AI Clauses in Commercial Contracts: A Practical Guide.

RELEASED JULY 2026

Table of Contents

AI Clauses in Commercial Contracts: Clauses Becoming Standard

AI Clauses in Commercial Contracts: Provisions Still Under Contention

AI Clauses in Commercial Contracts: **Clauses Becoming Standard**

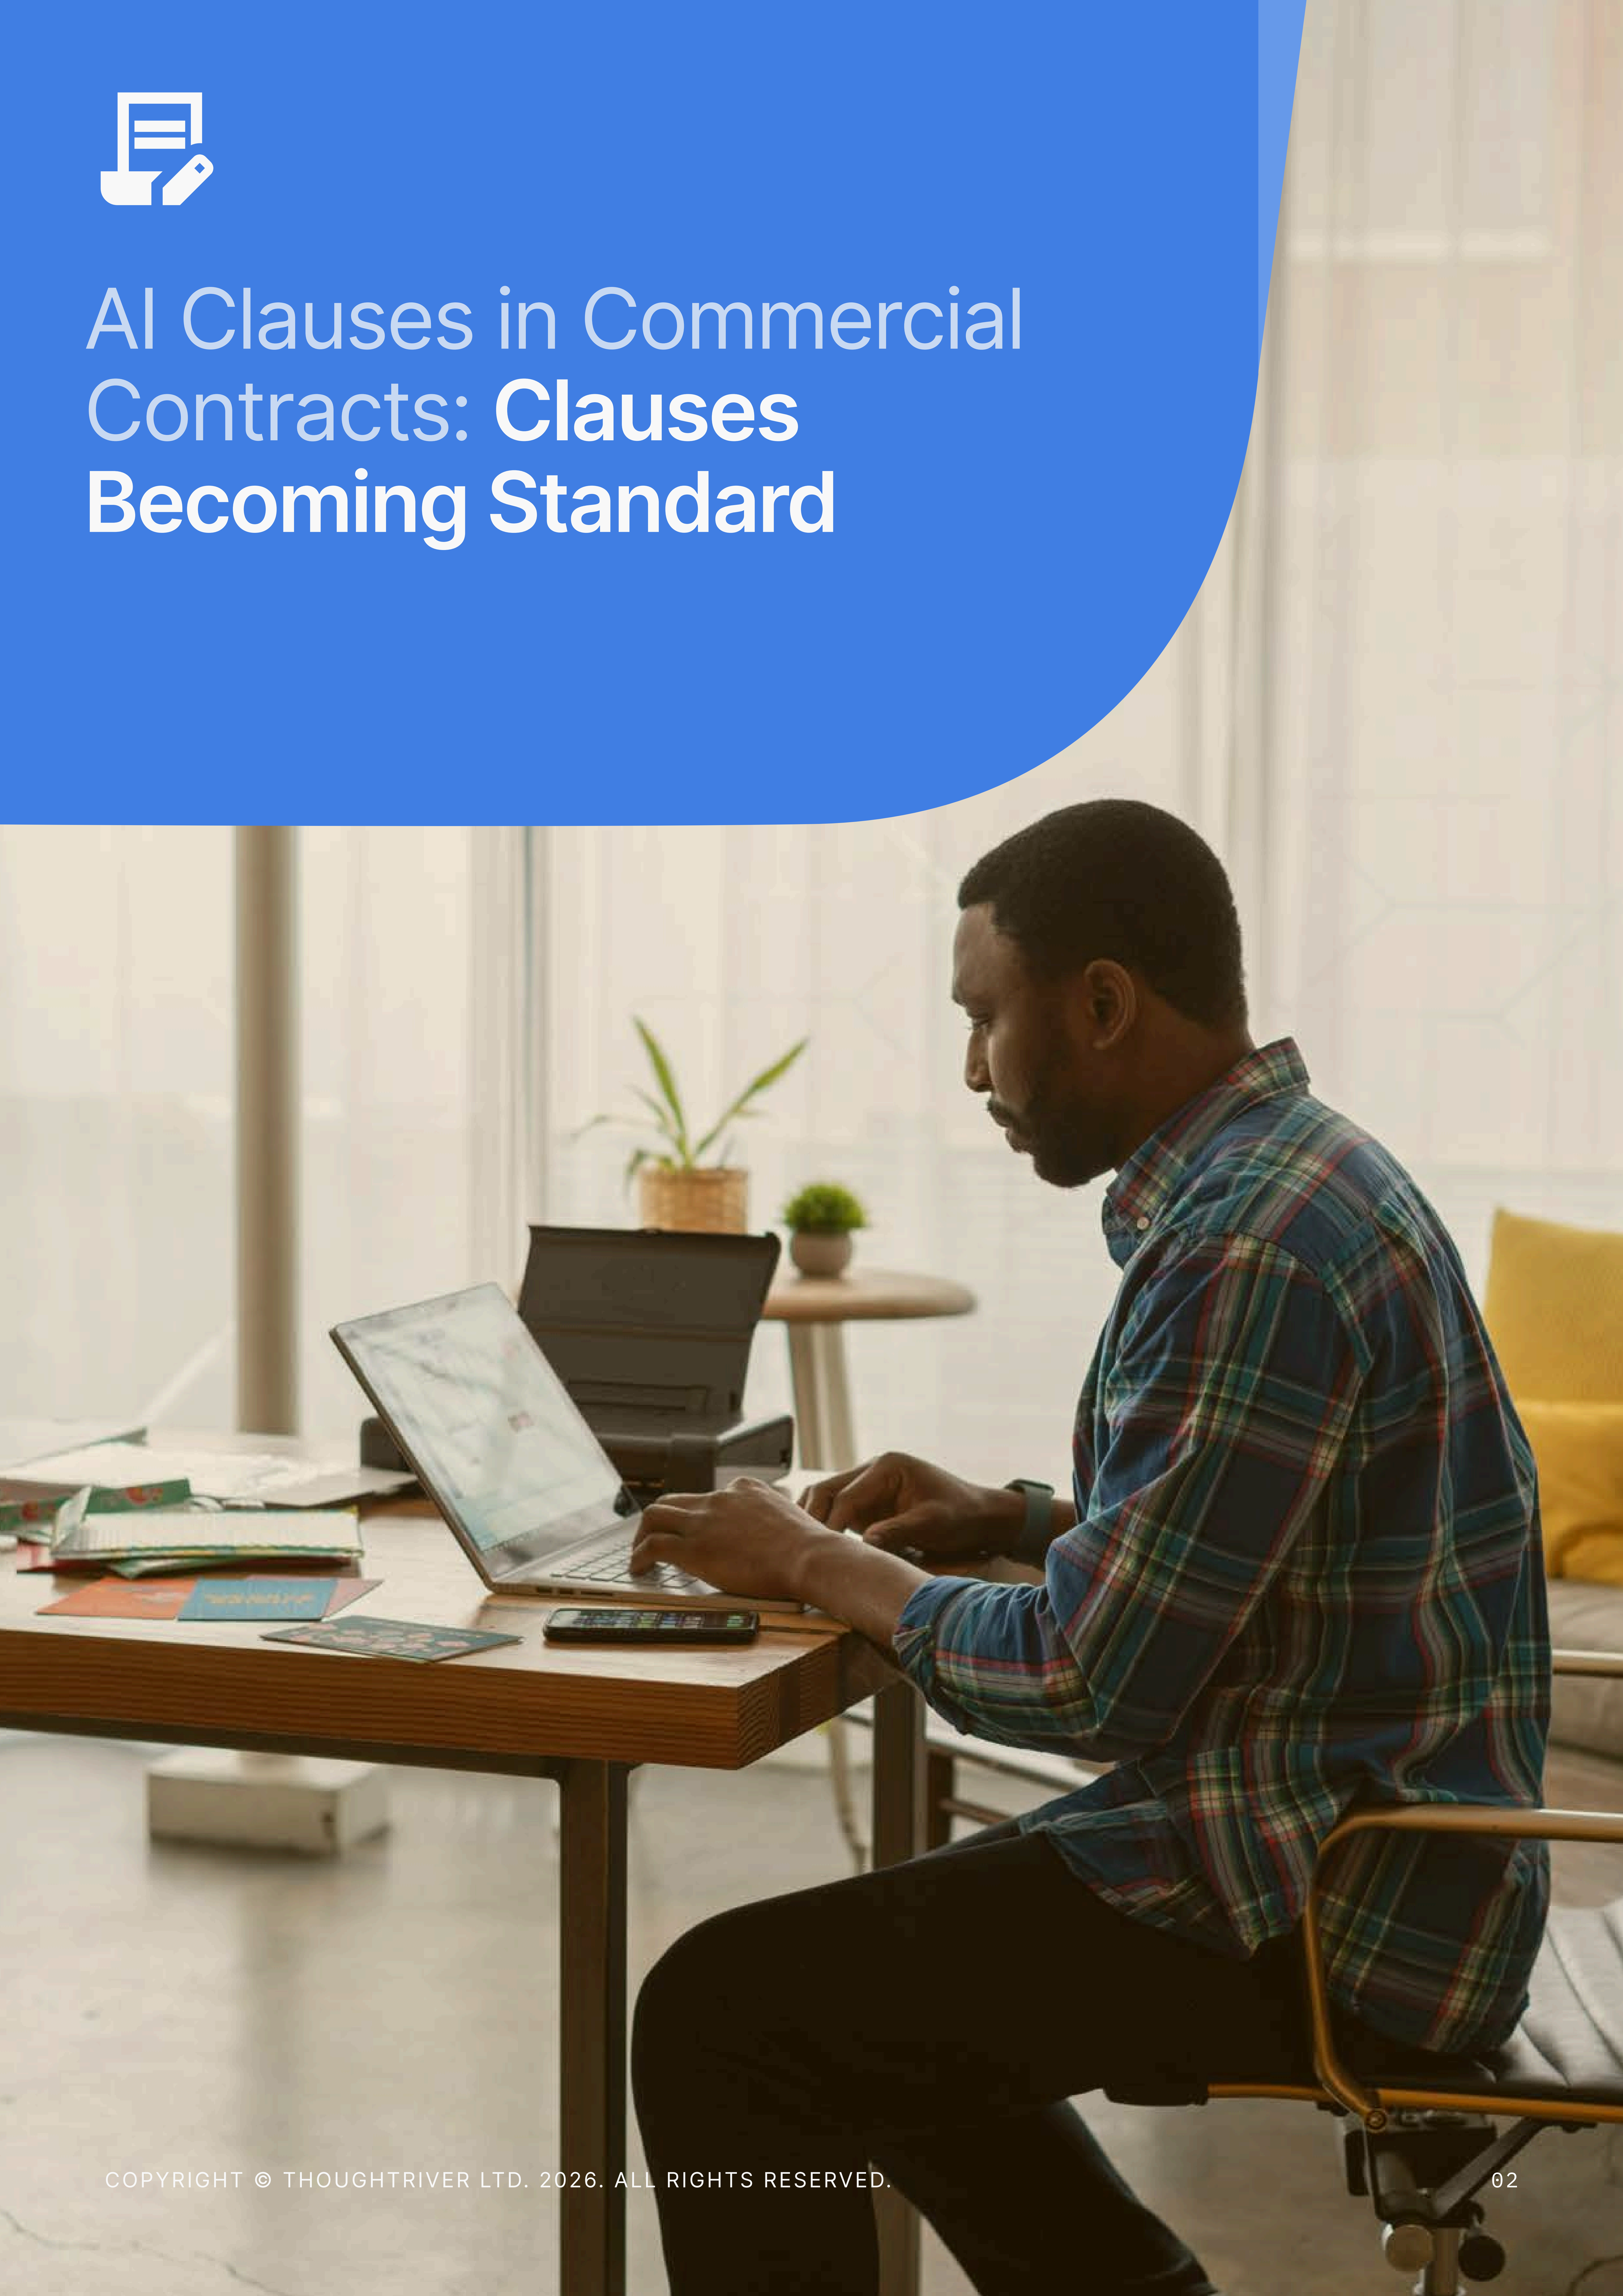

Summary

AI is now embedded in how businesses deliver products and services. Whether you're buying consulting services, licensing software, outsourcing manufacturing, or entering into a distribution agreement, there's a growing likelihood that your counterparty is using AI somewhere in the delivery chain. And that means your contracts need to address it.

This isn't about procuring 'AI services' as a distinct category. This is about **any commercial agreement** where one or both parties are introducing AI-related provisions, data usage restrictions, training prohibitions, output ownership, disclosure obligations into the standard terms.

At ThoughtRiver, we've been reviewing contracts with AI for nearly a decade. We've watched this shift happen in real time.

What follows is a practical guide to the AI clauses that are becoming standard, the provisions still being heavily negotiated, and the specific language and strategies that matter from both buyer and supplier perspectives.

Clauses Becoming Standard

Across commercial agreements of all types — SaaS, professional services, manufacturing, distribution and others — a core set of AI-related provisions is now expected.

