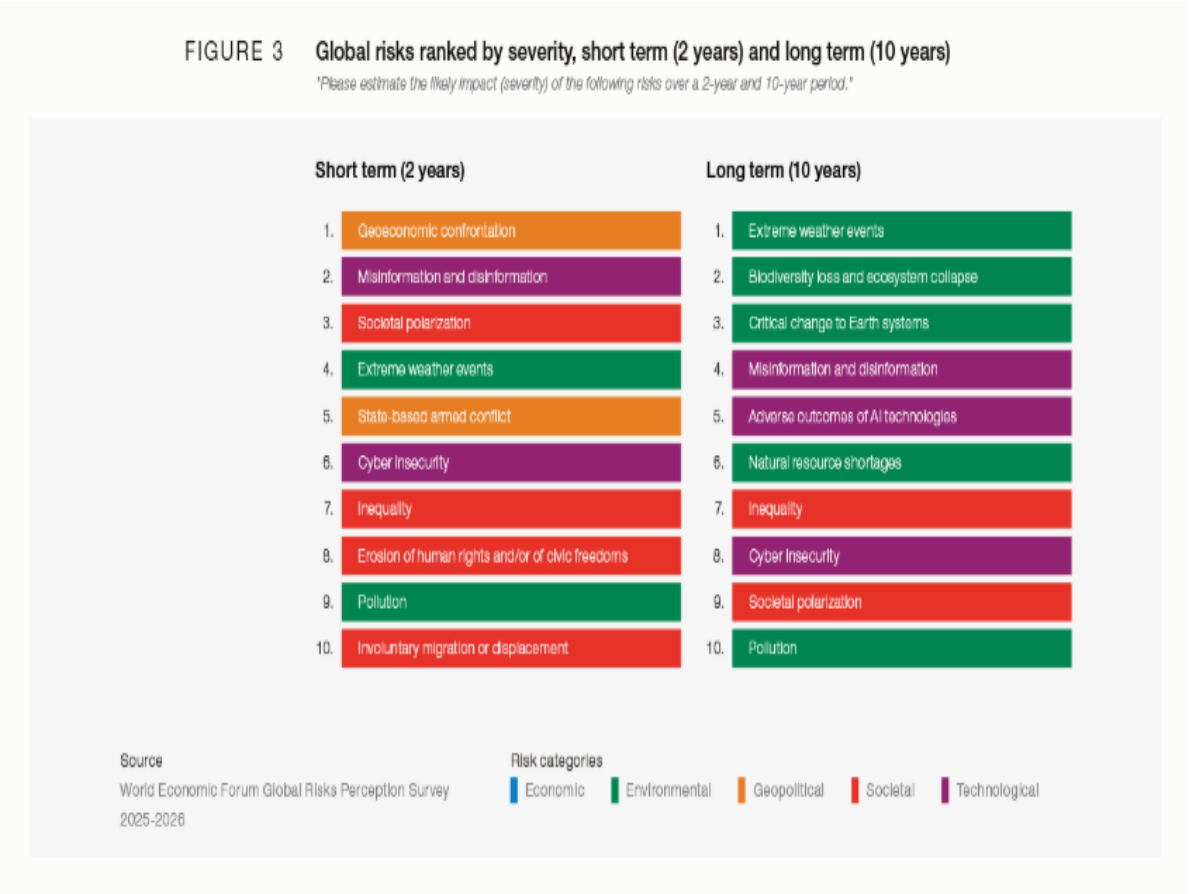
With the CSRD, by 2024 there was growing criticism that the system was too complex, required too many data points, was costly to implement and would damage European competitiveness. This led to the Omnibus changes that dramatically reduced the scope (some estimates suggest 80% of companies originally caught in the regime are now removed) and the reporting standards (ESRSs) have been simplified.

This pattern is mirrored in the business world's risk focus. The World Economic Forum's 2026 Global Risk Report notes that, in the near term, "not only do their [nature-based risks] rankings decline relative to other risk categories, but there has also been an absolute shift away from concerns about the environment." This is a different picture from 2023 when environmental risks were all over the top 10 short-term risks.



