



Agents of Change

How the Race to Deploy AI Agents is
Outrunning Finance Governance





Hugo Sarrazin

CEO, Avalara

“

Agentic AI has quickly moved beyond experimentation.

It's becoming part of how businesses operate, and finance leaders are under real pressure to figure out what that means for their organizations.

The opportunity ahead is significant. AI agents have the potential to help finance teams work faster, reduce manual effort, and make better decisions. But finance is different from most other functions. Every decision has compliance implications. Every transaction must stand up to scrutiny, leaving little room for error.

That's why the findings in this report matter.

What stood out to me is the tension finance leaders are managing every day. They're being asked to move quickly, prove value, and modernize critical processes while also protecting the controls and governance their businesses depend on.

The numbers tell that story. Nearly every CFO we surveyed feels personal pressure to demonstrate ROI from AI investments. At the same time, only a small percentage say governance is taking priority over speed, and many organizations are still building the capabilities needed to manage AI with confidence.

I don't believe the answer is to slow down. AI will continue to reshape finance, and the organizations that embrace it thoughtfully will be better positioned for the future. But successful adoption requires more than deploying new technology. It requires trust.

That means AI agents built on verified, domain-specific data. It means transparency into how decisions are made. It means governance and accountability designed into the system—not added after deployment. And it means choosing partners who understand the complexity of finance and compliance, not just the promise of AI.

At Avalara, that's how we approach innovation. Our goal is to help organizations automate tax and compliance with confidence, giving finance leaders technology they can trust for the work that matters most.

I hope the research in this report gives you a clear picture of where finance leaders are today, the challenges they're navigating, and the practical steps organizations can take to build AI strategies that deliver lasting value.”

Agentic AI is no longer a pilot program

Across financial functions - tax, compliance, financial closings, accounts payable, invoicing - it is becoming part of the operational core.

The Pressure Driving Adoption Is Real

Finance leaders are keen to start reaping the well-publicized cost and business optimization benefits that agentic AI offers. However, a growing number of deployments are also being driven by fear of missing out. In either case, moving fast without the right infrastructure carries significant risks in an environment with no margin for error.

Who We Surveyed

To understand what this looks like in practice, Avalara surveyed 1,500+ CFOs and senior finance leaders across the US, UK, India and Australia whose organizations have deployed, piloted or evaluated AI agents within financial processes in the past year. The majority surveyed represent midmarket companies, operating with tighter margins and less cushion than the largest enterprises.

Adoption Is Outpacing Oversight

The results suggest that what's happening is far messier than what sheer deployment numbers suggest. Three in four respondents say no one on their finance team is dedicated to understanding the operational mechanics of how these agents actually work. Audits, controls and lines of accountability are still being built in most cases, even at companies that already have a chief AI officer and formal governance function in place.

This Report Captures:

Where finance leaders actually are right now: caught between the push to deploy and the work of governing responsibly. Finance leaders aren't resisting agentic AI, and the technology isn't the problem. However, the pressure to deploy fast and show ROI is in danger of outrunning the controls meant to keep it accountable.


The pressure is real but ROI is more limited

It's clear that finance departments are being pushed to deploy agentic AI as quickly as possible.

In many instances, leaders are also being held personally responsible for both delivering ROI and keeping mistakes to a minimum, despite the compressed timelines. However, the much-needed proof that deployments are working effectively is limited in most cases.

A staggering **92%** of finance leaders said they felt personal career pressure to demonstrate their AI agent investment is delivering as intended against business targets and objectives.

Globally, half (**50%**) of those surveyed called the pressure they faced "significant." However, there are major regional discrepancies, with the figure rising as high as **57%** amongst US-based respondents and dropping to just **32%** amongst those in Australia, indicating that the pressure is not being felt evenly in every country.



Most respondents (**71%**) said the pressure they feel relates entirely or mostly to **the speed of agentic AI deployment**, even as their companies' governance frameworks struggle to catch up. Conversely, less than one in ten (**7%**) said their business is putting governance first in their deployment strategy. This paints a concerning picture of businesses rushing to move forward with deployments as quickly as possible, with potentially significant governance and control gaps.

