insightology.ai

The Future of Market Research is ‘Warranted Thinking’

What it is — and how to charge for it

A WHITE PAPER BY INSIGHTOLOGY

AUGUST 2026

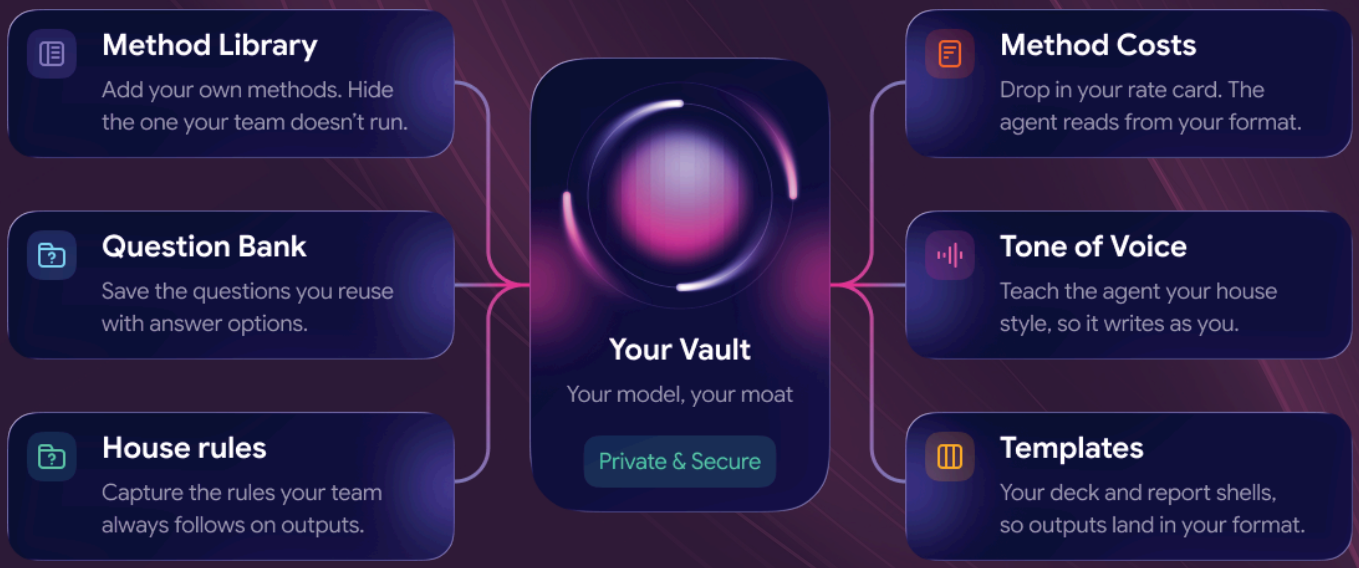

Part one — what is warranted thinking?

Plausible just became worthless

With generic AI tools, today anyone can produce a plausible research plan and a compelling deck. Your client's intern can do it with ChatGPT in an afternoon: a credible-looking methodology, professional charts, fluent recommendations. Five years ago that output would have taken a project team a fortnight. Today it costs nothing, which tells you exactly what it's worth.

This changes the market for strategic research more than any other single fact about AI. For decades, looking rigorous was expensive, so looking rigorous was a reasonable proxy for being rigorous. That proxy is dead. When the surface of good research can be generated on demand, the surface stops carrying information, and clients — slowly at first, then all at once — stop paying for it.

What they will pay for is the thing underneath the surface: thinking they can trust. Proven thinking, usable for the strategic decisions where being wrong is very expensive and may not be survivable. This paper gives it a name and treats it in full: warranted thinking. It has two halves. The warrant is everything that makes your agency's great work provable. The thinking is the set of skills that only the real people in your agency can supply. An agency needs both, because each half is worthless without the other — a point we'll return to, since it explains most of what's going wrong in the market right now.

One clarification before we start. This argument applies to strategic research specifically, and the next section explains why the word is doing so much work.

A strategic research project, by definition, informs a decision that affects the very future of the company: a brand repositioning, a market entry, a portfolio restructure, a pricing overhaul, an acquisition case. These are make-or-break calls. They get made once, they're expensive to reverse, and the organisation lives with the consequences for years.

That creates an asymmetry every agency should quote in every pitch. A strategic study might cost £200k. The decision it informs might move £100m. Getting the research wrong doesn't cost the client the research budget — it costs them the decision. Against stakes like that, the difference between a plausible answer and a proven one isn't a quality nicety. It's the entire value of the exercise.