Geoeconomic confrontation and state-based armed conflict now dominate the near-term risks (two years) with nature-based risks being driven out. The Extreme Weather Events entry sits alone in the top 5, at number 4 on the list of risks ranked by severity, with Pollution at 9.

The long-term view of Nature-based Risks remains the same. They still top the table 10 years, taking the first three slots (Extreme Weather Events, Biodiversity Loss and Critical Change to Earth System, respectively) the sixth spot and the tenth. The problem with short term vs. long term risks is that it simplifies the issues. It inherently assumes those risks can be managed at that point in time. Failure to Mitigate Climate Change being presented as a 10-year risk does not, however, seem to consider that we might cross a tipping point/threshold before then. Meaning that it becomes financially material for companies in 10 years' time, but by then they will not be able to reduce that risk. They are more like 10-year probabilities if they are dressed-up as future risks but we do nothing now.



How Do We Make Sense of This?

It feels like there are a few things at play, but there is a common denominator: the absence of a systems-based approach. To solve these huge systemic problems, we need business managers and service providers to have a systems mindset.

This isn't by any measure a new idea. It is nearly 40 years ago that Peter Senge wrote *The Fifth Discipline* in which he made the case that organizations fail not because their people lack talent, but because they are constrained by fragmented thinking, short-term reactions and unexamined assumptions. Long-term success requires becoming a learning organization.

That is, one that continually seeks to understand systems, adapt and innovate. At the heart of this is systems thinking, which enables leaders to see patterns, relationships and underlying structures rather than isolated events. It shifts organizations away from reactive problem-solving and toward generative learning, where people change how they think, not just what they know.

Environmental issues are seen separately, as context and lesser, the first to be deferred
While acknowledged as important, environmental risks are often framed by business as context rather than as something to be acted on.

"Boards continue to assess threats in isolation, separating financial risk from environmental risk and geopolitical risk from corporate strategy. Climate and nature are often delegated to sustainability functions, detached from core decisions on growth, capital allocation and risk appetite." (Governing environmental risk in an era of geopolitical instability, Emily Farnworth and Alice Ruhweza, WEF, 2026)

They are future issues, secondary concerns to economic and political instability to be managed later. In fact, they should be seen as cross-cutting drivers of all risks:

"Climate change and nature loss are not parallel challenges to conflict or economic volatility; they are deeply intertwined, creating reinforcing feedback loops that operate across borders and time horizons." (Emily Farnworth and Alice Ruhweza, WEF, 2026)

Even though WEF (as the quote above demonstrates) extensively publishes on embedding environmental risks and on the importance of interconnectivity, the risk report is primarily known for its tables of near and far terms risks, eagerly awaited each year. There is a sense in which this reinforces the idea of them as distinct and separate on a list, rather than drivers of the uncertain and chaotic world we know.

This is compounded by a compliance mind-set and what has become known as framework fatigue
There is a body of academic literature on how frameworks and standards drive a compliance mindset

The last two years has seen considerable commentary emerge on "framework fatigue" in sustainability and nature reporting. Because the term is so new, it is mainly in the non-academic, practitioner literature that has been discussing this problem extensively. The theme appears in consultancy reports, professional bodies, assurance providers and business press commentary.

The term framework fatigue is increasingly used in this context to describe organizations' exhaustion with the proliferation of standards and frameworks. Although it is non-academic, it sits within established academic literature that covers what might be called the creation of a compliance approach or mindset. These include terms such as institutional isomorphism (DiMaggio and Powell, 1983): tendency of organizations operating in the same environment to become increasingly similar over time; regulatory overload; and symbolic compliance: organizations comply with the form of a requirement without fully embracing its substance (Boiral, 2013), (Cho et al., 2015) and Michelon, Pilonato & Ricceri, 2015). And I'd add to that Bebbington & Larrinaga (2014), which discusses the risk of sustainability accounting frameworks becoming "technocratic compliance regimes."

The term framework fatigue is increasingly used in this context to describe organizations' exhaustion with the proliferation of standards and frameworks.

This compliance mindset, coupled with framework fatigue, has I would argue compounded the motivation to put off acting on nature loss and climate crisis.

