


The destruction of the billable hour

AI's impact on accounting firm pricing in 2027

Intelligence is becoming a direct cost of every engagement. What that means for firm economics, and how firms will price, deploy and run in the year ahead



CONTENTS

00	Labor, capital and now intelligence	3
01	AI's success breaks the pricing model	5
02	Deployment is the hard part	7
03	Five ways to price when agents do the work	9
04	Deploying a new pricing model	15
05	Where this goes, and why Instead	18
06	A note from Andrew	21

INTRODUCTION

For most of its history, an accounting firm has run on two inputs: people and capital. A third input has arrived, and it changes how every engagement should be priced.

That input is intelligence, delivered by AI agents. It is taking on a growing share of the work people used to do in accounting, tax, audit and consulting. On the cost side it is a new direct cost that varies by client, tool and model. On the delivery side it means work no longer happens one client at a time.

This paper explains why hours times rate stops working, why deployment is the real bottleneck, five pricing models firms are adopting, what it takes to roll a new model out across thousands of engagement letters, and where we think the firm goes next.

For managing partners and the leaders who own pricing and operations

Covers accounting, tax, audit and consulting

Horizon the 2027 engagement cycle


Labor, capital and now intelligence

Every service business turns labor and capital into revenue. Accounting firms are no different, and for decades their economics have barely changed shape.

When a firm serves a client it collects revenue and spends it in three places: cost of sales, operating expenses and profit. Every firm is different, but a common structure has been roughly one third each: about 33% cost of sales, 33% operating expenses and 33% partner profit.

Cost of sales has meant people. The hours of staff, seniors, managers and partners were the cost of the work and, through hourly rates, the basis of the price.

That is now changing. Next to capital and labor sits a new input, intelligence. It is already consuming large portions of the labor inside an engagement, and it is on track to outpace labor.



Where the revenue dollar goes. Historical shares are the common one third rule of thumb. The split inside cost of sales is illustrative: it varies by firm, service line and client.

00

What an engagement costs now


Look at a single engagement today and its cost has three parts.

Human cost

Other capital: technology and other direct costs

Intelligence cost

The first two are familiar. The third is new, and it does not behave like the others. Intelligence cost varies with the complexity of the client, the tools the firm uses and the models behind those tools. Two clients with the same fee can carry very different intelligence costs, and the same client can cost less next year as models improve.



Engagement cost by client complexity, illustrative. Intelligence grows fastest with complexity, while human cost grows more slowly once agents do the preparation.

The hour is still the unit

The majority of accounting firms in the United States still price on an hourly basis. Many of those that don't still track hours to understand how much time goes into the work. That is true across accounting, tax, audit and consulting.

78%

of firms still use hourly billing, the dominant pricing model. Cost, price and time have always moved together. Intelligence pulls them apart

Thomson Reuters Institute, October 2026


AI's success breaks the pricing model

Firms have tried to move away from hourly billing for decades. AI does more than add pressure. It scrambles the operating model behind engagement letters, time tracking, billing and gross margin.

From four-hour blocks to batches

A professional has typically worked one client at a time: a four-hour block on one client in the morning, another client in the afternoon, with the time recorded against each. Some people will still work that way, orchestrating several agents on one client all day.

Many will not. In a 15 minute block, one person can kick off agentic workflows across six clients, then watch them, iterate, give feedback and keep orchestrating through the day. The same pattern scales to 60 or 150 clients worked in batches. In that block of time the person is not working on a client. They are working on many clients at once, and a time entry against any one of them barely makes sense.




One professional's day, illustrative. Before, time maps to one client per block. Now short bursts of orchestration start and steer agent work across many clients, and the agents keep working in between.

A new direct cost, and fewer hands on each return

While that person orchestrates, the work consumes intelligence: a direct cost that is completely variable with the models and products in use and with the complexity of the work the agent performs.

The next implication is the simplest: the same work needs fewer people. A tax return has traditionally moved down a chain. An offshore team prepares it, US staff review it, a senior reviews it and a partner reviews it. AI can already do all of the preparation and get the return ready for review.



From reviewing records to reviewing the system

Over the next year, in many parts of the firm, AI will get work to the point where it does not need to be reviewed record by record. Take a firm with 20,000 partnership returns on Form 1065. If reviewers go through 250 in a row and make no edits to the return and no edits to the outstanding requests to the client, the firm may stop reviewing every return.

That becomes possible when the agent has the context a professional would have: every call and email, all of the documents, and the firm's standard for how the work should be prepared. The human role shifts from doing and reviewing each return to running a control framework. It checks that client information is processed under the laws and regulations, that the AI performs proper supervision and review, and that people supervise the AI system rather than each record.