The cheap-data paradox

Here's the strange thing about the moment we're in. Data has never been cheaper or more abundant — every organisation is drowning in dashboards, trackers, social listening and now synthetic respondents by the thousand — and yet big decisions are not getting noticeably easier or better. If anything the noise is making them harder. More cheap data raises the odds of finding a pattern; it does nothing to raise the odds of finding the right one.

Which points at the real scarcity. In a world of cheap data, the scarce and valuable things are true insight and the thinking that produces it — the ability to say, of all the options available, this is the one that will move the dial most, this is why, and this is how to execute. That is what is appreciating; everything else in the research value chain is deflating.

The warrant: what actually makes your research work provable?

Start with the half most agencies underestimate, because it feels like plumbing. The warrant is the answer to a client's hardest question — why should we bet the company on this? — and it's built from five elements. For each, it's worth knowing what failure looks like, because failures here are quiet: nothing goes visibly wrong until the day everything does.

Rigorous, validated methodology

Proven research logic, applied the same way every time — not improvised prompts that produce a different study depending on who typed them and when. Failure looks like a plausible methodology that does not really answer the question, or answers the wrong question, a sample frame that quietly excluded a specific customer type who mattered, a segmentation with questions that can't discriminate between consumer types. The test: could you show a sceptical client why this method was right, and evidence that it was followed?

Professional-grade instruments

Survey design, discussion guides and analysis tools of professional standard, producing defensible outcomes — built so they don't introduce hidden mistakes or gather the wrong type of data — data that's too shallow for true insight. This is where generic tools fail silently: a subtly biased or wrong scale, or a methodology that can't answer the questions fully or answers the wrong questions is dangerously and expensively flawed.

Real humans, carefully selected and interviewed

The right respondents answering the right questions — recruited properly, screened rigorously, incentivised ethically, answering carefully designed surveys — not synthetic shortcuts standing in for the actual customers whose behaviour the real decision depends on. That's not us dismissing the value of synthetic data — there's no question about its tactical uses. But consumer/customer behaviour and attitudes are changing faster than ever, and synthetic respondents capture a historical view of these attitudes and behaviours. Basing critical strategic decisions on synthetic responses makes strategic research just an expensive form of guessing. Yes, the client will save money on recruitment and incentives, but the true cost of getting it wrong could be 100 times greater than that.

Data security that survives an audit

Respondent and client data handled to a standard your clients' auditors will actually pass, and that complies with GDPR without a lawyer holding their breath. This element has teeth now. Generic AI cannot fully protect client or respondent data. Vendor security assessments increasingly ask AI-specific questions, and "the team uses various tools" is a failing answer. The warrant includes being able to say precisely where data goes, what it trains (nothing), and who can see it.

AI used with care

Not AI avoided — AI governed: validated workflows, an audit trail from brief to debrief, and humans checking every step that matters. Done this way, AI strengthens the warrant rather than threatening it, because a governed workflow produces exactly the evidence trail the other four elements need. The audit trail is the warrant made visible.

Notice what the five have in common: every one of them can be shown. That's the property that separates a warrant from a reassurance. "We're very experienced" is a reassurance. "Here is the method, the trail, and the person who signed it off" is a warrant.

Why the warrant can't come from tactical firms or generic AI

If the warrant were easy to supply, it wouldn't be worth much. It isn't easy, and the two cheapest alternatives to a strategic agency each fail to supply it — for structural reasons, not for lack of effort.

Generic AI can't warrant anything

A generic assistant is an extraordinary plausibility machine, and that is precisely the problem. To see why warranting strategic work with one is so hard, look at what a warrant needs and what the tool provides.

- **A warrant needs a method; a generic tool has a manner.** It predicts plausible text based on the very limited amount of market research know-how that trained it. It does not have a validated research logic to follow, so there is no method layer to check the output against — only a style of sounding methodical.
- **A warrant needs repeatability; generic output varies.** With generic tools the same brief produces different studies depending on who typed it, how, and when. You cannot warrant the quality of something that varies each time without explanation.
- **A warrant needs error signals; generic AI fails silently.** A subtly wrong methodology reads identically to a sound one — fluent, confident, formatted. There is no point at which the tool knows, or shows, that it's wrong.
- **A warrant needs evidence; the workings are a chat log.** No respondent trail, no approvals, no record of what was checked by whom. When the client asks how you know something is true, a transcript of prompts is not an answer.
- **A warrant needs compliant data handling; consumer tools fail the audit.** The security element of the warrant — vendor assessments, GDPR — breaks before the work even starts.
- **A warrant needs someone accountable; there's nobody home.** You cannot cross-examine a chat window in front of a board.