It has become a significant theme in the practitioner community. Businesses are experiencing framework fatigue as they attempt to comply with a landscape characterized by overlapping frameworks, shifting expectations and increasing assurance demands.

Across Big Four publications, fatigue is identified as a top operational risk. EY's *Global Sustainability Reporting Barometer* (2025–2026) reports widespread "reporting congestion," with organizations struggling to keep pace with the volume of new requirements. PwC's *State of Sustainability Reporting 2025* highlights "framework fatigue" as a major barrier to effective governance, noting that companies are layering TNFD on top of TCFD and ISSB rather than redesigning strategy. KPMG's *ESG Assurance Maturity Index 2026* describes "client exhaustion" and finds nature disclosures to be the most compliance-driven. Deloitte's *Sustainability Reporting Outlook 2026* warns that the pace of regulatory change is "unsustainable" for many teams.

Other think tanks/organizations reinforce these findings. A4S (2025–2026) reports that businesses are treating TNFD as a reporting exercise rather than a strategic tool, with CFOs delaying investment until ISSB and TNFD consolidate. WBCSD's *TNFD Early Adopter Insights* (2025) shows companies overwhelmed by the number of frameworks they must align with, often implementing TNFD's LEAP process as a template-driven workflow. PRI (2025–2026) notes that investors are receiving boilerplate nature disclosures due to reporting overload and calls for a single coherent baseline.

Regulators and standard-setters acknowledge the issue. ISSB's nature practice statement consultations (2025–2026) reveal concerns about reporting burden and framework overlap. EFRAG's ESRS feedback (2025) highlights significant fatigue, especially for nature standards. The UK FRC (2025) reports similar concerns during UK SRS consultations.

There is a broad consensus that framework fatigue is driving a compliance mindset, particularly in nature reporting, rather than a strategic response. Capability gaps and issue complexity make strategic integration more difficult. Instead, businesses are building reporting teams, hiring consultants, mapping overlaps, creating compliance matrices, focusing on deadlines and assurance and, thereby, making reporting ever more sophisticated, while strategic decision-making changes very little.

Reasons to be Cheerful

While sustainability as a topic may have lost its shine and moved from the centre ground, and while there may be much less public discussion by businesses about sustainability and while anecdotally there is talk of a tighter market, that doesn't mean the issues of climate change and loss of nature (as well as growing inequalities) aren't being addressed. The absence of evidence should not be taken as evidence of absence. In fact, the silver lining to the public scaling from sustainability maybe that it is being integrated into organizations as business issues. And, while that doesn't mean a wholesale shift to a systems-based approach, it feels like a move in the right direction.

Several analyses suggest that organizations are increasingly moving sustainability work into mainstream functions such as finance, risk, procurement, operations and strategy, rather than maintaining separate ESG departments (KnowESG *The Global ESG Jobs Crunch: What It Means for Companies and Professionals*; BCG *The Climate Transition Is Redefining Work. Do You Have the Right Talent Strategy?*).

In particular, BCG states that 25% to 50% of roles across sectors will need to evolve to meet sustainability priorities, but that most of these are existing jobs rather than new "green jobs." BCG also found that specialist sustainability roles account for only about 5% of sustainability-related positions and that sustainability work is increasingly concentrated in core functions such as finance, risk, procurement, operations and technology.

So, the market for ESG/Sustainability managers has softened, but more specific sustainable finance, carbon accounting, reporting and sustainability-finance roles are in demand. Many businesses are building the capacity internally, replacing "Head of ESG with a team of specialists" models with "finance, risk and operations staff who have sustainability expertise." I'd actually say the market has shifted from "ESG advisers" toward "people who can connect sustainability, risk, finance (P&L, balance sheet, cashflow, valuation), governance and transition planning."

This is backed up by the LinkedIn *Global Green Skills Report 2026*, which shows “green” recruitment is growing faster than the supply of green talent; that 74% of green hires are being recruited into roles without traditional green job titles; and workers with green skills are hired at a higher rate than the wider workforce.