Where the pressure sits:

Pressure is about speed

71%

Put governance first

7%

Despite the aggressive deployment pace, few respondents said the financial and operational returns have been everything they were hoping for. While **88%** of respondents reported at least some measurable ROI from their agentic AI strategy so far, half (**50%**) said these returns have been limited, while **12%** said they've seen no clear return at all, or that it's too early to tell just yet. This figure rises to **24%** amongst finance leaders based in India, suggesting businesses there are lagging behind other regions in the race to realize a true ROI on agentic AI initiatives.

Career pressure to prove ROI is nearly universal



- 50%** Called the pressure significant
- 42%** Felt moderate pressure

The Governance Gap: AI agent deployment in finance is ahead of governance maturity

The explosion in AI adoption in recent years has seen agents quickly being integrated into nearly every aspect of business operations.

“Security and governance around AI agents’ behavior and usage is top of mind right now for many organizations.”

Frank Cirone,
VP Commercial Strategy



However, while most business departments are free to utilize new technology however they see fit, finance processes are both highly complex and tightly regulated. This makes adoption of something like agentic AI – which is capable of making dozens of autonomous decisions per minute and executing complex workflows – a much more daunting prospect for finance leaders.

Over three quarters (**76%**) of the leaders Avalara spoke to said they currently lack dedicated expertise within their finance department that understands how their AI agents work.

76% Lack dedicated in-house expertise to understand how their AI agents work

Instead, they have to either lean on IT for help (**20%**) or rely on their AI vendor (**20%**). Furthermore, while a fifth of respondents are actively recruiting for experts to fill internal roles within their departments, **16%** have nobody responsible for it at all.

“You can’t give a tax team an AI tool and expect transformation, just as you can’t ask AI engineers to solve complex tax problems on their own. The strongest results come when tax professionals and AI specialists work together to build solutions that are accurate, practical, and scalable.”

James Erwin,
Head of Indirect Tax



This willingness to implement such a powerful new technology despite not having the necessary expertise in place raises numerous flags, summarized by one financial services leader in India, who stated:

“My biggest concern is that AI agents could make inaccurate or non-compliant decisions without timely detection, leading to financial, regulatory, or reputational risks for the organization.”

On top of this, 30% of respondents said they have not updated their internal control framework within the last year to account for AI agents taking or recommending actions, while almost half (46%) have AI incident response plans that are either untested or still in development.


Signs of the Governance Gap

30%

Haven't updated their internal control framework in the past year for AI-driven actions

46%

Have AI incident response plans that are untested or still in development



Even for the most motivated organizations, the speed at which the technology continues to evolve makes it extremely difficult for governance to keep up:

“As AI agents begin taking actions instead of simply generating recommendations, organizations need new governance models that clearly define what an AI agent can do, who approves its actions and how those actions are monitored. AI tools are evolving faster than corporate policies and control frameworks.”

Mark Gaeto,
CFO, Anderson Process



This is a cause for concern amongst many of those surveyed. Another respondent, who works for a professional services organization in the US that has deployed agents in last 12 months, said:

“One of our biggest challenges is making sure AI agents follow all financial and compliance rules without taking shortcuts.”

While another US-based leader working in financial services noted the challenge of agent speed vs the company’s ability to control it in a recent agent pilot project:


“The agent acts in milliseconds, but our compliance reviews take hours. We are always playing catch up.”

The Speed Mismatch

Agent action and governance response sit on completely different timescales.

AI agent acts in:

Governance responds in:



These governance gaps extend to how organizations vet the vendors supplying their agents in the first place.


01 No safeguards are formally required by a third of those surveyed before a vendor's AI agent is allowed to operate.

02 Just **31%** require documented human review or escalation thresholds, while the same proportion (**30%**) require independent third-party security or compliance

03 Only **28%** require documented audit logs showing how the agent reaches its decisions.

04 Almost a third (**29%**) say they rely primarily on the vendor's reputation and existing contract terms, with no formal technical or compliance requirements.

How much autonomy do CFOs allow AI agents, by process?



Accountability in the dark: Who's responsible when agents fail?

One of the key questions surrounding the use of AI agents in financial processes is where accountability lies if/when those agents make mistakes. The answer to this question seems unresolved for many.

In the event of a significant agent error, almost a quarter (23%) of respondents said accountability would be unclear or sit with no one (rising as high as 27% amongst respondents based in India).