Firms already work this way elsewhere. No professional reviews every transaction on a bank feed. They check the control totals.

Hours times rate stops computing. Depending on how far a firm scales these tools, it may not compute for any client.

Deployment is the hard part

The technology to do the work exists. Deploying it across a firm is a different problem.

Most accounting firms have used Microsoft Copilot, ChatGPT or Claude. Most have not deployed AI systems that complete the full set of tasks that make up a client engagement. Those systems exist today for many areas of work, and there is a strong argument that AI can already do anything a professional accountant does on a computer.

Deployment will be uneven. In a typical firm, 90% or more of clients may receive tax services, 10 to 20% accounting services and 5 to 10% audit, while audit-only firms and consulting practices vary widely. Next year many firms will scale agents to only part of that client base. Some people will work almost exactly as before, perhaps with an AI assistant. Others will hand whole engagements to agents. A few firms will move completely. Most will run both ways of working at once.

Systems of record built for a different era

A major bottleneck is the core software firms run on. Many systems of record in tax, audit and accounting were designed before the internet, or as Windows desktop applications, long before agents existed.

Area	Systems firms commonly run
Tax	Wolters Kluwer and Thomson Reuters suites, such as CCH Axcess and GoSystem
Audit	Audit products from the same companies, and Caseware
Accounting	QuickBooks, NetSuite, Sage Intacct, Xero

These products differ in how well agents can work with them. Cloud products with strong APIs tend to be the most accessible. Desktop software is often hard or impossible for an agent to operate. That will slow how quickly firms can move clients and teams to agent-run work.

The choice every firm will make

Firms will have to decide how their agents meet their systems.

One path lets agents roam through existing products, even where those products make it hard to change records, track versions or control what is done. The other moves the work to agentic systems that control the experience, version everything and track every change end to end.

Agents operate existing tools

Works inside the systems the firm has today

Changes are hard to version or trace

Limited control over what the agent did, and why

Parts of the workflow stay manual

Move to agentic systems

Every change is versioned and tracked end to end

People can observe, train and control the agent

Supervision and review are built into the work

The whole workflow can run, start to finish

The second path is how a firm knows what an agent is doing, can observe and correct it, and gets the full benefit of agents completing the entire workflow for the services it offers.



Every change, versioned. The current version on top, each earlier version kept, so any agent edit can be traced and reversed.

Five ways to price when agents do the work

If these tools can do the work, and the real constraint is deploying them on modern systems, the question becomes how to price.

Here are five patterns we see firms on Instead using, or planning to use, over the next year. None is the single right answer, and they can be combined.

Model 1 of 5

A multiple of intelligence cost



Price each engagement as a multiple of what the agent work costs to deliver it.

How it works	Take the intelligence cost per client and apply a firm multiple. If an agent costs \$100 to prepare a small tax return, a 20x multiple prices it at \$2,000.
Human labor	People still review and advise. That may bring the total cost of the return to \$300 to \$500, and the margin stays strong.
Fits best	Smaller, well-defined engagements where the agent does most of the work
Watch for	Model costs change. Revisit the multiple as they move, so the price follows the value, not the vendor

Five ways to price

Cost plus a gross margin target




Add up what the engagement really costs, then price to the margin the firm wants.

How it works	For each client, add the human labor cost and the agent cost, then set the price so the engagement meets a target gross margin.
Why firms like it	Margin is visible client by client, and intelligence becomes a planned line of cost rather than a surprise.
Fits best	Firms that already manage to gross margin and want the same discipline with agents
Watch for	Both costs must be measured per client. That needs systems that record agent cost and human time by engagement

Five ways to price

Complexity of inputs and outputs




Price from what goes into the engagement and what comes out of it.

Inputs	How many entities the client has, and how many documents come in
Outputs	How many forms go out, and how many jurisdictions they cover
Scenarios	How many complex scenarios the work has to handle
Where it applies	Audit, accounting and tax engagements alike
What changed	This analysis used to take too long to run across every client. Today it can run with one click on one client or on 20,000

Five ways to price

Value-based pricing



Define the value of the service, and price to that value.

How it works	The firm sets the price against the value the client receives, whether measured against the client's alternative or the nature of the service.
In tax	A simple measure is the tax savings the firm delivers for the client.
Elsewhere	Value is harder to define in other services, and sometimes it is the wrong basis. Pricing an audit on its opinion would not be appropriate.
Fits best	Services where the firm can define the value and the client agrees with the definition

Five ways to price

Complexity add-ons



Start from any base price and add for the factors that make a client harder to serve.