In addition to this trend, and while practice around the TNFD framework and associated tools (TNFD Tools Catalogue lists more than 200) is early and emerging, using the framework holistically will feed into strategic thinking through the use of scenarios and integrating nature into corporate risk frameworks. A measure of how the market is moving in this direction is Lancaster University’s business briefing on how these frameworks and tools can be used together in a systematic way (Business Briefing: Tools and Tool Stack for Nature Reporting Neytullah Ciftci). The Pentland Centre at Lancaster has also developed a suite of resources for nature scenario analysis (Biodiversity scenario guides for business - Lancaster University), also a measure of increasing sophistication in the business community.

To Drive a System-Based Approach in Markets, We Need the Right Policy Conditions
We also have agency, we are not hapless victims and we can change the rules of the game.

Climate and nature policy is heavily reliant on market-based instruments such as carbon pricing, voluntary markets, disclosure frameworks, ESG ratings and private finance mobilization. This tends not to drive strategic transformation. Firms optimize for reporting requirements, ratings methodologies and transactional carbon practices, while underlying business models and real-economy activities remain largely unchanged.

There is plenty of commentary on the need for collective action and the right policy environment. One example is *The Architecture of Transition The case for Transition Governance*, a report from the new Finance Transition Centre led by Steve Waygood. Another is *Bridging the Climate Risk Disconnect* by the Environmental Defense Fund and D A Carlin & Company.

But what would this look like? Again, thinking from a system-based approach helps.

Roland Kupers' work on complexity, resilience and public policy is helpful for understanding why this compliance mindset arises (Kupers, 2014; Kupers & Wilkinson, 2014). Kupers argues that markets are complex adaptive systems, not equilibrium-seeking mechanisms, and therefore cannot reliably coordinate long-term collective action or steer socio-ecological systems toward sustainability. Climate and nature systems are non-linear, threshold-driven and interdependent, yet market instruments assume rational responses to price signals or risk disclosures. This mismatch between tool design and system behaviour makes compliance a predictable organizational response: firms follow the rules of the system they are given.

Kupers' complexity approach points to a different set of arrangements. Effective climate and nature action requires state-enabled collaboration, shared narratives and expectations, scenario-based strategic coordination and adaptive policy frameworks that reflect the dynamics of complex systems. In this view, governments act not merely as price-setters or information regulators, but as architects of collective action capable of shaping markets, convening actors and guiding long-term transitions. This would result in a strategy-led system in which reporting isn't compliance but a necessary information flow to the markets.

Mariana Mazzucato's work in *The Entrepreneurial State* feels to me like it supports this argument. Using the four key technologies of the smart phone each independently developed by the US military, she demonstrates that innovation has historically emerged not from markets responding to price signals. Instead, they emerged from a mission-oriented state that created and shaped markets, coordinated actors and absorbed risks that private firms would not embrace.

Mazzucato has recently extended this critique of individualistic and market-centred conceptions of value. In *The Common Good Economy,* she argues that:

"Communitarian, biological and Indigenous perspectives reveal that individuals do not exist in isolation, but within interdependent social and ecological networks. Identity, wellbeing and even the generation of value in the economy are inseparable from these relationships. This challenges the idea of value as something produced only by individuals and instead points to collective value creation, in which markets are shaped by shared institutions and common objectives."

When the state limits itself to correcting market failures or providing information as in market-led climate governance, it produces weak, incremental incentives.

Value is not generated by firms acting independently within markets, but emerges through relationships, institutions and the health of the wider social and ecological systems on which economic activity depends. Climate change and nature loss are, therefore, not externalities to be managed alongside value creation; they are integral to the conditions that make value creation possible in the first place.

When the state limits itself to correcting market failures or providing information as in market-led climate governance, it produces weak, incremental incentives. Firms rationally respond through reporting, risk management and compliance rather than structural change. Market-led tools ask markets to deliver missions that only states have ever been able to lead.

Taken together, Kupers and Mazzucato offer a coherent alternative to the current arrangements. Kupers offers an analysis of why climate and nature crises require coordinated, systemic interventions that markets cannot deliver on their own. Mazzucato provides the institutional logic for a mission-oriented state capable of shaping markets, mobilizing innovation ecosystems and enabling long-term transitions. The present market-led architecture has reached its limits: a complexity-informed, mission-oriented governance model offers a pathway beyond compliance toward genuine transformation in both climate and nature outcomes.