None of this makes generic AI useless — it can inform thinking brilliantly. But warranting is exactly the layer it doesn't have, and at strategic stakes that layer is the product.

The lawyers have already run this test

If the analysis sounds theoretical, look at the profession that met AI two years before research did and shares its stakes: a subtly wrong clause, like a subtly wrong finding, looks fine right up until it's ruinous. When generic AI arrived, law firms tried the prompts-and-a-general-assistant route — and the top of the profession then moved, rapidly, onto specialist legal platforms built to warrant the work: validated method, enforced process, auditable trail, enterprise security. Most top legal firms now use specialist AI systems, built specifically for the legal profession. Harvey is the best-known, and the move has been stress-tested since: when Anthropic shipped legal plugins for Claude, the specialist platforms didn't empty out — within weeks of one such launch, Harvey raised new funding in January 2026 at an $11bn valuation. The market settled into a telling shape: generic models underneath, specialist platforms as the professional layer on top, because that layer is where the warrant lives. Major firms didn't adopt specialist AI out of fashion. They adopted it because they, too, sell warranted thinking — and they discovered generic tools couldn't supply the warrant half.

Tactical firms aren't built for it either

This one needs saying carefully, because tactical research firms and DIY platforms are good at their actual job. Fast concept tests, trackers, quick polls — work where the stakes are modest, the method is templated, and speed genuinely is the value. Nothing in this paper argues otherwise.

But their economics are the economics of volume: templated methods, thin margins, junior-heavy delivery, minimal senior attention per project. A warrant is expensive in exactly the currency that model can't spend — senior methodological judgement, applied to this project's specific risks, with a senior name accountable for the outcome. A tactical firm asked to warrant a repositioning study isn't being asked to work harder; it's being asked to be a different kind of business. The occasional tactical player will make that leap. The category, structurally, cannot.

Put the two failures together and the strategic agency's position becomes clear. The warrant is the moat: too architectural for generic AI, too expensive-per-project for the volume players. It is the thing a strategic agency can sell that neither cheap alternative can counterfeit — provided the agency can actually show it.

The thinking: what only great humans can supply

Now the second half, and the reason the first half matters. Nobody warrants their way to a brilliant answer. The warrant proves the thinking; the thinking is the product. Seven uniquely human skills sit at its core and in a world of cheap data they are, more or less, all that matters. Our companion paper, The New Division of Labour in Market Research, takes each of them in turn.

These can't be delegated to AI — not because the models aren't clever, but because the skills run on things a model doesn't have: lived understanding of humans, accountability, and a stake in being right. What AI has done, though, is change the economics of deploying them — and that is where the thinking half gets commercially interesting.

Thinking time becomes a value driver

On a pre-AI project, thinking happened in the gaps around the insight process, because production of documents, materials and outputs consumed the budget. Now the AI tools can handle most of the production work, reducing the time it takes to prepare proposals, write discussion guides, design complex survey tools and uncover deep insight by up to 50%. This change has released hours of senior time that was previously dedicated to this production grunt work.

And the need for those hours is rising at exactly the moment they arrive. With data cheap and ubiquitous, the shallow findings surface themselves; the deep insight — the finding that differentiates and transforms a client's business — takes more thinking to reach, not less, because it has to be dug out from under everything the dashboards already say. The pincer works in the agency's favour: AI hands you more time to think and mine for insight, just as the work starts demanding it. The agencies getting this right are spending those hours deliberately rather than letting them drain into more email. Part two sets out how to structure that time, and how to charge for it.

Contextual intelligence — knowing the client, the category and the moment, so the answer lands in a way the organisation can actually use.

Insightfulness — seeing beyond the data to what's really going on: the throwaway comment that reveals the real purchase driver, the pattern that reframes the category, connecting two seemingly unconnected points to see something everyone else missed.

Creativity — inventing the approach that doesn't exist yet, and applying the insights to create new products, propositions and opportunities.

Strategic thinking — connecting the finding to the decision and the opportunity it informs, two moves ahead of the data.