Other responses suggested accountability could lie with either the individual who deployed/manages the agent (20%), the team that deployed/manages the agent (19%), or even the executive who initially approved the AI investment (16%). This will likely come as a shock to many executives, most of whom have little or no input into the day-to-day running of AI agents within their financial environments.

Interestingly, just a fifth (22%) of those surveyed said the vendor would be held primarily accountable and only where contractually specified, meaning accountability is conditional, not assumed.

Avalara also asked about the extent to which they hold the vendors financially accountable for agent errors. Responses again highlighted the inconsistencies customers face around this topic.


In 17% of cases, the contracts state that the vendor bears only minimal accountability, while in a handful of cases liability sits entirely with the customer. The most common approach (42%) is a shared liability model between both vendor and customer, while 38% of contracts state vendors are fully liable in the event of agent errors.

"The contractual standard should make vendors more accountable for AI-related issues."

Australian Finance Leader



Who is responsible when agents fail?



- 23% Unclear/no one
- 22% The vendor, where contractually specified
- 20% The individual who deploys / manages the agent
- 19% The team that deploys / manages the agent
- 16% The executive who approved the AI investment

Another critical consideration is whether organizations have the capability to accurately explain errors to auditors or regulators if something goes wrong.

More than half of respondents (53%) were **very confident** that they could do so, while 44% said they were only **"somewhat" confident** in their ability to explain an AI agent’s actions. Amongst respondents from India, one in ten said they were either not very confident or not at all confident, which should be raising major alarm bells internally.


“Organizations are increasingly exploring how to use AI to do reconciliations, validations, and much more. The risk comes from the delegation of authority that happens when those things start to be infiltrated by AI. You may be indirectly giving an agent access to data that you do not want them to have.”

Frank Cirone,
VP Commercial Strategy

 snowflake®

Confidence explaining an agent to an auditor

Each circle = 1% of finance leaders surveyed



- 53% Very confident they could explain it
- 44% Only ‘somewhat’ confident
- 10% Not very confident

India 

1 in 10 are not confident at all

How CFOs are drawing lines around the technology


The majority of CFOs surveyed clearly believe the benefits on offer from agentic AI are simply too big to ignore, despite the challenges that accompany adoption.

As a result, most are attempting to embrace it in a controlled manner. Some of the ways they are doing this include only conditioning AI agents on verified data, and prioritizing vendors based on key factors like auditability and accountability wherever possible.

For example, finance leaders put significant limits on the use of general-purpose large language models (LLMs).

Interestingly, **29%** of respondents did believe LLMs could be used to support regulated workflows, but only if every output is reviewed by a human, highlighting the inherent mistrust most respondents still have towards the technology at present.

Where CFOs draw the line on general-purpose LLMs



Together, these responses highlight the inherent mistrust most respondents still have toward the technology.

When asked what capabilities would increase confidence in deploying/expanding AI agents across financial environments, responses were similarly varied. The top five most popular answers:

23% Documented audit trails showing what the agent did and why

24% Contractual accuracy, performance or accountability commitments from the vendor

25% Evidence that outputs are tested against known compliance requirements

25% AI outputs being grounded in verified tax, compliance or financial data sources

27% AI agents having the ability to operate within the rules, permissions, and controls of existing systems of record

Ultimately, for most CFOs, the adoption of agentic AI still boils down to the issue of control.

Many say that unleashing agents across their environments will leave them at the mercy of a technology they don't fully understand, which can quickly lead to major repercussions in such a strictly regulated environment.

As a result, it's no surprise to see the top "proof" they want from vendors offering agentic solutions: auditability and explainability evidence (**18%**), measurable ROI from real deployments (**18%**), and governance and control documentation (**18%**).

The capabilities CFOs would find most impactful are also heavily skewed towards control, with audit-ready documentation for every AI action (**30%**) and real-time regulatory monitoring (**30%**) both tied at the top of respondents' wishlists.





Jayme Fishman

EVP, Chief Strategy & Product Officer,
Avalara

From Insight to Action

“The findings in this report reflect what we hear consistently from finance leaders: the pressure to adopt agentic AI is real, the desire to get it right is genuine, and the gap between the two is wider than most organizations expected.


Nearly every leader we surveyed feels personal pressure to prove these investments are delivering. Meanwhile, the controls, expertise, and accountability meant to sit underneath them are still being built. That gap isn't a reason to slow down, but it is a reason to be sure about the foundation.

