PROBLEM? RELAX. YOU'VE GOT THIS!

The Client Wants AI, but Doesn't Know What They're Responsible For

How to Agree on the Scope of AI Use in a Project, Rules for Handling Data, and Ways to Verify Results Without Promising That the Tool Will Take Responsibility for Mistakes



Max Paradox

The Client Wants AI, but Doesn't Know What They're Responsible For

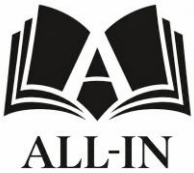
How to Agree on the Scope of AI Use in a Project, Rules for Handling Data, and Ways to Verify Results Without Promising That the Tool Will Take Responsibility for Mistakes

Max Paradox

PROBLEM? RELAX. YOU'VE GOT THIS!

ISBN: 9788384839201

ALL-IN Rafał Płonkowski



CONTENTS

INTRO 5

Chapter 1 - The Client Wants AI. But What Are They Actually Buying? 10

Chapter 2 - It Was the AI's Fault. Why That Doesn't Answer Who Is Responsible 18

Chapter 3 - The Demo Was Brilliant. Why That Doesn't Mean You Can Promise a Miracle 26

Chapter 4 - The Client Sent the Files. Does That Mean You Can Feed Them to AI? 34

Chapter 5 - Just One More Little Thing: How AI Turns a Small Project into a Very Large Invoice 43

Chapter 6 - A Brilliant Prompt Is Not a Business Process 52

Chapter 7 - Everyone Has an Opinion. Who Actually Makes the Decisions? 61

Chapter 8 - The System Is Online. Unfortunately, Nobody Knows Whether It's Still Right 69

Chapter 9 - The Contract Says "AI May Be Used." What Does That Actually Mean? 78

Chapter 10 - From the Upload Folder to the Final Deliverable: Building a Safe AI Data Workflow 87

Chapter 11 - Check the Answer, Not the AI's Confidence 96

Chapter 12 - The AI Works. But Is the Project Actually Finished? 104

Chapter 13 - AI Made a Mess. How to Handle a Complaint Before It Becomes a Crisis 113

Chapter 14 - AI Costs Pennies. Why Your Service Shouldn't 122

Chapter 15 - The Client Wants AI to Make Every Decision. When Should You Say No? 131

Chapter 16 - The AI Provider Goes Down. What Happens to Your Business Now? 139

Chapter 17 - The AI Project Is Falling Apart. How to Rescue the Work Without Pretending Everything Is Fine..... 148

Chapter 18 - It Worked Perfectly Until the Expert Left: Keeping AI Useful After the Launch 156

Chapter 19 - Nothing Broke. So Why Is the AI Suddenly Doing Things Nobody Approved? 165

Chapter 20 - The Client Wants AI. Excellent. Now You Know How to Set the Terms..... 174

INTRO

At 4:47 on a Thursday afternoon, a business owner sends a message to a consultant: “We want AI to handle our customer complaints. Can you have something ready by Monday?” The consultant asks what the system should actually do. “Read the messages, work out who's right, offer refunds when necessary, and send replies. You know, automate the whole thing.” There is a brief silence. The business owner adds, “Obviously, it needs to be accurate. We can't afford mistakes.” Apparently, the plan is to replace several human responsibilities with software, keep the existing budget, and eliminate uncertainty by adding the word “obviously.” It's an ambitious strategy. If it worked, project management would consist entirely of stating desirable outcomes in a confident voice.

The consultant now faces a familiar dilemma. Saying yes could secure an attractive contract. Asking difficult questions might make the project appear more complicated than the client expected. Yet the questions are precisely what separates a workable service from an expensive misunderstanding. Can the system access customers' purchase histories? Who determines whether a refund is justified? What happens when a complaint involves damaged goods, disputed delivery information, or a customer threatening legal action? Are responses sent automatically, or does an employee approve them? Each answer changes the nature of the assignment. Until those details are settled, nobody really knows what is being purchased. They know only that artificial intelligence will be involved, which is roughly as informative as announcing that a building project will involve electricity.

This is where many AI projects get into trouble. The technology is presented as the central feature, while the actual work remains poorly defined. Clients imagine the results they want, providers

imagine the tools they could use, and everyone assumes the missing details will become clear during implementation. Sometimes they do. Unfortunately, they often become clear shortly after somebody receives a result they dislike. Then the client insists the provider promised a complete solution, the provider explains that the system was only intended to assist, and someone produces a meeting summary containing the immortal phrase “AI-powered workflow.” It sounded perfectly adequate when the contract was signed. Now it explains almost nothing.