We Need to Become System Thinkers and System Actors, Developing New Ways of Thinking and Business Models that Recognize an Interconnected World; Business Needs to See the World Differently

The ability/appetite to think in systems, to reach for an understanding of connectivity and feedback loops, embracing uncertainty and to let go of a linear way of perceiving the world is needed to grasp the importance of nature to economic success. Businesses need to be able to connect the dots between big systemic issues and day-to-day strategic and operational decisions. They have to be both micro stewards (of their businesses) and macro stewards (of the system), recognizing environmental risks as objects of corporate risk management and as drivers of systemic change. In so doing they must not prioritize adaptation and resilience over rapid mitigation of climate change and nature loss (see Salla Kaarina Pasula, 2024). Also, see “From a business risk to full-scale crisis: the understanding of climate threat” in the *World Economic Forum’s Global Risks Reports*, Globalizations, 21:7, 1197-1214, for a critique of a business dominant narrative). If they cannot marry these two positions, then the problem – nature-based risks – will remain obscure and distant.

There is a wealth of commentary on the need for new thinking and business models and new system-based thinking. At the heart of this is the point that we cannot continue to assume the business model is a given and ask what we can do better given the crises we face. That is an argument of incremental improvement and efficiency. Instead, we need to understand the crises as givens and ask how we change the way we do business to solve those crises. For example, see the EU project A-Track led by the Capitals Coalition, the *Accounting for Sustainability Essential Guide to Managing Future Uncertainty*, and the Global Accounting Alliance sustainability materials, in particular the *Connectivity in Reporting* series and *Why Nature Matters to Accountants*.

One example that really takes a systems-based approach, aimed at the management accountant is *Reframing Business Models for a Sustainable World*, by Martin Farrer at CIMA. The context is significant here: CIMA represents the management accountants, one important group who in all probability will be making nature risk visible and embedded in business.

This report argues that traditional business models, build around short-term financial performance and efficiency, are no longer fit for an era of interconnected global crises. Organizations need to move beyond a narrow focus on shareholders, efficiency and short-term financial performance and recognize that they operate within a wider system, with prosperity ultimately dependent on healthy societies and planet. Instead, business models must be redesigned to create value for people, planet as well. The report calls this process "reframing" business models.

The report introduces concepts such as polycrisis, systemic risk, macro-stewardship and antifragility (the ability not just to with stand shocks but to evolve) to help organizations navigate an increasingly uncertain future. It argues that finance professionals have a critical role in leading this transformation by integrating sustainability into strategy, risk management and decision-making. Rather than simply helping organizations survive disruption, the goal is to create business models that can adapt, evolve and become stronger through uncertainty while creating long-term value for both the organization and society.

This all requires new skills and a change in mindset (a term with a lot of currency these days). As this article has frequently returned to, businesspeople of all types, not just leaders, need to develop systems thinking as a core skill, along with strategic foresight, critical thinking, communication and influencing skills alongside their technical expertise. The CIMA report encourages a move away from deterministic and compliance-driven approaches toward greater curiosity, adaptability and creativity. Leaders must become comfortable working with uncertainty, engaging diverse stakeholders and considering long-term societal and environmental outcomes. Ultimately, the report calls for a mindset centred on stewardship, continuous learning and purposeful value creation, enabling organizations to respond effectively to complexity and build more sustainable and antifragile business models.

A core issue throughout this essay is the need for businesses to move beyond understanding and managing organizational risks, what Aviva calls micro-stewardship. Businesses need also to focus on and manage systemic risks, macro-stewardship (Aviva). This may well not be possible as it may also require operating with government (see above on policy) and in collaborations. These collaborations may not just be sector-specific but could be from across the economy involving the public and third sectors. This is an approach that recognizes businesses, investors, regulators and society as interconnected participants in shared economic, social and environmental systems.

