



A DRIFTWOOD CREW DISCUSSION PAPER

When Execution Gets Cheaper, *What Are Clients Paying For?*

*AI, judgement and the changing economics of professional services.
A Driftwood Crew discussion paper by [Alex Rankin](#).*

15 October 2026

www.driftwoodcrewconsulting.com



AI can now produce a first version of many kinds of professional work, whether code, analysis or a draft report, in a fraction of the time it used to take. Most professional services leaders already accept that judgement still matters when this happens. The harder question is commercial. If clients pay for fewer hours of execution, how does a firm fund the judgement it still needs?

I put versions of that question to three people running technology businesses that are further down this path than most professional services firms. Cyrus Kiani is founder and CEO of TouchZen Media, which designs and builds mobile and web apps. Ali Syed is founder of NoCodeHero, which builds automation and custom software for startups and small businesses in Texas. Jonathon Shaevitz is Founder and Head of Product and Strategy at GrowthCode, which provides first-party signal and identity infrastructure for digital advertising.

All three are clear that expertise is still necessary. The open question for professional services firms is whether clients will pay for it, and whether firms can capture its value. Cyrus's experience suggests that won't happen automatically: fewer delivery hours and smaller project values, while discovery, architecture and checking still need senior people. A firm can end up with less billable execution and just as much expert work to fund.

This paper argues that when the first version gets cheap, three things have to change: how firms separate discovery from delivery, how they decide what must be checked and by whom, and how they price, staff and train. It ends with a test you can run on one of your own deliverables.



A firm can end up with less billable execution and just as much expert work to fund.

Why start with technology firms?

Software is an early warning for the rest of professional services. Developers adopted AI tools early and widely: Google's [2025 DORA report](#) found 90% of nearly 5,000 technology professionals use AI at work. What software agencies are dealing with now is likely to reach other advisory and technical firms soon.

Driftwood's focus is management and professional services consultancies, and the same pressures are easy to picture there. These are illustrations, not findings from the interviews:



- If AI produces a first draft of a strategy review or board paper in a day, what happens to the fee, and to the budget for senior review?
- If a tool can sort hundreds of stakeholder submissions and trace each claim to whether the decision-maker has the legal authority to consider it, the analysis becomes auditable. Will staff use the time saved to work out how the project should change in response, or to find faster reasons to set difficult concerns aside?
- If junior consultants do less of the first research and drafting, how will they learn to recognise weak evidence, conflicting sources and conclusions that don't follow?

The first and third questions point to the same risk: an output can look complete well before the work needed to justify it is complete. The second shows what a well-designed tool makes possible. Once the evidence is compiled and traceable, the judgement that matters is deciding what to do about it.

01

1. Faster production doesn't resolve uncertainty

One of Ali's recent clients wanted a tool that would generate newsletters automatically. He built a working prototype quickly. The client then used it for two months before they knew what they actually needed, and development started again from scratch. "That discovery journey didn't get any faster," he said. AI had shortened the build. It hadn't shortened the time it took the client to understand their own problem.

Cyrus sees the same pattern from product strategy. Because AI makes features cheap to add, founders are tempted to aim straight for something the size of Instagram or Uber. His advice runs the other way. Find the one feature that creates the value, which he calls the X factor. Confirm people want it, work out how you'll reach them, and build the smallest product that tests it. Every extra feature adds code, data, possible vulnerabilities and testing. His rule on privacy is equally simple: if the product doesn't need the data, don't collect it.

Ali's clearest example of this came from a lead-generation client. Its customers wanted leads to book straight into their calendars, but wouldn't give API access to their systems. The client proposed an AI agent that would visit each calendar and pull the available times. The page needed to show availability within about ten seconds, and an agent might take 15 minutes, cost far more to run and fail whenever a website changed. Ali built what he called a "super dumb" scraper with a



fixed set of instructions instead. "People solved this 15 years ago," he said, "and we can use the same thing here." The saving came from understanding the requirement, not from choosing a more sophisticated tool.

There is a trap here for clients too. They see what a chat assistant can do and expect their own product to behave the same way, without realising how many tools and connections sit behind it. As Ali tells them: "We're not looking to build Claude. We're looking to build this solution."

“*That discovery journey didn't get any faster.*”

What this means for a firm: treat discovery as its own piece of work, with its own budget and deliverable. A fast prototype is a good way to learn what a client needs. It is not proof that you understand the problem, and it shouldn't be priced as though the hard part is done.

02

2. Cheaper output moves the burden to checking

Ali's sharpest point was about people who build an app with AI and later find it is insecure. "It's not that the AI is skipping things," he said. "If you never told it, it's not going to do it." An experienced developer thinks about login, where credentials are stored, version control, testing and what happens when two requests arrive at once. Ali calls this "the periphery stuff that you learn after years and years of developing". Newcomers don't ask for it because they don't know it exists.

He explains the gap with a publishing comparison. Building alone with AI makes you the writer, editor and publisher of a book at once, but it doesn't let you skip the editor's job. Change one word on page 199 and something breaks on page 10. Without tests in place, neither you nor the AI will know where to look.



The survey evidence points the same way. DORA describes AI as an amplifier that makes strong teams better and a struggling team's problems worse. A [later DORA analysis](#) found that time saved in creating work is often spent auditing and verifying it. Output gets cheaper, but someone still has to check it.