Artificial intelligence can provide substantial practical value. It can help organize information, produce drafts, identify patterns, retrieve relevant material, and automate selected activities. Depending on the task and implementation, it may reduce repetitive work or make previously impractical services more accessible. But generating an answer is not the same as establishing that the answer is correct. Producing a recommendation is not the same as being authorized to make a decision. And giving software access to a document does not automatically mean that sharing the document was permitted. These distinctions can seem inconvenient when everyone is excited about a demonstration. They become considerably more interesting when the demonstration turns into a service that real people rely on.

Consider the information involved in that customer-complaint project. Messages may contain names, addresses, order details, payment-related information, and other personal or commercially sensitive material. A consultant cannot reasonably assume that receiving access to those records automatically permits uploading them into whichever AI service happens to be convenient. Different services, configurations, and contractual arrangements can involve different rules for processing and retaining information. Some uses may require additional safeguards or specialist review. The practical question is not whether an application has an impressive privacy statement on its website. It is whether the proposed use of

particular information is appropriate, permitted, and consistent with the obligations of the people involved. A reassuring padlock icon is not a substitute for that assessment.

Responsibility becomes even more confusing once the system starts producing results. Suppose it incorrectly approves a refund, invents a store policy, or sends a reassuring message that contradicts the company's actual position. The technical explanation might be that the model misunderstood the available information. That tells us something about how the error occurred, but it does not settle the responsibilities of the client, the consultant, or any other participant. Those depend on the agreed work, the circumstances, and applicable rules. The useful questions concern who supplied the information, who designed the process, what checks were required, and who authorized the resulting action. An AI tool is not a convenient substitute for identifying those obligations. It cannot attend a difficult meeting and explain which contractual promise it intended to keep.

You might think the solution is simply to tell clients that AI can make mistakes. Unfortunately, that warning is about as useful as telling someone that driving involves roads. It's true, but it doesn't help them plan a journey. A client needs to understand what kinds of errors matter, how they might be detected, and what happens when something goes wrong. An inaccurate first draft of a routine promotional message might require a straightforward correction. An incorrect instruction affecting somebody's finances, rights, or safety can have very different consequences. Effective verification should reflect those differences. It also needs to be practical. A process requiring every employee to spend an hour checking each thirty-second AI response may be wonderfully cautious, but it is unlikely to deliver the efficiency everyone was promised.

There is another uncomfortable issue: money. Clients sometimes assume that if a model produces text in seconds, the entire service should cost approximately the price of those seconds. Providers

may respond by emphasizing their expertise without explaining what the expertise actually contributes. Neither approach creates much clarity. The cost of producing a raw output can be very different from the cost of delivering an agreed, usable result. Preparing information, setting requirements, checking accuracy, handling exceptions, and correcting problems all require work. At the same time, a provider should not pretend to have spent ten hours manually completing a task that was performed in twenty minutes, especially under an hourly arrangement. AI can change the economics of a service. It does not make honest pricing obsolete.

This book is for people who find themselves negotiating that uncomfortable territory: freelancers, consultants, agency owners, project managers, business clients, and anyone expected to deliver something useful with AI while remaining answerable for what happens afterward. You do not need to become a machine-learning engineer or memorize every technical term in circulation. You need to become better at recognizing unclear expectations, turning them into workable agreements, identifying important limitations, and deciding what must be checked before a result can be used. Those skills matter whether the technology is sophisticated or relatively simple. They also remain valuable when the product names, interfaces, and popular promises inevitably change.

The aim is not to turn every small assignment into a legal proceeding with a committee, a twelve-stage approval process, and ceremonial signing pens. A modest project may need only a short, clear agreement and a sensible review procedure. More consequential work may require stronger controls and professional advice. Legal, privacy, security, and specialist obligations cannot be resolved by a universal disclaimer copied into every contract. What you can do is recognize when those questions arise, understand what needs clarification, and avoid promises that go beyond your authority or knowledge. Caution should be proportional to the

actual stakes, not to the number of alarming articles somebody read over breakfast.