As the CIMA report notes, many of today's greatest risks, including climate change, biodiversity loss, resource constraints and social instability, are systemic in nature and cannot be addressed through conventional enterprise risk management alone. Organizations, therefore, need to adopt systems thinking, scenario analysis, stress testing and long-term foresight to understand how risks interact and cascade across systems.

Table 5 from the CIMA report helpfully captures the balance between entity risk and system risk that needs to be struck but also connected.

Table 5. The difference between enterprise risk management and systemic risk management

Characteristics	Enterprise risk management	Systemic risk management
Scope	Organisation specific	Organisations, multi-sector, global or national
Focus	Strategic and operational objectives	System-wide stability and antifragile
Risk type	Known, quantifiable risks	Complex, uncertain, interdependent risks
Approach	Risk mitigation and control	Antifragile, adaptation, and transformation
Responsibility	Internal stakeholders	Cross-sectoral collaboration (internal and external to an organisation)

This Will Have Implications for How We Account for Our Business Activities

Accounting is the means by which businesses understand the world. It isn't a passive exercise: it constructs the world and it drives action. Through accountancy businesses understand nature in terms of assets, liabilities, risks and materiality. As a result, people do things and make decisions. It is an interpretation and, like any interpretation, it is a function of power rather than truth. It serves an interest and, in so doing, constructs a particular version of, in this case, nature. (Bebbington and Rubin, *Accounting in the Anthropocene*; Bebbington et al., *Accounting and Accountability in the Anthropocene*; Hines, *Financial Accounting: in communicating reality, we construct reality*; Gouldson and Bebbington, *Corporations and governance of environmental risk*).

There is a whole set of frameworks and projects within the accounting/business space dedicated to making nature visible. The TCFD, TNFD, TPT, the Capitals Coalition, the World Benchmarking Alliance, the ISSB and so on. But all of these construct nature from a business perspective. There is a substantial body of literature both academic (*Accounts of nature and the nature of accounts*, S Russell et al; Sullivan and Hannis, *Mathematics maybe, but not money*) and civil society (for example Georges Monbiot) that is sceptical of this, essentially because it does not account for nature at a systemic level nor from different perspectives.

As business professionals, if we have the agency to construct a representation of nature that has helped drive its degradation, then haven't we the agency to create a new reality, with new business models and leadership that drives a sustainable world? One in which we maintain the resources we depend on for our prosperity.

Limitless Us: A Book Review

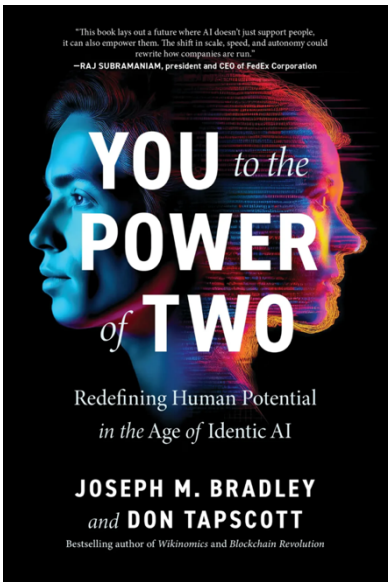
By Robert Edison Sandiford



Robert Edison Sandiford is the author of several books, among them the story cycle Fairfield*, the novel* And Sometimes They Fly*, and graphic novels with NBM Publishing. He has been shortlisted for The Frank Collymore Literary Award and is a recipient of Barbados' Governor General's Award of Excellence in Literary Arts. In 2003, he founded, with the poet Linda M. Deane, the Barbadian cultural resource ArtsEtc Inc. Sandiford has worked as a publisher, teacher and, with Warm Water Productions, producer. His stories have appeared in journals, magazines, newspapers and anthologies. Currently, he's busy with another novel, this time about fathers, sons and dementia*

A robot may not injure a human being or, through inaction, allow a human being to come to harm.
A robot must obey the orders given it by human beings except where such orders would conflict with the First Law.
A robot must protect its own existence as long as such protection does not conflict with the First or Second Law.
A robot may not harm humanity, or, by inaction, allow humanity to come to harm.