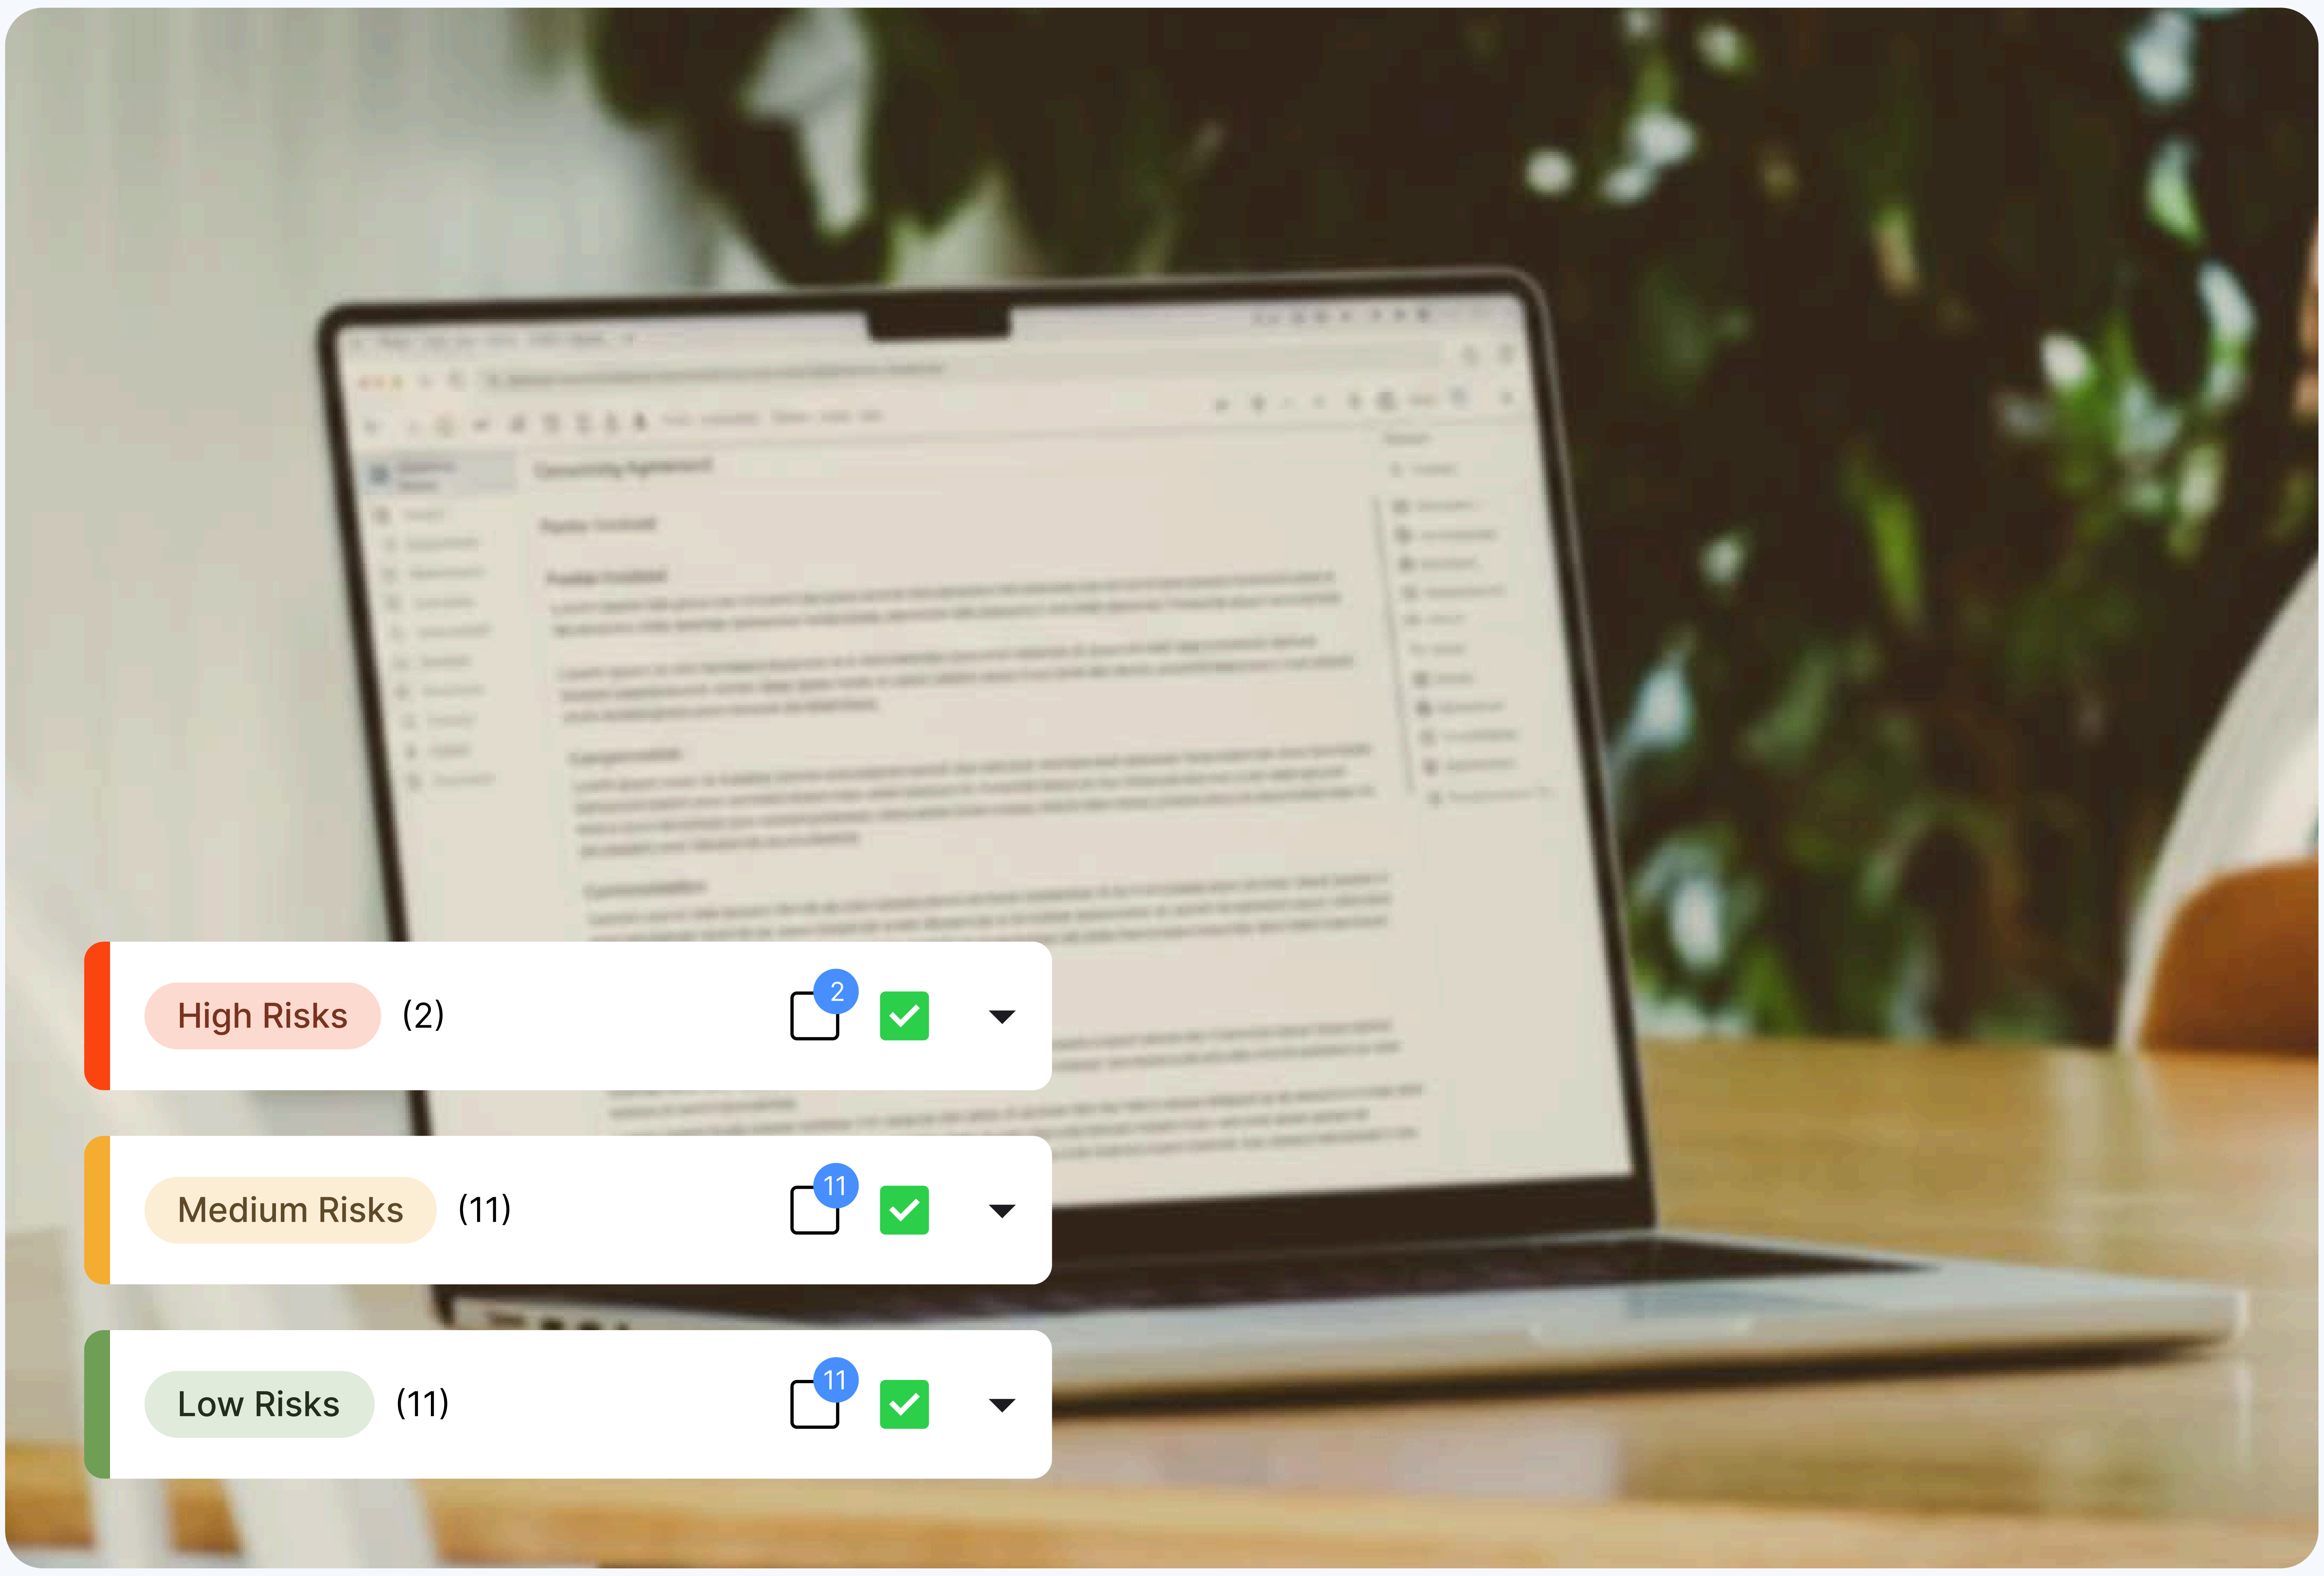

1. Definitions

What buyers want: A clear, broad definition of AI that doesn't become outdated as technology evolves.

It should clearly define how AI is used, such as in the service of the buyer, for the benefit of the seller, as well as how any of the buyer's data is used for the service and model training.

Be wary of general purpose or catch-all statements as that can be an open door to use your data in ways you may not want.

Example language (buyer-friendly):

"Artificial Intelligence" or "AI" means any automated system, software, algorithm, or technology that performs tasks requiring human-like perception, reasoning, learning, decision-making, or prediction, including but not limited to machine learning models, natural language processing systems, computer vision systems, generative AI, and any successor or related technologies incorporated into the Services.

What suppliers resist: Definitions so broad they capture basic automation or rule-based systems (e.g., Excel macros, basic if/then logic).

Compromise position:

Add a carve-out for basic automation and AI that is not connected to any inputs provided by the Buyer.

Red flags:

Definitions that are overly narrow (e.g., "AI means only systems explicitly marketed as artificial intelligence by Supplier") leave too much room for evasion.