Framing — finding the real question behind the brief, and being willing to tell a client that the question they asked is not the question they have.

Judgement — choosing between readings under uncertainty and putting a name to the choice, answerable for it in front of a board paid to be unconvinced.

Relationships — the trust that earns honesty from respondents and candour from clients, compounded over years and impossible to instruct.

What it looks like on a real project

Definitions are cheap, so here's warranted thinking traced through the anatomy of a strategic project — the version a client would actually experience.

The brief gets interrogated, not just accepted: what's the business problem behind the research question, and what decision will this inform? (Thinking: strategic, contextual.)

The method is chosen from validated options, with the trade-offs documented — an audit trail that starts before fieldwork does. (Warrant: validated methodology.)

The fieldwork instruments — discussion guides and surveys — are professionally designed, checked against the objectives, and free of the quiet biases that corrupt downstream. (Warrant: instruments.)

Fieldwork — runs with real, carefully selected humans, their data handled to auditable, GDPR-clean standards. (Warrant: real humans, data security.)

Analysis uses AI for what AI is good at — complete coverage, every source read as one joined-up set — with humans checking every step and making the interpretive leaps. (Warrant: AI with care. Thinking: insightfulness.)

The debrief is a working session on implications, not a deck reading — senior thinking, live, with the client. (Thinking: time spent with the client.)

The recommendation carries a name: someone chose it, someone will defend it, and the full trail from brief to conclusion is there to be inspected. (Both halves, together.)

And the commercial consequence runs right through it: in the AI world, only fully warranted thinking has any true value. A project built like the one above justifies the fees attached to the decision it protects — and gives the client a simple test to put to any supplier, ourselves included: show me your workings. A chat log fails that test. The project above passes it at every line.

What this means for your pricing — and why thinking time should now command a premium over doing time — is where part two of this paper turns next.

Part two — how to charge for warranted thinking

Whoops, your pricing model just broke!

AI hasn't broken the market research industry. Demand for insight is growing, the strategic end of the market is strengthening, and clients need great thinking more than ever. But AI has just broken your pricing model, and this part of the paper is about why — and what to build in its place.

Here's how agency pricing has always worked: add up your days, multiply by your day rate, then add the third-party costs with a markup applied. It's contractor pricing — exactly how a builder prices an extension — and there was nothing wrong with it. In a world where effort was the scarce input, effort was a fair proxy for value, and everyone could follow the maths. But contractor pricing only works while the hours of production work are still there to bill.

Specialist AI will remove up to half of those hours. Desk research, proposal development, survey and discussion guide design, coding, first-pass analysis, deck building — the production work that fills 50 to 70% of billable time (GreenBook's 2025 GRIT Report) is compressing fast. So the maths turns brutal. Keep charging day rates, and every hour AI saves you is a discount you just handed your client — without them even asking. On one strategic segmentation: 45 team days at a blended £960 was a £43,200 people fee. Take production at 60% of those hours and remove half of it, and the same study honestly runs to 31 days — £29,760 on the same rate card. Roughly £13,400 has vanished from a single project, and across a mid-size agency running forty strategic studies a year that is over £500,000 of fee income, gone at a constant win rate. Your margin percentage will look respectable all the way down; it's the absolute cash that disappears, and absolute cash is what pays salaries.

The market didn't do that to you. Your rate card did. Here are the four fixes

Fix one: share some of the saving — deliberately

Start with the part that feels most counterintuitive: give some of it away. Share a visible slice of the production-time saving with your clients, as strategy rather than generosity. If you don't, you'll lose work — because other agencies using specialist AI will share it, and a fifth off with faster turnaround is an easy switching argument for any procurement team. Being undercut by an equally good agency with the same tools and a keener price is not a risk worth running to protect a rate card that's already broken.

But you don't have to hand over all of it, and you must not. A useful planning rule: think of the saving in thirds — one third to clients as sharper prices, one third banked through natural attrition as your shape changes, one third kept as improved project margin. The exact split is yours; the principle isn't. Agencies that share some of the saving will take work from agencies that don't. Agencies that hand all of it back will go broke politely.