Asimov's Laws



This review of *You to the Power of Two: Redefining Human Potential in the Age of Identic AI*, by Joseph M. Bradley and Don Tapscott (BenBella Books, Hardcover, 9781637747841, 400pp, 2025) makes very clear that the presence and use of artificial intelligence generate so many prickly questions. Joseph M. Bradley's and Don Tapscott's individual dedications to their joint exploration of AI raise a few before we even reach Chapter 1. Their words are for their real-life family, with gratitude for teaching them about "becoming the fullest expression of who I am" (Bradley, a CEO) and with wishes for a "smarter, smaller world" that is "a better one" (Tapscott, a business technology author). Despite their enthusiasm for the subject, how AI will continue to be part of such human processes remains unclear.

I say "continue" because much of the technology they describe, and likely more than most of us suspect, is already here and part of our lives, hence some of our discomfort. Certainly, it is being dreamed up or imagined as you read this; and though we can't keep people from aspiring, we can pause every now and then to wonder at the wisdom of our progress.

Bradley and Tapscott open their enquiry with a scenario in which the use of AI is an everyday reality (for work, at least): “You expect the usual morning review at the office. But before you even sit, your personal AI has revised your agenda and responded to queries that you’ve *authorized*.” That last word is important. We will come back to it. Because the whole question of if we still control AI or if it is gaining ground on us will very much determine how beneficial the brave new world the authors are positing will actually be.

If AI is a tool, and possibly the best tool we’ve come up with in a long time, then we should be using it, not the other way around. I’ve yet to see a hammer build a house, if you know what I mean. Nor have I seen so many fear a hammer.

The “Why?” of our fear we all grasp. “A few years ago,” the scenario rolls on, “these decisions would have taken you weeks of meetings and emails. Now, your personal AI anticipates them, gathers and processes all the necessary inputs, and makes them just in time.” And there comes the frisson, that cold-sweat, anxiety-building realization that maybe we have become too busy to manage our own lives. Is the solution to our busyness AI, or more specifically identic AI, which Bradley and Tapscott variously define as a “personal agent,” “digital consigliere,” “companion,” or “copilot”— that is “an extension of you — thinking, deciding, shaping outcomes” on your behalf?

It can and will tell you when “[y]our stress markers are up. A gentle suggestion: ‘Take a moment to reset, perhaps a walk.’ You hadn’t noticed, but your AI had. You take a breather….” For many, this will sound like giving AI greater control over us than any tool we use should have. For Nobel Laureate Geoffrey Hinton, the godfather of A1, scenarios such as these are reminders of the control we’ve already ceded, no prompting required.

Yet the authors’ musings about the use of AI in our societies are beyond idyllic; they are convincing. What happens when “[t]echnology is no longer a tool. It is an infinite extension of you and your human potential…‘you to the power of two’”? The book covers a range of responses, from the optimistic to the alarmed, offered by familiar (and not always familiar) corporate, governmental and academic players like Eric Schmidt, Elon Musk and Mary Chayko. All involved, in their own way, acknowledge the iffyness of AI.

For one, it’s not infallible. “If you can apply your values to guide its behaviour, and it always acts according to your values and in your best interests, then what a decisive power that would be!” To be fair, many are working to ensure such applications — the authors included, with an enquiry that grows more measured in its optimism across eleven chapters. They acknowledge, “Social media has hurt our kids, balkanized society, and enabled unscrupulous people to spread false information. Our privacy is being undermined. Legions of us seem to be addicted to technology.” The last thing they want for AI is its abuse as a tool, or its evolution into an entity beyond our ability to control or comprehend.

Either outcome seems unlikely Bradley and Tapscott believe, as long as we have ownership, or sovereignty, over our digital versions of ourselves. So much of the current race to integrate AI into all facets of our lives is really about the ownership and control of data as a commodity as precious as oil or gold. At risk is our understanding and protection of personal identity. Every

plugged-in move we make, every switched-on breath we take, is recorded and converted 24/7 into some "digital mirror image" of who we are.

"[I]ncreasingly, digital technologies are capturing this behaviour as data" for the trading benefit of third parties who may be benevolent, mercantile or mercenary. In 2025, "the global AI market was valued at approximately US$200 billion"; it's forecast to "reach over US$1.8 trillion by 2032."

