QuantSpark

A 90-day action plan for partners, advisory leads and heads of practice

The AI Inflection Point in Accounting, Advisory and Audit

Choose 3 best stats from page 3

I suggest 41%, 5hrs, $37.6bn

XXX	**XXX**	**XXX**
xxxxxxx	**xxxxxxx**	**xxxxxxx**
xxxxxxxxx4	xxxxxxxxx4	xxxxxxxxx4

Contents

Capitalise all titles

The state of play, and your next move

In one year...

AI adoption inside accounting firms went from 9% to 41%

Wolters Kluwer, 2025

The global AI in accounting market is forecast to grow from $6.68bn to $37.6bn by 2030

Mordor Intelligence

The gap between using AI and having a clear, pragmatic, actionable strategy for AI is now the single largest source of variance in firm outcomes.

46% of accountants now use AI every day, 95% of firms have deployed some form of automation

64% plan to increase AI investment this year

Intuit QuickBooks, 2025

Only approximately 25% of firms have a defined, pragmatic AI strategy and roadmap. Those that do are three to four times more likely to capture revenue and efficiency gains

Thomson Reuters, 2025

This paper is a 90-day action plan for partners and heads of practice. It maps where your service lines are exposed, which AI fits which job, what is working in deployment, and what to do this quarter. If you take only one thing from it, take this. The action is at the service line, not the firm.

With three of these stats moved to the cover, lower this paragraph so that “The gap between using ai…” can be full width and above stats have breathing room

41%

of accounting firms now use AI, up from 9% in 2024

Wolters Kluwer, 2025

46%

of accountants use AI daily

Intuit QuickBooks, 2025

5hrs

saved per professional per week, worth around $19,000 a year

Thomson Reuters, 2025

3-4x

higher revenue and efficiency gains for firms with a defined AI strategy

93%

of firms now offer advisory services, up from 83% the prior year

Wolters Kluwer, 2025

$37.6bn

projected AI in accounting market size by 2030, from $6.68bn

Mordor Intelligence

Section I

Where is your firm exposed?

Firm level AI scores hide more than they reveal. Inside one accountancy firm, audit and SME bookkeeping face opposite exposures. The action is at the service line, not the firm.

Audit is largely protected. Regulatory underpinning, signing partner liability and audit committee preference for human accountability keep buyers paying for the human signature. SME bookkeeping is being directly substituted. Vertical SaaS and AI agents are already winning low-end work. Tax compliance sits in between. Advisory grows on the back of capacity released elsewhere. A single firm-wide AI score will tell you none of this.

Lower exposure

Audit

Lower/align this title with the one on the right

Why is this more protected

Regulatory underpinning. Signing partner liability. Trust and reputation drive buying. Audit committees prefer human accountability.

Where AI bites

Substantive testing and review can be augmented through faster sampling and anomaly detection. Augmentation, not substitution.

Strategic move

Use AI to expand capacity and take share. Lower cost to deliver while pricing holds.

Posture

Defend and augment

Higher exposure

Subject Matter Expert book keeping

Why is this more exposed

Repeatable workflows. Low regulatory friction. Price-sensitive customers. Vertical SaaS already substituting at the low end.

What is being substituted

Categorisation, reconciliation, statement preparation. Increasingly: routine advice and basic compliance prompts. Direct substitution.

Strategic move

Migrate clients up to advisory. Or restructure the cost base and compete as a utility.

Posture

Migrate or restructure

ACTION

Score every service line on its own merits

Rate each service line on four dimensions: substitutability, regulatory friction, client willingness to pay for human accountability, and the strategic move available. Manage them as a portfolio. A firm-level AI dashboard with one number is the wrong unit of analysis.

Section II

Service line action map

Split into 3 ‘sections’ or tables to reduce reading fatigue. black lines indicate where each ‘break’ should be.

Service line	Substitution risk	Do this quarter	First action
SME bookkeeping and close	Very high	Migrate clients up to advisory	Map top 20 SME clients; book advisory-fit conversations within 30 days
Tax compliance (individual, high volume)	High	Restructure preparer model	Trial an agentic prep platform; redefine staff role as reviewer
Payroll and compliance services	High (at low end)	Restructure or partner	Quote a build, buy or partner decision to the partnership
Tax compliance (corporate, complex)	Medium	Augment senior judgement	Deploy a domain-tuned LLM for research; mandate cited-source review
Internal audit and risk advisory	Low to medium	Augment testing	Deploy continuous controls monitoring for one client
Statutory audit and assurance	Low	Augment to take share	Pilot full-population anomaly detection on two engagements
Advisory and CAS	Low (grow)	Scale advisory propositions	Productise three advisory offers built on AI-generated insight
Forensic and fraud	Low	Augment investigator capacity	Pilot full-population anomaly detection on the next engagement
AI assurance (new line)	N/A	Build if positioned	Scope an EU AI Act readiness service for one regulated client

Section III

Which AI for which job

Not every problem has the same AI solution.

Picking the wrong model architecture is the single largest source of disappointment in firm AI programmes.