Most importantly, you will learn to have conversations that make responsible use of AI easier rather than more intimidating. You can ask a client what a successful outcome looks like without sounding obstructive. You can explain why certain data cannot simply be uploaded, why a particular result requires human verification, or why an additional feature changes the price. You can also recognize when an ambitious project should begin with a limited test instead of a grand commitment. None of this requires treating clients as foolish or technology as inherently unreliable. It requires understanding that commercial enthusiasm and operational readiness are two different things, even when they arrive at the same meeting wearing matching corporate presentations.

The next time somebody says, “We want AI to do everything,” you won't need to respond with either blind enthusiasm or a lecture about the end of civilization. You'll know how to bring the conversation back to the work: what should happen, what information can be used, what people must decide, and how the result will be judged. That is where a useful project begins. Because when something goes wrong, the question will not be whether the AI sounded intelligent. It will be whether the people involved understood what they had agreed to do.

Chapter 1 - The Client Wants AI. But What Are They Actually Buying?

The meeting starts with a sentence that sounds remarkably like a business requirement: "We need an AI solution." Everyone nods. The managing director looks pleased, the operations manager opens a spreadsheet, and the consultant begins imagining the possibilities. Ten minutes later, the discussion has covered chatbots, automated reports, document analysis, predictive systems, and something described as "an intelligent digital employee." Nobody has explained what problem the company wants to solve. Apparently, purchasing artificial intelligence is now a business objective in its own right, rather like deciding that the office needs more innovation and ordering seventeen additional monitors.

The difficulty is that AI describes a broad collection of capabilities rather than a particular business outcome. A client might want faster administrative work, fewer repetitive tasks, better access to information, or an entirely new service. Those needs can involve very different solutions, even when they all appear under the same fashionable label. A provider who begins by discussing models and platforms may unintentionally encourage the client to choose technology before understanding the task. The conversation becomes a tour of impressive features rather than an examination of what the organization needs. By the time someone asks about the actual work, the client may already be emotionally committed to a system capable of doing approximately everything except making sensible decisions about its own project budget.

Imagine a regional landscaping company that prepares proposals for commercial property owners. Its sales team receives requests containing site photographs, rough measurements, maintenance requirements, and sometimes elaborate descriptions

of what a customer considers "a natural but extremely tidy garden." Preparing a proposal involves reviewing the request, selecting suitable services, estimating materials, checking staff availability, and producing a professional document. The owner wants AI to "handle proposals automatically." That sounds like a clear instruction until the consultant asks what handling a proposal actually involves. Does the system draft descriptions? Calculate quantities? Recommend plants? Determine prices? Promise completion dates? Send the finished proposal directly to a potential customer? The owner initially treats these questions as minor technical details. They are, in fact, the difference between several distinct projects.

A useful discovery conversation starts with an ordinary piece of work, preferably one the client recently completed. Ask to follow it from arrival to completion. What information came in? Who interpreted it? Which decisions required experience? Where did employees spend time waiting for information? What did the customer ultimately receive? This approach is more revealing than asking the client to describe an ideal AI system, because ordinary work contains awkward details that disappear from ambitious presentations. Perhaps the landscaping company already uses a pricing spreadsheet that performs calculations reliably. Perhaps the largest delay comes from waiting for a supervisor to confirm equipment availability. Neither problem necessarily requires a generative model. AI may still help, but it should be assigned a task rather than invited to improve the entire business through sheer enthusiasm.

The consultant discovers that the sales team spends considerable time rewriting similar service descriptions and transferring information from customer requests into proposal templates. Estimating materials, however, depends on site conditions that photographs cannot always establish. Scheduling requires current information from several teams. Suddenly, the

original ambition can be divided into useful and less certain possibilities. AI could help organize the customer's requirements and prepare an initial draft using approved descriptions. Existing calculation tools could continue handling defined pricing rules. Specialists would still confirm estimates that depend on incomplete site information. This arrangement may sound less revolutionary than an autonomous proposal department. It has the considerable advantage of being understandable, testable, and connected to work people actually perform.

One reason clients struggle to describe their needs is that they know their desired improvement better than the underlying process. A business owner may say, "I want proposals completed in five minutes," without knowing whether the current delay is caused by drafting, calculation, internal approval, or missing information. It would be unreasonable to demand a perfect process map before agreeing to an initial discussion. The provider can help identify the relevant steps by asking employees to describe a recent difficult request. If the answer is that they spent two days waiting for photographs of an inaccessible area, no amount of faster text generation will remove that dependency. Technology cannot manufacture missing evidence simply because a project plan describes the information as essential.