The race we're witnessing in this market is similar to past ones toward fresh frontiers like the so-called New World, its western coasts, the moon, and Mars. It's about money for sure, the expectation of exploitable resources, and some other instinct lodged within us that can result in as much devastation as prosperity…with the odds favouring those in the monied minority. "The deeper issue is agency itself," note the authors, our right to privacy ("the foundation of freedom"), and the safe preservation of our unique, authentic identity for our own use. Open-source blockchain technology used in a decentralized space like Web3, say Bradley and Tapscott, could help to secure this "shared digital commons." It should be understood: "an employer cannot claim ownership over an individual's digital identity."

This leads to questions about trade-offs. The authors and a number of those quoted, like Peter Diamandis, founder of the XPRIZE Foundation, talk about identic AI freeing up people for more engaging and satisfying pursuits. Often, however, those pursuits in *You to the Power of Two* are described as work related, involving endless upskilling/reskilling, and geared toward achieving ever greater productivity. Even with their proposed, people-centered Manifesto for a Digital Age, including mention of a "Universal Fabulous Income," the mindset is far less 21st-century than it should be regarding how our time might be spent in leisure, family, creative or philanthropic activity once freed from outmoded work.

Another concern about the spread of identic AI is the fortifying of personal walls, silos and echo chambers. With viable applications across the fields of healthcare, education, social services, transportation, logistics and countless others requiring people to act together, Bradley and Tapscott ask, "Are we genuinely enhancing human connection, or are we replacing it?" Worse, do we risk losing it altogether?

Amid the hurly-burly imposition of AI on us, it sometimes seems we've forgotten how to say "Stop" and "No." At present, we can't have a coffee, cuddle, take a leisurely walk or make meaningful eye contact with AI. But should we want to? What if that changed? What would that AI look like? How would it look back at us? "[W]here do we draw the line between assistance [to us] and control [of us]?" Do we know? These are crucial ethical and moral questions we should take our time to discuss.

Maybe we're looking for a saviour in a more advanced, autonomous artificial general intelligence (AGI), yet unwilling to admit it. Maybe we're playing God, and we want so badly to say we've created life that we don't care to what extent that creation may compromise our own existence. Maybe we're looking for something/someone else to blame for all that goes wrong with our social systems, climate, workplaces, homes, bars and shopping malls, our city streets and community centres.

Through reiteration and an admittedly constricted Western focus, Bradley and Tapscott suggest another option. It lies somewhere between a giddy-headed, PR-like embrace and paralyzing, apathetic fear. “[W]e are not technology determinists. Rather than forecasting a future, we need to choose it and achieve it, each of us and together. Our goal is to contribute to a promise fulfilled and a peril unrequited.” Their contribution, “Seven Principles for Identic AI,” recall speculative fiction writer Isaac Asimov’s four core laws of robotics. The authors contend these guidelines should drive any development and deployment of AI. They are reliability of the product; transparency of its processes; fairness in its outputs; accountability on the provider’s part; agency for human users; adaptability of its technology; and safety for all.

Pioneering computer scientist Douglas Englebart predicted in 1962 that AI could someday “assist humans in making better decisions” in a collaborative way. Seven decades on, are we there yet? If we are, does that mean AI is as smart as us or smarter? Like questions about the existence of the soul, the answer to the last part may be best left a mystery. To borrow that famous phrase about training, we’ve taught AI everything it knows, not everything we know. And what do we know, honestly? Encyclopedic information and immense computational speed, or even the ability to spout seemingly endless probabilities or masterfully mimic art and emotions and autonomy, do not account for the sum of all intelligence.

Unless we have finally understood all there is to know about the human machine, such statements are disingenuous. More than that, they are dangerous. They make us feel we should give up on ourselves: on what AI engineer Harper Carroll might call the unplugged, imperfect, creative, passionate, messy and unpredictable reality of “human experience.” Bradley and Tapscott argue neither for retreat nor surrender. Now is the time, they say instead, working with AI and employing all our known and unknowable humanity, to believe in our limitless potential.