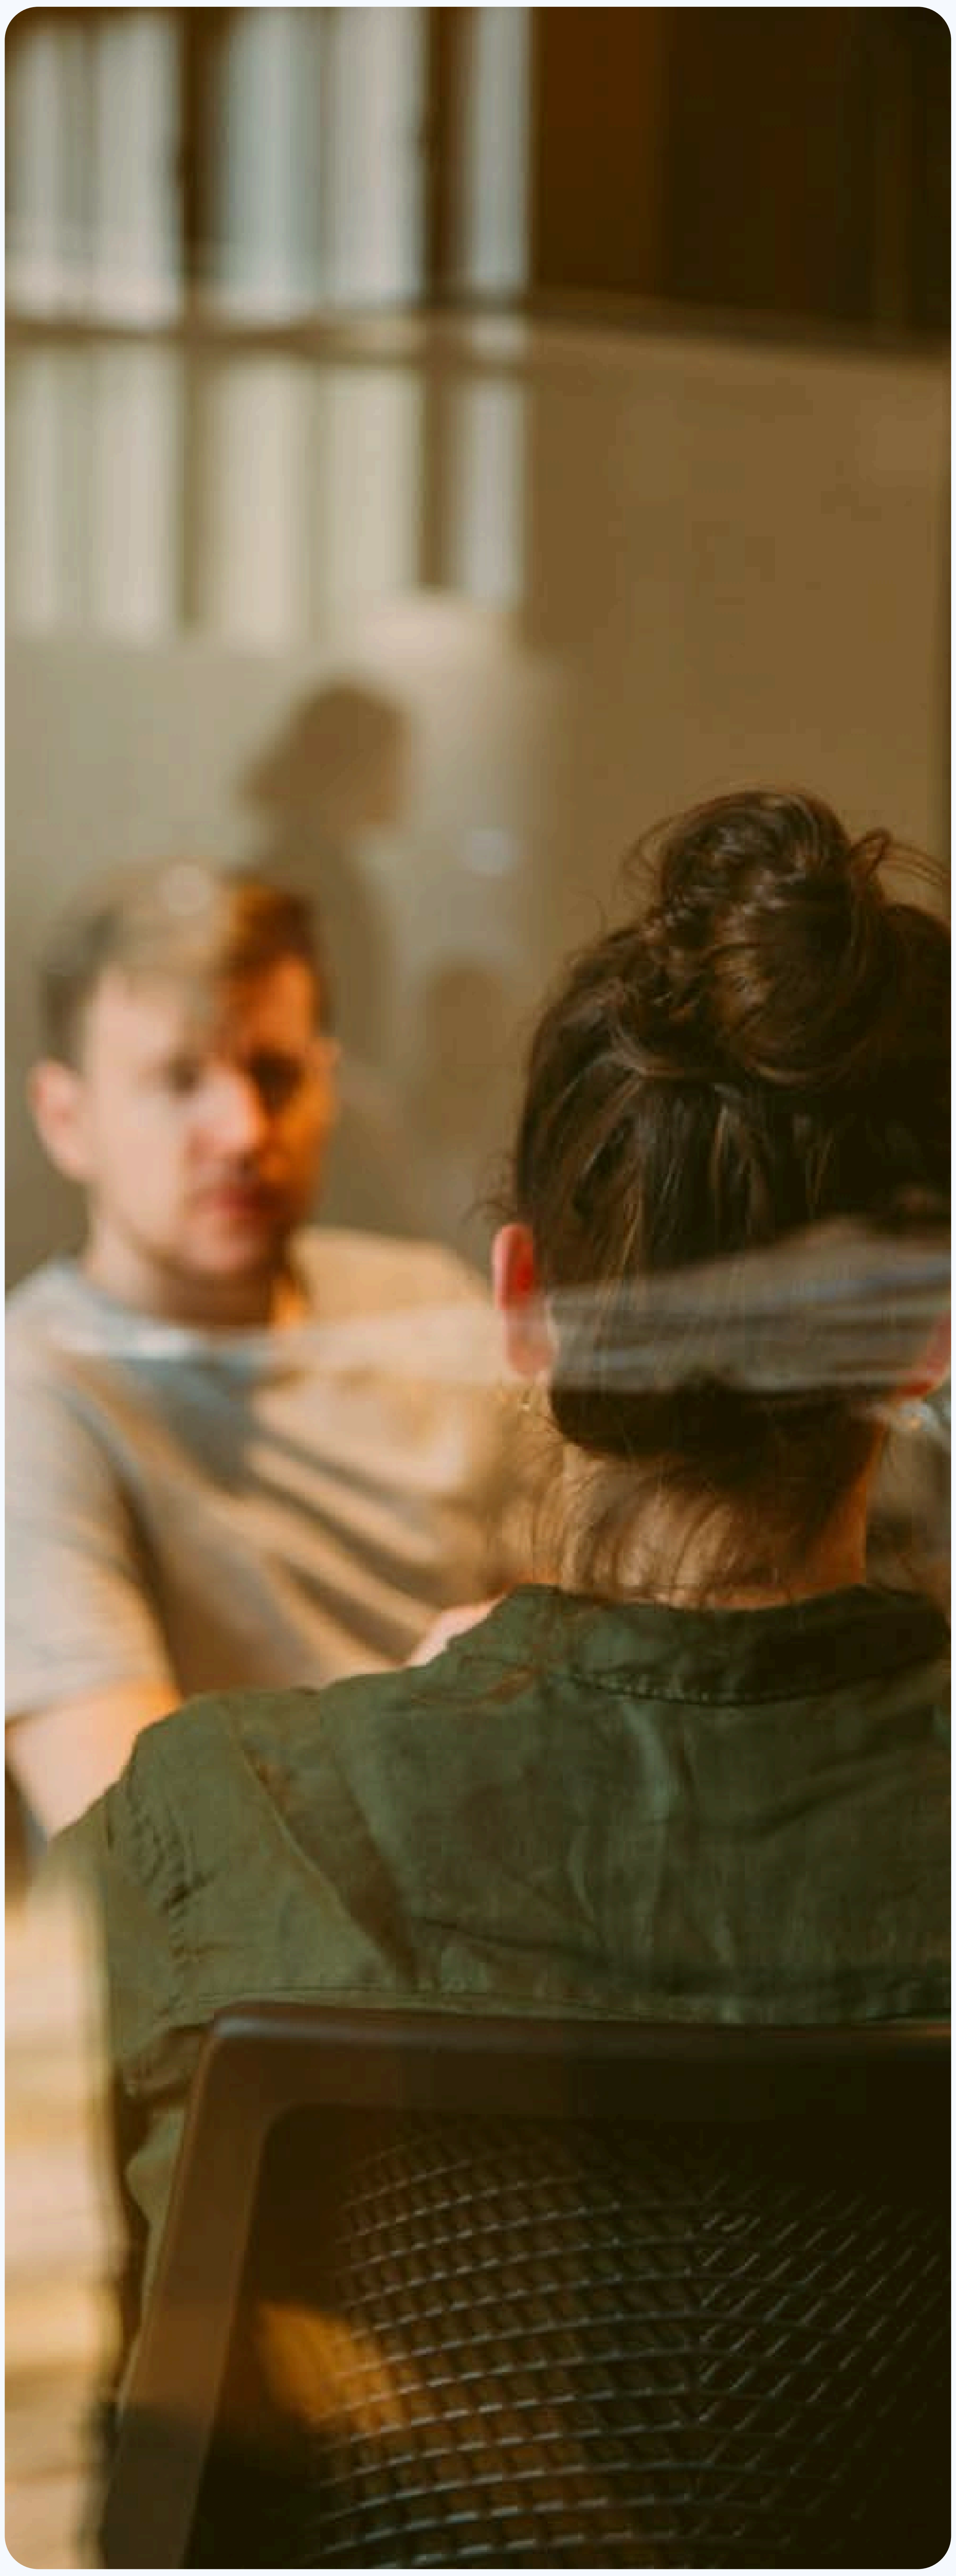

2. Disclosure and Transparency

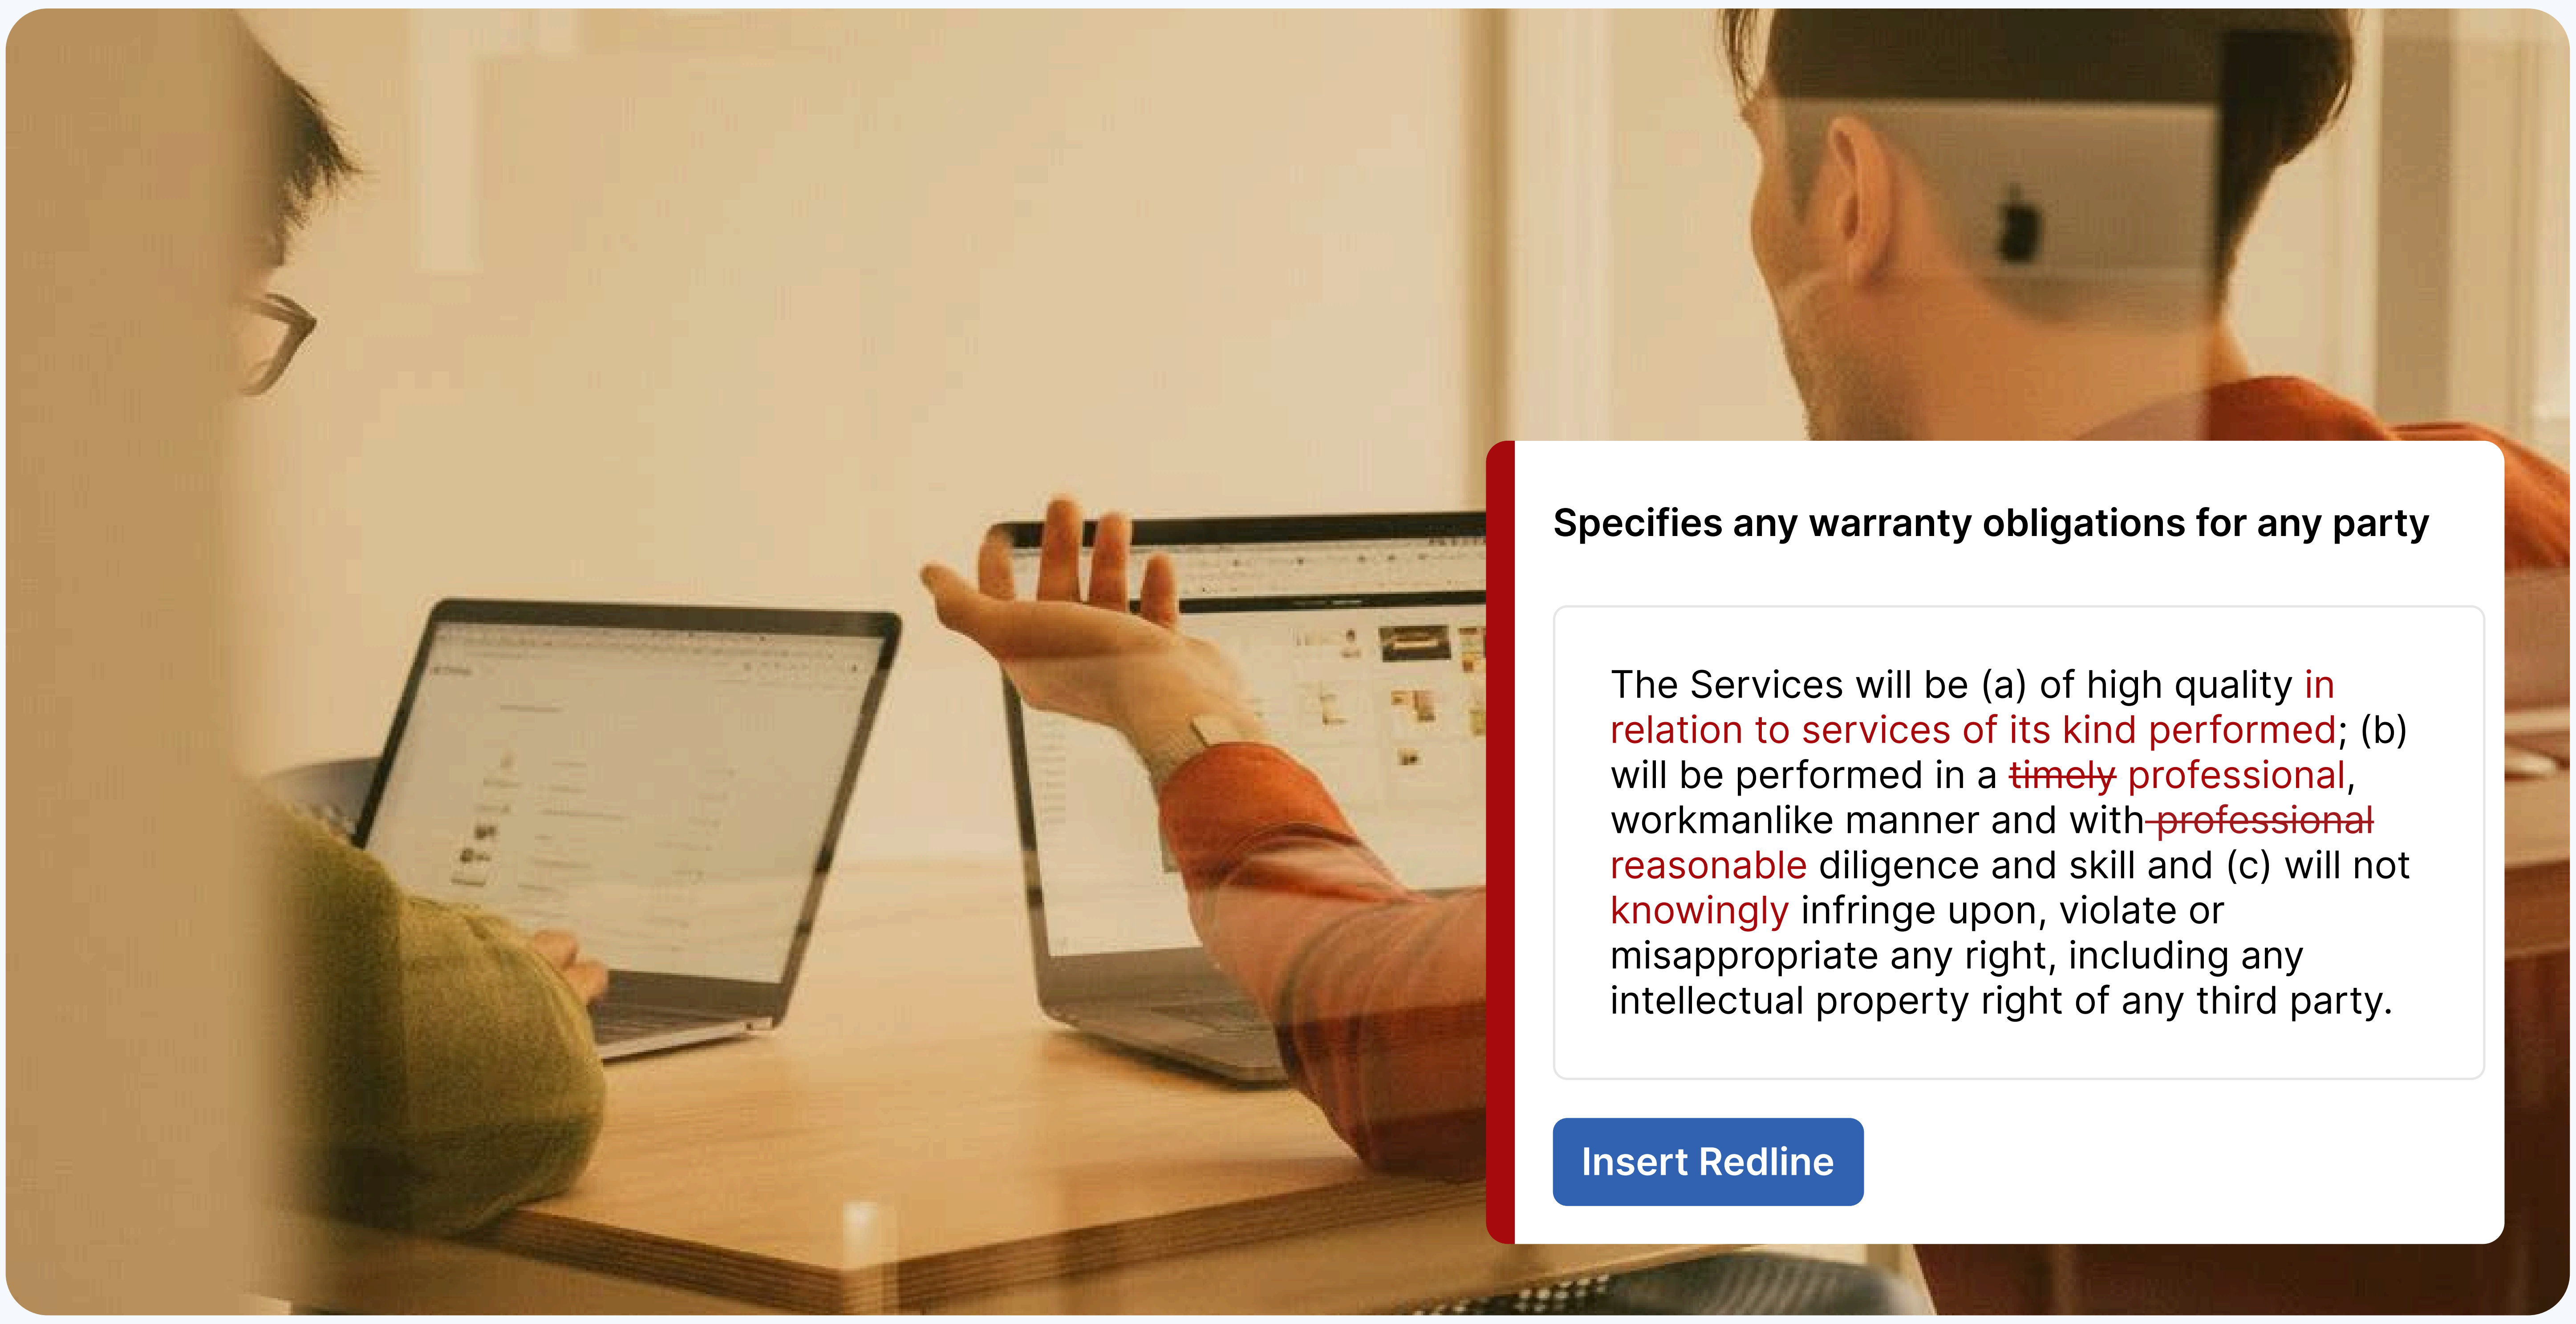

What buyers want: Upfront disclosure of which AI tools are being used, and advance notice before any changes.

Example language (buyer-friendly):

Supplier shall provide Customer with a written list of all AI systems used in connection with the Services within 30 days of the Effective Date and shall update such list within 30 days of any material change.

Supplier shall provide Customer with at least 60 days' advance written notice before deploying any new AI system that will process Customer Data or materially affect Service delivery.

What suppliers resist: Detailed, tool-by-tool disclosure, model disclosure (citing competitive sensitivity as well as the administrative burden) and lengthy advance notice periods that slow down innovation.

Compromise position:

Disclosure at a category level with a specific purpose essential to its utility (e.g., "natural language processing for customer support," "machine learning for fraud detection") rather than naming specific vendors/ models.

Reduce notice period to 30 days for non-critical changes but maintain 60+ days for changes involving customer data or high-risk use cases.

Red flags:

Clauses that permit AI deployment "at Supplier's discretion" with no notice obligation whatsoever.