Fix two: emphasise the warrant

If clients are paying less for the production, what are they paying more for? The first answer is the warrant: everything that makes your work provable, which the cheap alternatives cannot supply. Most agencies deliver this every day and never put it on the invoice. Start reminding your clients, explicitly, of the value they get from:

- **A proven methodology** — precisely designed and accurately executed by experienced professionals, not improvised prompts.
- **Real human respondents** — and the right ones — carefully recruited and screened, not synthetic shortcuts standing in for the customers the decision depends on.
- **True insights traceable back to the raw data** — an audit trail from finding to respondent, so the answer survives challenge.
- **Strategic thinking** and interpretation only humans can provide — judgement with a name on it.
- **Recommendations your client can trust** as the basis for their big, make-or-break decisions.
- **Built with the help of professional AI** — specialist tools, governed workflows, humans checking every step. AI used with care is part of the warrant, not a threat to it.

This list is the reason the price holds when the hours fall. Anyone can produce a plausible deck now — the client's own team can do it in an afternoon — which means plausible is worthless, and warranted is what's scarce. Part one of this paper unpacked the concept in full; for pricing purposes the point is simple: the warrant is what lets you charge for thinking rather than hours. An agency that can show how its answer was reached is selling demonstrated judgement. Everyone else is still selling days.

Fix three: sell the quality and depth of your thinking

Part one set out why thinking time became the scarce input: AI hands you the hours back just as cheap, ubiquitous data makes the breakthrough insight harder to reach. That is the opportunity. Two disciplines turn it into revenue.

First, build the thinking time into your internal processes. Thinking used to happen in the gaps — evenings, taxis, the night before the debrief — because production consumed the budget. Now give it the structure production always had: named stages in every project plan — a hypothesis session when first data lands, a structured insight workshop, a challenge round where someone senior attacks the leading interpretation, an implications sprint before the debrief — each with an owner, a slot and a standard. When timelines squeeze, the unstructured thing gets skimped, and the unstructured thing must never again be the thinking, because it is precisely what clients will pay the most for.

Second, bill it. Thinking time goes into the project plan and into the price — openly. More senior hours on interpretation. Working sessions on implications with the client in the room. Co-analysis. That time used to be unaffordable; it is now the most valuable line on the proposal, and clients notice it more than anything else you've changed.

Fix four: move up the thinking value chain

Research is an input to solving a business problem — so don't price your thinking as if the research were the end of the job. Run a strategic segmentation and most agencies treat the deliverable as the segmentation model: segments, sizing, golden questions, invoice raised. Ask the client and the segmentation is the starting pistol: it exists to drive the brand repositioning, the portfolio calls, the media reallocation, the proposition work. Right now, somebody else — usually a strategy consultancy, at three times your rates — is often paid to turn your thinking into their decisions.

Picture the chain as five rungs, each closer to the client's money: findings, implications, recommendations, activation and embedded partnership — staying in the room while the decisions get made, quarter after quarter. Most agencies price at rungs one and two and give rung three away in the debrief. The value, and the pricing power, concentrates at rungs three to five.

Climbing is a capability decision as much as a pricing one: make building your business and marketing consulting muscle a priority, through hiring, training or partnering. The destination is worth naming plainly, because it is where this whole argument points: a consulting firm with insight at its heart — charging consulting-shaped fees for consulting-shaped value, with warranted research underneath that the strategy houses cannot match.

How to charge properly and fairly for warranted thinking

The four fixes tell you what to sell. Now the question of what to charge for it — and it starts with an uncomfortable admission. You cannot get away from time as the basis of cost. An agency's costs are overwhelmingly its people; its people are paid for their time; so time is what every project consumes and what every price must ultimately recover. Any pricing model, however it's dressed, sits on that fact. The real question was never whether to charge for time. It's what an hour of your time is worth — and that depends entirely on what's happening in it.

Thinking is worth more than doing

Split agency time honestly and it comes in two kinds. Doing is the production work — fielding, processing, charting, deck building. Necessary, professional, and deflating fast as AI absorbs it. Thinking is the strategy, the insight and the creativity — plus the warrant of reliability that makes it provable. Thinking is what your client's make-or-break decision actually depends on, it cannot be delegated to a machine, and it is becoming scarcer and more valuable as everything around it gets cheap.

Two kinds of time, wildly different values — and the blended day rate prices them identically. That was always a quiet mistake. In the AI world, with the doing deflating and the thinking appreciating, it's a commercial disaster. Thinking has more value than doing, and your pricing has to say so.

The consultancies already know this

Strategy consultancies typically charge far more per day than market research agencies — routinely two or three times as much, and often for work built on an agency's own data. They are not smarter than you. They have simply understood one thing, and had the confidence to act on it: great thinking has great value, and clients facing big decisions will pay for it. Nothing stops a research agency claiming the same ground except its own rate card — and its own nerve.