It also helps to distinguish a tool the provider uses while delivering a service from an AI product the client will operate after delivery. The consultant might use AI internally to prepare proposals, check wording, or organize project notes, while remaining responsible for the agreed deliverables. Alternatively, the landscaping company might commission software that its employees will use every day. That second arrangement introduces questions about user access, training, support, updates, and operational limitations. The client may not recognize the distinction because both options are described as "using AI." Yet a consultant delivering twenty approved proposal templates is

undertaking something quite different from a developer building a system expected to produce thousands of future drafts. Confusing the two creates expectations that become expensive long after the original invoice has been paid.

The next distinction concerns how independently the proposed solution should act. There is a meaningful difference between suggesting a price and approving it, between drafting a proposal and sending it, and between recognizing a customer preference and committing the company to satisfy it. Each additional power changes what can happen when the system makes a mistake. A drafting assistant might describe a maintenance package too enthusiastically, allowing a salesperson to correct the wording. An autonomous sending system might promise weekly maintenance for the price of a monthly visit. The error is similar at the text level, but its business consequences are not. Before discussing how advanced the technology can become, identify which actions it should be allowed to perform without further authorization.

A practical way to reveal hidden assumptions is to ask the client what should happen when information is missing. Suppose a customer requests a landscaping estimate but provides no dimensions and only one photograph of a large, mysterious patch of greenery. Should the system request measurements, prepare a clearly labeled preliminary proposal, or transfer the case to an employee? Each is potentially reasonable. What would be unreasonable is allowing the system to invent an area, calculate a confident price, and present it as an approved quotation. The fact that a model can complete a sentence does not mean it possesses the missing dimensions. The client's answer to this question often reveals more about the desired service than an hour spent choosing between competing AI products.

The people who will use the solution deserve a voice in the discussion. Senior managers may want one-click automation because they see the time consumed by a process. Frontline

employees often understand why that time is consumed. A salesperson may explain that customer photographs rarely show drainage problems, or that a particular maintenance package is available only in certain locations. These details affect whether an automated proposal can be trusted. They also reveal which tasks employees would welcome help with. Consulting users need not involve an expensive workshop filled with sticky notes and people inventing new names for ordinary meetings. A focused conversation with two experienced employees and several representative requests can uncover requirements that a management presentation entirely misses.

Another common source of confusion is the difference between delivering a feature and delivering a business improvement. The provider may demonstrate that AI can draft a proposal in thirty seconds. The client concludes that the entire proposal process will now take thirty seconds. Unfortunately, someone may still need to verify quantities, obtain approval, confirm availability, and respond to the customer's questions. Those activities belong in the evaluation of the proposed solution. If drafting previously took forty minutes and now takes five, the improvement may be worthwhile even when the overall process still requires several hours. But the provider should not sell a five-minute service merely because one component operates quickly. A stopwatch placed beside a chatbot is not a complete business case.

The conversation should eventually produce a description of the intended result in ordinary language. For the landscaping company, that might be a system that receives approved customer-request information, organizes requirements into standard fields, and produces an editable proposal draft based on the company's current service catalogue. Employees would review estimates and approve documents before they are sent. This description is not yet a complete specification. It is a shared understanding of what the proposed solution does and where people remain involved. It also

makes later discussions about data, testing, cost, and responsibilities much easier. Without such an understanding, every meeting risks returning to the original declaration that the company wants something intelligent.

It is equally useful to describe what the solution is not expected to do. This need not be an enormous catalogue of imaginary features. Focus on exclusions that a reasonable client might otherwise assume are included. If the system will not independently inspect site conditions, approve quotations, or confirm scheduling, say so. If a later phase could add certain capabilities, distinguish those possibilities from current commitments. Clients generally appreciate flexibility, but flexibility should concern how the project develops rather than the meaning of an already agreed price. A statement that "additional intelligent functionality may be explored" should not accidentally become an obligation to build an entirely new sales department inside a browser window.

When the client cannot yet identify the appropriate solution, a limited discovery or trial stage may be worthwhile. It should have a defined purpose, such as determining whether representative customer requests contain enough information to generate useful drafts. The provider can agree on suitable examples, identify what will be evaluated, and explain what decisions the results will support. This is different from offering an unlimited free demonstration until the customer feels sufficiently impressed. A useful trial may reveal that the proposed system works for routine maintenance requests but struggles with complex design projects. That finding is valuable because it supports a narrower, more realistic decision. A trial that establishes a limitation has not necessarily failed. It may have prevented a much more expensive failure.