AI Type

Rules-based automation and RPA

Best for

Deterministic, high-volume tasks where the same input always produces the same output

Where to use it

Payroll, standard filings, data movement between systems

What it must not do

Interpret ambiguous data or handle exceptions

AI Type

Classical Machine Learning and anomaly detection

Best for

Pattern recognition on structured data with known labels

Where to use it

Audit anomaly detection, fraud screening, transaction categorisation, risk scoring

What it must not do

Generate prose or reason about novel situations

AI Type

Domain-tuned LLMs

Best for

Synthesis grounded in cited professional authority

Where to use it

Tax research, technical standard interpretation, complex advisory synthesis

What it must not do

Replace professional judgement on contested or novel matters

AI Type

Document AI (OCR plus extraction)

Best for

Turning unstructured documents into structured data

Where to use it

Receipts, invoices, K-1s, 1099s, contracts, bank statements, prior-year returns

What it must not do

Be trusted without sample-based human verification on high-risk extractions

AI Type

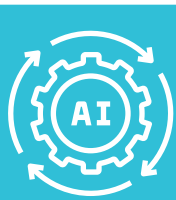

Agentic AI and multi-agent systems

Best for

Orchestrating multi-step workflows across systems

Where to use it

End-to-end engagement workflows (intake, preparation, review), month-end close

What it must not do

Run without logged, auditable handoffs and human review at named control points

ACTION

Pick three use cases that span three model types

If you deploy only one new capability this quarter, choose a solution-specific AI.

Pair it with one documented AI use case and one anomaly-detection pilot.

Three pilots, three model types, three controlled risks.

Section IV

What is working. What is not.

Eighteen months of deployment evidence now lets us separate the use cases that consistently deliver from the ones that consistently disappoint.

Deploy now (high confidence)

Document intake and extraction.

Manager-review quality output from receipts, K-1s, contracts, bank statements. Vendors: Veryfi, Dext, Filed, Canopy, Accrual.

Advisory scenario modelling.

Cash flow, pricing and what-if scenario modelling on client-specific data.

Tax research with cited authority.

Three to five times faster on routine technical questions in published case studies. Mandatory: verification against cited source.

Anomaly detection in audit and forensic work.

Full-population analytics replacing sampling. MindBridge and Big Four internal platforms show consistent gains in risk identification.

Written client communication and meeting workflow.

The leading AI use case among accountants at 64% (Karbon, 2025), ahead of task automation (41%) and research (40%).

Approach with caution

End-to-end agentic tax preparation.

This is promising at scale (Accrual's $75m raise in early 2026, adoption by H&R Block, Armanino, Creative Planning). It requires mature review controls and clarity on who signs the return.

Avoid

Final regulated output directly from general-purpose AI.

The top recurring source of client-quality failures.

Autonomous agent action across nuanced areas without human checkpoints.

Regulatory and reputational risk regardless of model maturity.

AI for contested standards interpretation.

Revenue recognition, lease classification, going concern. The model can structure the analysis. It cannot make the call.

Section V

Risks you must control before you scale

The firms that get this wrong will face client-quality, regulatory and reputational consequences that outlast the technology cycle.

The regulator is already here

In June 2025, the UK Financial Reporting Council issued its first formal guidance on AI in audit, alongside a review of the Big Four, BDO and Forvis Mazars. With one exception, none of the firms had defined KPIs to monitor the AI's contribution to engagement quality.

The PCAOB, AICPA and IESBA have signalled similar concerns.

ACTION Add more variation in colours to this section for skimming

Five governance controls you need

Rate each service line on four dimensions: substitutability, regulatory friction, client willingness to pay for human accountability, and the strategic move available. Manage them as a portfolio. A firm-level AI dashboard with one number is the wrong unit of analysis.

1. Named AI policy with engagement-letter language disclosing AI use to clients.
2. KPIs that measure AI's contribution to engagement quality (the FRC standard).
3. Data classification scheme stating which data classes go to which model tier.
4. Mandatory human verification against cited source for any regulated output.
5. Logged, auditable handoffs at every agent control point.

Section VI

Your 90-day action plan

The firms capturing the strongest returns share a recognisable operating model.

This is the sequence we use to build it.

Diagnose

Weeks 1 to 2

What to do

Map each service line to one of four postures (Defend, Augment, Migrate, Restructure). Inventory current AI use. Most firms underestimate by 40 to 60%. Pick three to five priority use cases.

Output

Service line exposure map. Prioritised use case shortlist.

Govern

Weeks 3 to 6

What to do

Stand up AI policy, engagement-letter language, data classification, vendor risk framework, KPI definition. Sign-off by the partnership.

Output

Governance framework ready for regulator inspection.

Pilot

Weeks 7 to 12

What to do

Deploy each priority use case as a controlled pilot inside a named team. Measure against the KPIs from Phase 1. Plan the rollout decision before the pilot ends.

Output

Validated, ready-to-scale use cases with documented controls.

Scale

Months 4 to 6

What to do

Roll validated configurations across the service line. Train. Monitor KPIs monthly.

Output

Production deployment at service-line scale.