Timing	Urgency and deadlines
Accounting	Number of entities, foreign currencies, debt, public company reporting, investor reporting
Tax	R&D credits, fixed assets, complex partner allocations, complex K-1s, multi-state and international work

Combine them

These models do not need to stand alone. A firm can set a base from complexity of inputs and outputs, check it against a gross margin target, and layer add-ons on top. What matters is that the firm defines one policy, and every price can be traced back to it.

The five models side by side


Each model answers a different question. Most firms will use more than one.

Model	Priced from	Data the firm needs	Strongest when
Intelligence multiple	Agent cost per client times a firm multiple	Intelligence cost per engagement	Small, well-defined engagements
Cost plus margin	Human plus agent cost, priced to a gross margin target	Human time and agent cost per client	Firms run to gross margin
Inputs and outputs	Entities, documents, forms, jurisdictions, complex scenarios	Client records the agent can read	Large client bases priced consistently
Value-based	The value the client receives	An agreed measure of value, such as tax savings	Planning and advisory work
Complexity add-ons	A base price plus defined add-ons	A list of triggers per service line	On top of any other model

What every model has in common

None of them is computed from hours. Each needs data about the client that used to be too expensive to gather for every engagement, and each can now be calculated by an agent for every client the firm serves. That moves the hard part from calculating prices to deploying them.

Model 1, worked



A simple example on a small tax return. An agent cost of \$100, a total cost of \$300 to \$500 once people review and advise, and a price of \$2,000 at a 20x multiple.

Deploying a new pricing model

Choosing a pricing model is a strategic decision. Rolling it out is an operations problem, and historically a messy one.

Consider a firm with 70,000 engagement letters that all need to go out in a short window, plus new letters every month for new clients, existing clients and new engagements. Visibility and control at that scale have always been hard.

Agents change what is possible. They can follow the firm's pricing policy and render each engagement letter dynamically from it. They can work with the full context of each client: historical and current records, email, and call transcripts. The result can be a completely dynamic pricing system built on the firm's own vision.

Instead, Client Grid lets a firm prompt hundreds or thousands of clients at once and run these steps in one agentic workflow: analyze each client's situation, map it to the pricing methodology and generate the engagement letter.

One workflow, every client

- 01 **Pricing policy.** The firm defines its models, multiples, margins and add-ons in writing
- 02 **Client context.** Records, documents, email and call transcripts for each client
- 03 **Price and terms.** The agent maps each client to the policy and calculates the price
- 04 **Engagement letter.** Rendered for that client from the firm's template
- 05 **Controls and review.** People check the prices the policy says to check
- 06 **Signature.** Sent, signed, and the engagement begins

One request, every client

The same pattern already runs in Instead today. The firm agent takes one request from the firm and dispatches it to each client's workspace, where the work is done with that client's own documents and context, then reports back per client.

Workspace tasks1 dispatched

Tax return preparation1

Birchline Studio Inc. (Training)1120-S

This is the Birchline Studio Inc. (Training) S corporation client. I need to know what open items remain before we can file the 1120S return for tax year 2025. Please read the computed 1120S return and its diagnostics/validation errors, check any workpapers and...

Show more

Done

Workspace tasks1 client

Birchline Studio Inc. Training

Done

The Birchline task is dispatched and running. I'll wait for it to complete before dispatching Harborview, per the sequential rule.

Birchline Studio Inc. (Training)

Task completed

Go to client

Birchline's review is complete. Now let me check Harborview.

Reasoning

Workspace tasks1 dispatched

Tax return preparation1

Harborview Capital Partners LP (Training)1065

This is the Harborview Capital Partners LP (Training) partnership client. I need to know what open items remain before we can file the 1065 return for tax year 2025. Please read the computed 1065 return (if one exists) and its diagnostics/validation errors, chec...

Show more

Done

Workspace tasks1 client

Harborview Capital Partners LP Training

Done

Harborview's review is running. I'll wait for it to complete before summarizing both.

Harborview Capital Partners LP (Training)

Task completed

Go to client

Both reviews are complete. Here are the open items per client.

The firm agent in Instead. One firm request dispatched as tax return preparation tasks to two client workspaces, each completed and reported back. Sandbox, training clients.

Controls, terms and the first year

When the price was hours times rate, firms were comfortable sending out 70,000 letters. A dynamic price needs a different kind of review.

Firms need a method of control that confirms the agent calculated each price correctly: which prices a person reviews, which are sampled, and which triggers send a letter back for a closer look.