Charge the thinking at a premium

So separate the two in your costing, and price them differently. Doing time at a doing rate, shrinking as it should. Thinking time — senior interpretation, insight workshops, implications sessions, strategy in the room with the client — at a thinking rate, and be bold about it: twice the doing rate is a sensible starting point. Fix three puts the thinking stages into every plan and every price; this is what those stages are worth. And the warrant is what makes the premium defensible: you are not charging more for opinions, you are charging more for proven judgement with the workings to show for it.

Be willing to defend it

A premium rate only survives if you defend it, and the defence has two parts. First, show the thinking. Named stages, visible in the plan and the price, with the audit trail behind them, so you can show exactly what thinking was done, by whom, and what it changed. Clients will pay for what they can see.

Second, when procurement asks for a cut, do not quietly trim the thinking time to hit the number. Challenge the request, politely and head-on: "OK, you want to pay less — so that means you want me to spend a lot less time helping you solve your problem. Which part of the answer would you like to leave out?" Cutting the doing is efficiency. Cutting the thinking is cutting the answer — and a client who understands what the thinking is for will stop asking.

In summary: how to price work in the future

AI will drive significant changes in how work is priced and valued in our industry, but managed properly both sides win. Clients pay less than they used to and get better thinking — deeper insight, more senior attention, recommendations they can bet on — while your margins still rise, because production costs fell further than prices did. That is the whole trade, and it only fails for the agency that keeps pricing the hours. Which is why this is a pricing decision, not a technology one — and the biggest one most agencies will make this decade. The technology will keep improving whether you act or not. The only question is who captures the value of the disappearing days: you, or a rate card negotiating against you on your client's behalf.

Every force this paper has described also converges on the same point. Production is deflating toward free. Data is deflating toward free. Plausibility already is. What's left — what clients with make-or-break decisions will pay a premium for, precisely because everything around it got cheap — is thinking they can trust and prove: warranted thinking.

Supplying it takes both halves. A warrant built from validated method, professional instruments, real humans, auditable data handling and governed AI. And thinking built from insightfulness, creativity, strategic thinking, contextual intelligence, framing, judgement and relationships — given room to breathe by the production hours AI released, spent increasingly with the client, and carried up the chain to the business problem itself. Neither the volume players nor the chat windows can assemble that combination. A well-run strategic agency can, and the ones that work it out first will own the next decade of strategic research.

The commercial half follows directly: share some of the saving deliberately, emphasise the warrant, sell the quality and depth of your thinking, move up the value chain — and charge for thinking time at the premium it now commands, with the warrant to defend it.

Insightology.ai exists to supply the warrant half at platform level — validated workflows, an audit trail from brief to debrief, enterprise-grade security — with humans in control of every step that matters. What it deliberately doesn't do is the thinking. That's yours, and it's the point: the platform exists to give your thinkers their hours back and the proof their judgement deserves. Has your pricing moved yet — or is your rate card still doing the discounting for you?

About Insightology

Insightology.ai is the first AI platform built specifically for market research agencies, rather than generic AI with research prompts bolted on. It is method agnostic — qual, quant and mixed-method in one platform — rather than a point tool for a single technique, and it is built for the complex, multi-faceted programmes strategic agencies actually sell: multi-market, multi-stage work managed end to end in one place. It covers the whole strategic project: decoding the brief, methodology design, proposal writing, survey and questionnaire design, discussion guides, fieldwork materials, analysis and reporting. Everything runs as a validated workflow with an audit trail, built to reproduce the quality of expert researchers consistently.

Go to www.insightology.ai to find out more and book a demo if it looks relevant to you

It is built to do seven things: professional-grade output from brief to debrief, the same work in a fraction of the time, tuning to your agency's own methodologies and house style, a better win rate in pitches, deeper insight than surface-level tools reach, data security your clients' auditors will accept, and lower production costs that let you sharpen prices and still make more margin.

Agencies using it report time savings of 50 to 60% across proposals, fieldwork materials and analysis. Pricing is per won project with unlimited users. Proposals are free and unlimited, and credits only apply once a project is won and reaches fieldwork, so your cost follows your revenue rather than your headcount.

The argument in this guide applies to agencies doing strategic work, for clients who audit them, where warranted thinking is the real product. If that describes you, the quickest way to test it is a proper pilot on live projects, with the numbers measured.