Reinvent

Months 7 onwards

What to do

Redeploy capacity released by AI into relationship building, advisory or expanded audit capacity.

Output

Reshaped service portfolio aligned to chosen postures.

The economics, plainly stated

Time savings per professional are conservatively five hours a week. That equates on average, in mid-sized companies, to around $19,000 a year (Thomson Reuters). Payback for well-scoped deployment in a mid-sized firm is typically inside twelve months, if governance and integration costs are budgeted alongside the licence fee. Firms that disappoint themselves budget for the licence and absorb the rest from operating margin.

Section VII

Where to start this week

If you do nothing else in the next five working days, do these seven things.

01
Pick one service line.

Not the firm. One service line where you have decision authority and engaged practitioners.

02
Run a 30-minute substitution test.

Ask which parts of this work are repeatable, regulated, judgement-led, or relationship-led. The repeatable parts are your AI candidates.

03
Audit shadow use.

Survey your team on which personal AI tools they already use for client work. You will discover more than you expect.

04
Pick one high-confidence use case from Section 4.

Document intake, anomaly detection, or tax research with cited authority. Pilot in one team

05
Draft engagement-letter language.

Disclose AI use to clients. This is becoming non-optional under FRC guidance.

06
Name three KPIs.

Without measures of AI's contribution to engagement quality, you cannot answer the regulator and you cannot defend the investment internally.

07
Book the build, buy or partner decision.

Big Four are building. Everyone else should buy at the application layer, own data and governance, and partner for deep specialism.

ACTION Add more variation in colours to this section for skimming

Self-assessment: can you answer these five questions today?

1. Which of our service lines are in Defend, Augment, Migrate or Restructure?
2. What is our defined AI strategy, written down?
3. Who is accountable for AI quality on each engagement?
4. Which KPIs prove AI is improving engagement quality (not just speed)?
5. Where is capacity being redeployed as AI takes routine work?

If you cannot answer three or more of these, you are in the 75% of firms without a strategy. Three to four times less likely to capture the upside.

Thomson Reuters, 2025

Conclusion

The shift from 9% to 41% adoption in a single year tells the story.

AI in accounting has crossed from advantage option to operating reality. What separates the firms capturing the upside from those still reacting is not access to technology, which is rapidly commoditising at the application layer. It is the discipline to score each service line on its own exposure, choose the right model for each job, govern AI tightly enough to answer a regulator who has already arrived, and redeploy released capacity into advisory rather than absorb it as cost saving. Three quarters of the profession is not yet doing it. The firms doing this work are three to four times more likely to grow revenue and margin.

Every partner reading this paper should start that work this week. Pick one service line. Pilot one high-confidence use case. Name three KPIs. The firms that move now will define the operating models the rest of the profession adopts in 2027 and beyond. The firms that wait will inherit them.

Take the next step

If the action plan in this paper maps to questions your firm is wrestling with, we will share our service line exposure diagnostic template and walk through how it applies to one of your propositions in a 45-minute working call.

No obligation. We will tell you candidly if we are not the right partner.

For more information

Visit www.quantspark.ai

Email contact@quantspark.com

Sources

Primary research

Wolters Kluwer, 2025 Future Ready Accountant report. Source for AI adoption growth from 9% to 41%, advisory penetration figures, investment intent.

Intuit QuickBooks, 2025 Accountant Technology Survey (n=700). Source for daily AI use, automation adoption, productivity and cognitive load findings.

Thomson Reuters, 2025 Future of Professionals report (n=2,275). Source for AI strategy adoption divide, time savings, the $32bn US opportunity figure.

Karbon, 2025 AI in Accounting Industry Report. Source for AI usage by task type.

CPA.com, 2025 AI in Accounting Report (with AICPA).

Mordor Intelligence, AI in Accounting Market 2025 to 2030 ($6.68bn to $37.6bn forecast).

Regulatory and standard setter sources

Financial Reporting Council (FRC), AI in audit guidance and 2025 firm review (Big Four, BDO, Forvis Mazars).

Public Company Accounting Oversight Board (PCAOB), 2025 investor advisory group discussion of AI in audit.

AICPA and Journal of Accountancy, AI risks and CPA-as-evaluator coverage (2025 to 2026).

IESBA Code of Ethics updates on AI hallucination.

Platforms and incidents

PwC Agent OS (2025); KPMG Workbench (June 2025); Deloitte Zora AI; EY Agentic Platform.

Deloitte and Australian government incident (October 2025).

Accrual launch and $75m Series funding (February 2026), with H&R Block, Armanino, Creative Planning as early adopters.

Vendor landscape: MindBridge, Vic.ai, Botkeeper, AppZen, BlackLine, FloQast, Dext, Veryfi, Xero, Intuit, Sage.

QuantSpark

Authors

Adam Hadley CBE
CEO & Founder
adam.hadley@quantspark.com

Wanda Haddock
Head of AI Transformation
wanda.haddock@quantspark.com

Contact

www.quantspark.ai

contact@quantspark.com

230 Blackfriars Road
London
SE1 8NW

Update job number

MK406