They also need terms that allow for a mistake. With that many dynamic instances in the first year, an agent can get a price wrong, perhaps where human review was limited or the reviewer did not understand what they were checking. If the letter goes out and the client signs, the firm should not be held to a price the agent got wrong. Engagement letters should be updated to allow that flexibility.

Done well, this gives firms more freedom than they have ever had to define a vision for pricing and run it, with prices tailored to each client under the policies of this new era.

The hard part moves from calculating prices to deploying them with control.

BEFORE YOU ROLL OUT

Are all of your clients in one system an agent can work across?

Is the pricing policy defined, and who owns it?

Which system sends letters and collects signatures?

What controls confirm the agent calculated each price correctly?

Do your letter terms let you correct a price in year one?

Where this goes, and why Instead

We believe agents will run the entire professional services organization, with people supervising throughout.



Agents will create engagement letters, set the price and terms, and send them for signature. Once a letter is signed, they will raise the requests to the client that the engagement needs and define the tasks inside the firm for people and for agents. As clients submit documents, agents will process them and route them into workflows the firm defines.

People will be placed throughout, supervising individual work areas and the systems acting across the firm. As engagements finish, people complete supervision and review, agents deliver results to the client in ways far beyond what we are used to, and agents send bills, follow up on collections and tell people when they need to step in to close the loop.

The firm as its own intelligence


When people hear this, they ask what the firm is worth. Using models off the shelf is one thing. Using Instead's models is another. Layer in the firm's own intelligence and the result is different. A firm is a kind of mini superintelligence: professionals with knowledge and skills, and documents, context and examples no one else has.

These systems will spread to many firms, but the experiences will not be the same. Not every firm has the same knowledge of complex tax situations or the same experience in budgeting and forecasting. Each firm's agents will reflect its own preferences, style and knowledge. They also keep that knowledge in the firm: when an employee leaves, the intelligence they built can keep serving clients.

That makes it important to choose who you partner with for this journey, and what vision that partner has for serving clients.

Going deeper with clients

We believe firms will use this technology to go deeper into their clients' businesses. Instead of a monthly P&L, balance sheet and spreadsheet, an accounting engagement might deliver a live forecast and budget, updated in real time, with assumptions drawn from every conversation with the client and connected to their systems of record to track cost, spend and activity across their teams. The firm can raise issues within the month, through agents or people, that the client has not yet seen. The same is true in every service line.



The cost of a journal entry, over time. Each step took more of the work to effectively zero.

Lose the bundle of tasks, keep the career

Many of the services firms sell today will go to zero, the way journal entries did.

Fifty or sixty years ago, accountants recorded journal entries in pencil on green ledger sheets. Then came spreadsheets, general ledger systems, bank feeds and automatic reconciliation, and the cost of that work went to zero. The cost of a journal entry today is effectively nothing. Many of the services professionals provide today will become free in the same way.

Losing the bundle of tasks that makes up a job today does not mean losing a career. It means finding new ways to provide value, going deeper for clients and solving new problems.

That is what we do at Instead. We help accounting firms move to the frontier of this technology and discover what they are capable of, often things they did not think were possible just a few months before they dig in.

WHERE TO START


- 01 Write down the pricing policy you want for 2027, even if it is a draft
 - 02 Check that every client is in a system an agent can work across
 - 03 Run the policy on a sample of clients and compare it with last year's fees
 - 04 Update engagement letter terms so prices can be corrected in year one
 - 05 Decide which prices people review, and which they sample
-

A note from Andrew

Thank you for reading.

Pricing is where AI's impact on a firm becomes concrete: in every engagement letter, every invoice and every conversation with a client about value. The firms that define their pricing policy now, and deploy it with the right controls, will set the terms for how the profession works in the years ahead.

We built Instead so that firms can make this move with confidence: agents that do the work across every client, with people supervising and every change tracked. If you want to work through what this means for your firm, we would be glad to help.



A stylized, handwritten signature of Andrew Argue, consisting of two large, overlapping 'A' characters.

Andrew Argue
Founder and CEO, Instead

NOTES

Hourly billing share: Thomson Reuters Institute, "If AI can do the work, then what are you actually charging for?", October 2, 2026.

Revenue split, engagement cost and timeline charts are illustrative, not measured. Every firm's numbers differ.

The \$100, \$300 to \$500 and \$2,000 figures are a simple example on a small tax return, not a recommended price.

Pricing models describe patterns firms on Instead are using or planning. They are not legal, tax or accounting advice. Review engagement letter terms with counsel.

Product image: Instead sandbox, training clients, read-only crop.