Rather than reading this data as hesitation, I read it as a set of standards finance leaders are right to demand.

When they tell us the capabilities that would most increase their confidence are agents that operate within the rules and controls of their existing systems of record, outputs grounded in verified tax and compliance data, and evidence tested against known compliance requirements, they're describing the minimum bar for AI in a regulated environment.

At Avalara, we've built our agentic AI platform specifically for the precision and accountability that tax and compliance demand – trusted AI that works across the full compliance lifecycle rather than in isolated tasks – and we've learned a great deal about what separates deployments that deliver from those that don't.

The survey data points to five places where finance leaders can close the gap, starting now:



Build AI agent accountability into your organizational structure.

The finding that **76%** of finance teams have no dedicated expertise to understand how their AI agents work is the most immediately actionable gap in this report.

Organizations that govern AI agents effectively treat it as an embedded capability: clear ownership of agent oversight distributed across finance, IT, and compliance functions, explicitly documented processes, and governance that survives personnel changes. And because teams are still building up that expertise in-house today, the depth and accountability of the partner you choose matters more than ever.

02



Insist on verified compliance data, not general-purpose AI.

The single biggest determinant of whether an AI agent performs reliably in a regulated workflow is the quality of the data it reasons from.

Nearly half of the leaders surveyed said general-purpose LLMs aren't appropriate for decisions or actions in regulated workflows – and they're right. Those models are both too slow for commerce and lacking in the precision that tax law and compliance obligations require. Before deploying or expanding any agent in financial workflows, ask your vendor directly: what data is this agent grounded in, how is it maintained, and how is compliance accuracy verified? **If the answer is vague, that's your answer.**

03



Keep humans in control of the decisions that matter most.

The clearest line finance leaders draw in this report is around autonomy. They're comfortable with agents that review invoices, flag discrepancies, and recommend next steps, but far less comfortable letting an agent act unchecked on high-stakes decisions like releasing a payment.

That instinct is sound, and it should be designed into how you deploy. The goal of agentic AI in finance isn't to remove people; it's to remove friction, so human judgment is reserved for the exceptions and approvals where it genuinely matters. Define those control points deliberately, such as which actions an agent can take on its own, which require human review, and how uncertain or anomalous outputs get escalated.

04



Test your incident response plan before you need it.

Nearly half of respondents have incident response plans that are untested or still in development. Run a tabletop exercise.

Define who is notified, who decides, and how an agent is paused or rolled back if something goes wrong. Doing this before a real incident is the difference between a recoverable event and a regulatory issue.

05



Build governance in, not on.

The most common mistake in agentic AI deployment is treating governance as a layer to add after the agent is running.

In tax and compliance, that means the agent is already making decisions before the controls exist to govern them. It's no coincidence that the proof finance leaders most want from vendors all point to the same thing: the ability to know what an agent did, why it did it, and stand behind it if questioned. The platforms and vendors worth investing in are those where audit trails, humans in the loop, and permission boundaries are native to the architecture from day one.

Summary

The emergence of agentic AI is fundamentally reshaping the business landscape, rapidly changing the way organizations of every size operate. The resulting goldrush is causing many business leaders to push ahead with aggressive adoption timelines despite most early adopters reporting modest ROI at best from their initial investments.

The financial industry is certainly not immune to this, and despite the inherent challenges associated with implementing powerful new technologies in such a heavily regulated sector, it's clear that CFOs and finance leaders are under significant pressure to embrace it, just like everyone else.

In some cases, this is leading to scenarios where powerful agentic tools are being implemented before the guardrails needed to govern them can be put in place, creating high levels of risk and uncertainty amongst those tasked with utilizing them on a daily basis.

Significant question marks around accountability in the event of agent errors are also prevalent, with financial leaders wanting to see vendors take more accountability for issues, but few vendors keen to do so in a contractual capacity. This is leading to confusion about where the buck would ultimately stop in the event of a major incident taking place.

Despite these risks, the potential benefits are simply too big to ignore. As a result, many CFOs are choosing to move ahead cautiously – prioritizing vendors and strategies that offer high levels of accountability and control over those who can't or won't.



Avalara

✦ Avalara Agentic Tax and Compliance™

Start simplifying compliance today

With a platform designed to do the heavy lifting for you.

[Get started](#)