How much checking a decision needs depends on what a mistake costs, how often it repeats and whether anyone would notice. Jonathon draws that line inside his own business. One mistargeted ad costs very little. An identity error is different. A bad match or a missing consent signal repeats across millions of bid requests, often for weeks before anyone sees it, and a privacy error carries regulatory cost, not just wasted spend. So GrowthCode automates the high-volume decisions and puts people on monitoring and on anything expensive to reverse. It uses AI to help write job descriptions, but people interview candidates and decide who's hired, because on a small team one bad hire is expensive.

This is the part of the argument closest to my own work in project assurance. In a gateway review, the most useful questions are usually not about whether the right documents exist. They are about whether the governance team actually uses them to manage the project or produces them because it has to, whether everyone shares the same understanding of what each document is for and what it says, and how accountability is enforced and measured. AI makes polished documents much cheaper to produce, which makes it easier than ever to have the paperwork without the practice.



AI makes polished documents much cheaper to produce, which makes it easier than ever to have the paperwork without the practice.

What this means for a firm: before an AI-assisted output is used, decide what evidence it needs, who will check it, and whether they have the time and skill to do so. Three questions help set the level of checking. What would an error cost, including if it were repeated? Could it be reversed? Would anyone notice?



03

3. Firms must rethink how they earn and renew expertise

Cyrus has seen the economics shift first-hand. By his estimate, a project TouchZen once scoped at around 800 development hours now takes 300 to 400. A six-month MVP can become a three-month one, and a team of four engineers can become three. Hourly rates haven't fallen, but total hours have, and clients now ask why a project should take six months when AI-assisted teams can do it in three. TouchZen's response has been to sell judgement more directly: it now leads with discovery and senior advice, and is moving into AI code review and fractional senior engineering.

Jonathon describes a similar gain at GrowthCode, where he estimates AI has lifted what the engineering team can deliver by three to four times. He expects roles on the buy side to shift. For example, a junior media buyer doing structured optimisation is doing exactly the work AI handles well, while a media analyst, who explains what happened and whether it makes commercial sense, is in a stronger position. Even in a highly automated business like GrowthCode, engineers spend significant time watching their algorithms and code for drift.

There is a challenge running through all three interviews. Each of these leaders depends on expertise built through years of hands-on work. Ali's "periphery" knowledge came from building and breaking systems. Yet the tools they describe remove some of the work through which that expertise was built. That isn't only a ten-year problem. It raises an immediate management question: what will a junior employee actually practise next month?

Here are some options worth testing:

- Juniors make their own assessment before they see the AI's answer, then compare the two.
- They investigate discrepancies between AI output and the source material, rather than just tidying the draft.
- They defend a conclusion against the evidence in front of a senior colleague.
- They sit in on senior reviews where the reasoning is explained, not just the corrections.

All of this costs money. Senior review takes time, and so does deliberate training. If clients expect cheaper delivery, firms need to work out who pays for both. Otherwise the claim that judgement is becoming more valuable may hide a business model that pays for less of it.



Jonathon made a related point about data. If every competitor has the same models, the models stop being the advantage. The advantage is the data they run on: whether you own it, can show where it came from, and have the right to use it. For a consultancy, that might be years of engagement reports, benchmarks and methods. Those records only become an advantage if they are useful, reliable, accessible and legitimately usable. Simply having them achieves little.

What this means for a firm: price around the judgement you provide rather than the hours of production, budget review and training explicitly, and decide what your junior staff will practise when AI does the first draft.

Test it on one deliverable

Pick one piece of work your firm produces regularly, such as a board paper, a strategy review, a due diligence report or a tender response, and work through four steps:

- 1 Identify which production steps AI can shorten, and by roughly how much.
- 2 Specify what must be checked before the result is used, and by whom.
- 3 Recalculate the economics of the job, including review and training time, not just drafting time.
- 4 Identify what junior staff would stop learning, and how they will learn it instead.

By the end you will know whether AI is improving that service or just producing drafts faster. This is the kind of workflow review Driftwood Crew does with management and professional services firms: working through a real deliverable to find where automation helps and where it only moves the work somewhere else.

Where this research goes next

These interviews show what happens when production gets faster in technology businesses. The next stage of this research looks at what it means for management and professional services firms, whose product is advice and whose deliverables, fee models and ways of developing people are different. I'll be exploring how CEOs and leadership teams can create more capacity, improve engagement and accountability, and adapt to changing ways of working without simply increasing workload, and how AI and automation are changing that equation. The questions I'll be putting to leaders in those firms, and back to the original interviewees, include:

- As more routine analysis and execution is automated, which human capabilities become more important?



- How should leaders decide what can be delegated to AI and what still needs clear human accountability?
- If AI makes execution cheaper and more abundant, what becomes the scarce resource?
- Where have faster production and fewer hours actually changed fees, margins or client expectations?

About this research

This paper draws on interviews Alex Rankin conducted for Driftwood Crew in 2026 with:

- Cyrus Kiani, Founder and CEO, [TouchZen Media](#)
- Ali Syed, Founder, [NoCodeHero](#)
- Jonathon Shaevitz, Founder and Head of Product and Strategy, [GrowthCode](#)

Quotes are lightly edited for readability.

Sources

- Nathen Harvey & Derek DeBellis. [Announcing the 2025 DORA Report: **State of AI-Assisted Software Development**](#), September 24 2025
- Jessica Baolin & Nathen Harvey. [Balancing AI tensions: **Moving from AI adoption to effective SDLC use**](#), DORA, March 10 2026