Some clients resist this process because they want immediate certainty. They may face pressure to present an AI initiative at an

upcoming board meeting or to justify a technology budget before the end of the financial period. A long series of open questions will not help them. The provider can respond by identifying the decisions that genuinely matter now and separating them from details that can wait. For example, the company might first agree that the solution will create editable drafts from approved information, with pricing remaining outside the initial scope. That is enough to establish a meaningful direction without pretending that every interface detail has been settled. Good discovery reduces uncertainty in a deliberate order rather than demanding complete knowledge before anyone can begin.

Before work starts, repeat the proposed arrangement in language the client can confirm or correct. Explain the purpose, expected output, major dependencies, and limits of the system's independence. Then ask the client to describe how employees will use the result. If the consultant says "draft for review" and the owner replies "excellent, so customers will receive their quotes automatically," there is still a disagreement to resolve. Discovering that disagreement during a conversation is considerably cheaper than discovering it after launch. A brief written confirmation preserves the shared understanding and gives everyone something more reliable than memories of a cheerful meeting.

The best outcome of this discussion is not necessarily a larger AI project. Sometimes a simple workflow improvement will solve the problem. Sometimes the organization needs better information before advanced automation becomes useful. And sometimes AI genuinely offers an excellent solution once its role has been defined. A competent provider should be willing to recognize all three possibilities. The goal is to help the client purchase work that addresses a real need, not to ensure that every business difficulty receives a decorative layer of artificial intelligence. When someone says, "We want AI," the appropriate response is curiosity about the

result. The technology can come afterward, when there is finally something specific for it to do.

Chapter 2 - It Was the AI's Fault. Why That Doesn't Answer Who Is Responsible

Imagine a coworking company preparing to launch a new membership package. It hires a marketing freelancer to produce a series of promotional emails using approved information about the service. The materials look excellent. The subject lines are engaging, the benefits sound attractive, and the emails are delivered ahead of schedule. Unfortunately, one paragraph announces that members receive unlimited access to meeting rooms. The actual package includes four hours per month. The campaign goes out to thousands of potential customers, and several people sign up expecting the advertised benefit. When the client discovers the problem, the freelancer explains that the AI generated the wording. The client responds with a question that no language model can conveniently answer: "And why did you send it to us as finished work?"

The freelancer might be correct about how the inaccurate statement originated. The model may have introduced the word "unlimited" while turning a collection of product features into promotional copy. But identifying the source of a sentence is not the same as determining responsibility for its use. The client supplied information, the freelancer prepared materials, someone approved the campaign, and the business distributed it. Each stage involved actions and expectations that may matter when examining what went wrong. Saying that the AI made the mistake explains only one part of the story. It is rather like a building contractor explaining that a wall is crooked because the measuring device displayed the wrong number. The explanation may be relevant, but it does not automatically settle whether the work met the agreed requirements.

Responsibility is often discussed as though it were a single package that must be placed on one person's desk. Either the client is responsible, or the provider is responsible, or perhaps the technology company can somehow be persuaded to accept the entire problem. Real projects are usually more complicated. Different participants may have different contractual duties, decision-making authority, and obligations under applicable law. The marketing freelancer may have agreed to verify promotional claims against the supplied information. The client may have undertaken to approve certain commercial terms before publication. A service provider may have separate obligations concerning the operation of its platform. These roles can overlap, and their legal consequences cannot be determined from a generic example alone. Clear agreements help establish what each participant was expected to do; they do not automatically override applicable legal requirements.

A useful way to examine the coworking incident is to reconstruct the path of the incorrect information. The original membership document stated that members could reserve meeting rooms for four hours each month. The AI-generated draft described "generous meeting room access." During a later revision, the wording became "unlimited access." The freelancer delivered the campaign, a marketing manager approved its visual presentation, and the emails were scheduled for distribution. This sequence raises several distinct questions. Why was the source information altered? Was the freelancer expected to verify claims? What was the marketing manager actually reviewing? Did anyone compare the final copy with the approved membership terms? Investigating these questions is more useful than gathering everyone into a conference room and asking who would like to volunteer for responsibility.

One important distinction is between responsibility for supplying information and responsibility for processing it. A client