3. Data and Training Restrictions

This is the most commercially sensitive area. Buyers want absolute control over whether their data is used to train or improve AI models. Suppliers often resist blanket prohibitions.

Example language (buyer-friendly):

Supplier shall not use Customer Data, including any inputs, prompts, queries, or outputs generated in connection with the Services, to train, fine-tune, or otherwise improve any AI model or system, whether proprietary to Supplier or provided by a third party, without Customer's prior written consent. This restriction survives termination of this Agreement.

What suppliers resist: Total prohibition on using anonymised or aggregated data to improve their platform.

Compromise position:

Supplier may use aggregated, de-identified usage data (not including any Customer Data content or input) to improve the Services only and for no other purpose, provided such data cannot reasonably be re-identified. Supplier shall not use any Customer Data content, prompts, or outputs for training purposes.

Be sure the lower-case term “data” is clearly understood in the agreement.

Remember: outputs can be data and still contain sensitive client information. In addition, outputs can trigger a reasonable guess as to the inputs as well.

Buyers may want to push for an annual right to audit how de-identification is performed and confirm it meets recognised standards for your industry or dataset.

But keep in mind that suppliers often view audits as a back door to trade secrets: how they create their algorithms and build their technology.

If audit rights are granted, they will be limited to reasonably account for confidentiality, and a third party (other than the client) may not be permitted to conduct an audit.

Red flags:

- Clauses that permit use of Customer data "to improve the platform" without defining what that means or limiting it to non-content metadata
- Any language allowing training "with anonymisation" without specifying the de-identification method or providing audit rights
- Suppliers who refuse to commit in writing that customer data won't be used for training

4. Intellectual Property Ownership of AI Outputs

When AI generates content such as code, designs, reports or marketing copy, contracts must clarify who owns it.

Example language (buyer-friendly):

All outputs generated by AI systems in connection with the Services, including but not limited to text, images, code, designs, and reports, shall be the sole and exclusive property of Customer. Supplier hereby assigns to Customer all right, title, and interest in such outputs.

What suppliers resist: Assigning IP when the AI model itself (or templates/prompts used to generate outputs) is proprietary to the supplier.

Compromise position:

Customer owns all AI-generated outputs. Supplier retains ownership of its underlying AI models, algorithms, training data, and any general-purpose templates or prompts. Supplier grants Customer a perpetual, irrevocable, worldwide licence to use the outputs for any purpose.

Red flags:

Clauses that give suppliers joint ownership of outputs or retain any rights to customer-specific outputs. Also watch for restrictions on using outputs (e.g., "for internal use only").

5. Regulatory Compliance and Change Management

The **EU AI Act** came into force in 2024. US states are passing their own AI laws. Compliance obligations are expanding, and contracts need to keep pace.

Example language (buyer-friendly):

Supplier warrants that its use of AI systems complies with all applicable laws, regulations, and industry standards, including but not limited to the EU AI Act, UK AI regulations, and any applicable US federal or state AI legislation.

If changes in law require modifications to the AI systems or Services, Supplier shall implement such modifications at no additional cost to Customer within 30 days of such new law going into effect.

What suppliers resist: Bearing the full cost of compliance with future, unknowable regulations, particularly when those regulations impose expensive technical requirements.

Compromise position:

Supplier shall comply with applicable laws.

Red flags:

Any disclaimer of regulatory compliance obligations (e.g., "Services are provided 'as is' with respect to regulatory compliance").

Much like AI itself, legal terms are evolving quickly. The more quickly critical clauses can be identified and tied back to other obligations within the agreement, the more effectively legal leaders can ask the right questions to manage risk and opportunities for their organisations.

6. AI Audit Rights and Transparency

Without the ability to verify how AI is being used, governance is theoretical. However, audits no longer tell the full story as they have in previous generations of software development.

On the one hand, understanding how a supplier uses AI – particularly with a buyer's data – is useful for confidential and IP right tracking.

On the other hand, use of that data may be intertwined in proprietary data orchestration or systems which the Supplier has a legitimate interest in protecting. It can be beneficial or a fool's errand – administratively burdensome to both sides.

Consequently, use sparingly with very specific terms to address the business issue that is clearly illuminated to both sides.

Example language (buyer-friendly):

Customer may, upon no less than ten (10) business days' prior written notice, audit Supplier's AI systems, models, training data, algorithmic decision-making processes, and any third-party AI tools or sub-processors used in connection with the Services provided to Customer. Customer shall use reasonable efforts to avoid material disruption to Supplier's operations during any audit conducted in the absence of a Security Incident or material compliance concern.

Supplier shall provide Customer or its designated auditor with access to relevant documentation, system logs, and personnel to verify compliance with this Agreement.

Supplier shall bear the cost of such audits unless non-compliance is identified, in which case Customer's reasonable audit costs shall be reimbursed.

What suppliers resist: Unlimited audit frequency and scope, particularly when it exposes proprietary model details or competitive information. Also, the "use of AI systems" is extremely broad. Narrowing it to precisely what is necessary to identify the requisite behaviour is more likely to yield agreement.

Compromise position:

Annual audits with scope limited to verifying compliance with data usage, training restrictions, a sufficient notice period and security obligations. Customer may use a mutually agreed third-party auditor bound by confidentiality. Supplier's proprietary model architecture and training data sources are excluded from scope unless there's evidence of a breach.

Red flags:

Complete absence of audit rights, or language limiting audits to "on-site only" (which can be impractical for cloud-based services).

7. Incident Response and Breach Notification

AI systems can be breached. AI outputs can cause harm. Contracts need clear notification and remediation obligations.

Example language (buyer-friendly):

Supplier shall notify Customer within 24 hours of becoming aware of any: (i) unauthorized access to or breach of any AI system processing Customer Data; (ii) AI-generated output that is materially inaccurate and has been or may be relied upon by Customer or third parties; or (iii) regulatory investigation related to Supplier's AI systems. Supplier shall provide a root cause analysis within 10 business days and implement corrective measures at no additional cost.

What suppliers resist: Very short notification timeframes (particularly for accuracy issues, which may require investigation to confirm) and bearing all costs of remediation. Also, Suppliers are not going to take liability for output yet.

Compromise position:

24-hour notification for security breaches; 48-72 hours for accuracy issues (allowing time for investigation); 5 business days for regulatory matters. Supplier bears remediation costs for incidents caused by Supplier's negligence or breach; costs are shared for incidents outside Supplier's reasonable control.

Red flags:

No incident notification obligations, or notification timelines of "reasonable time" without defining it.

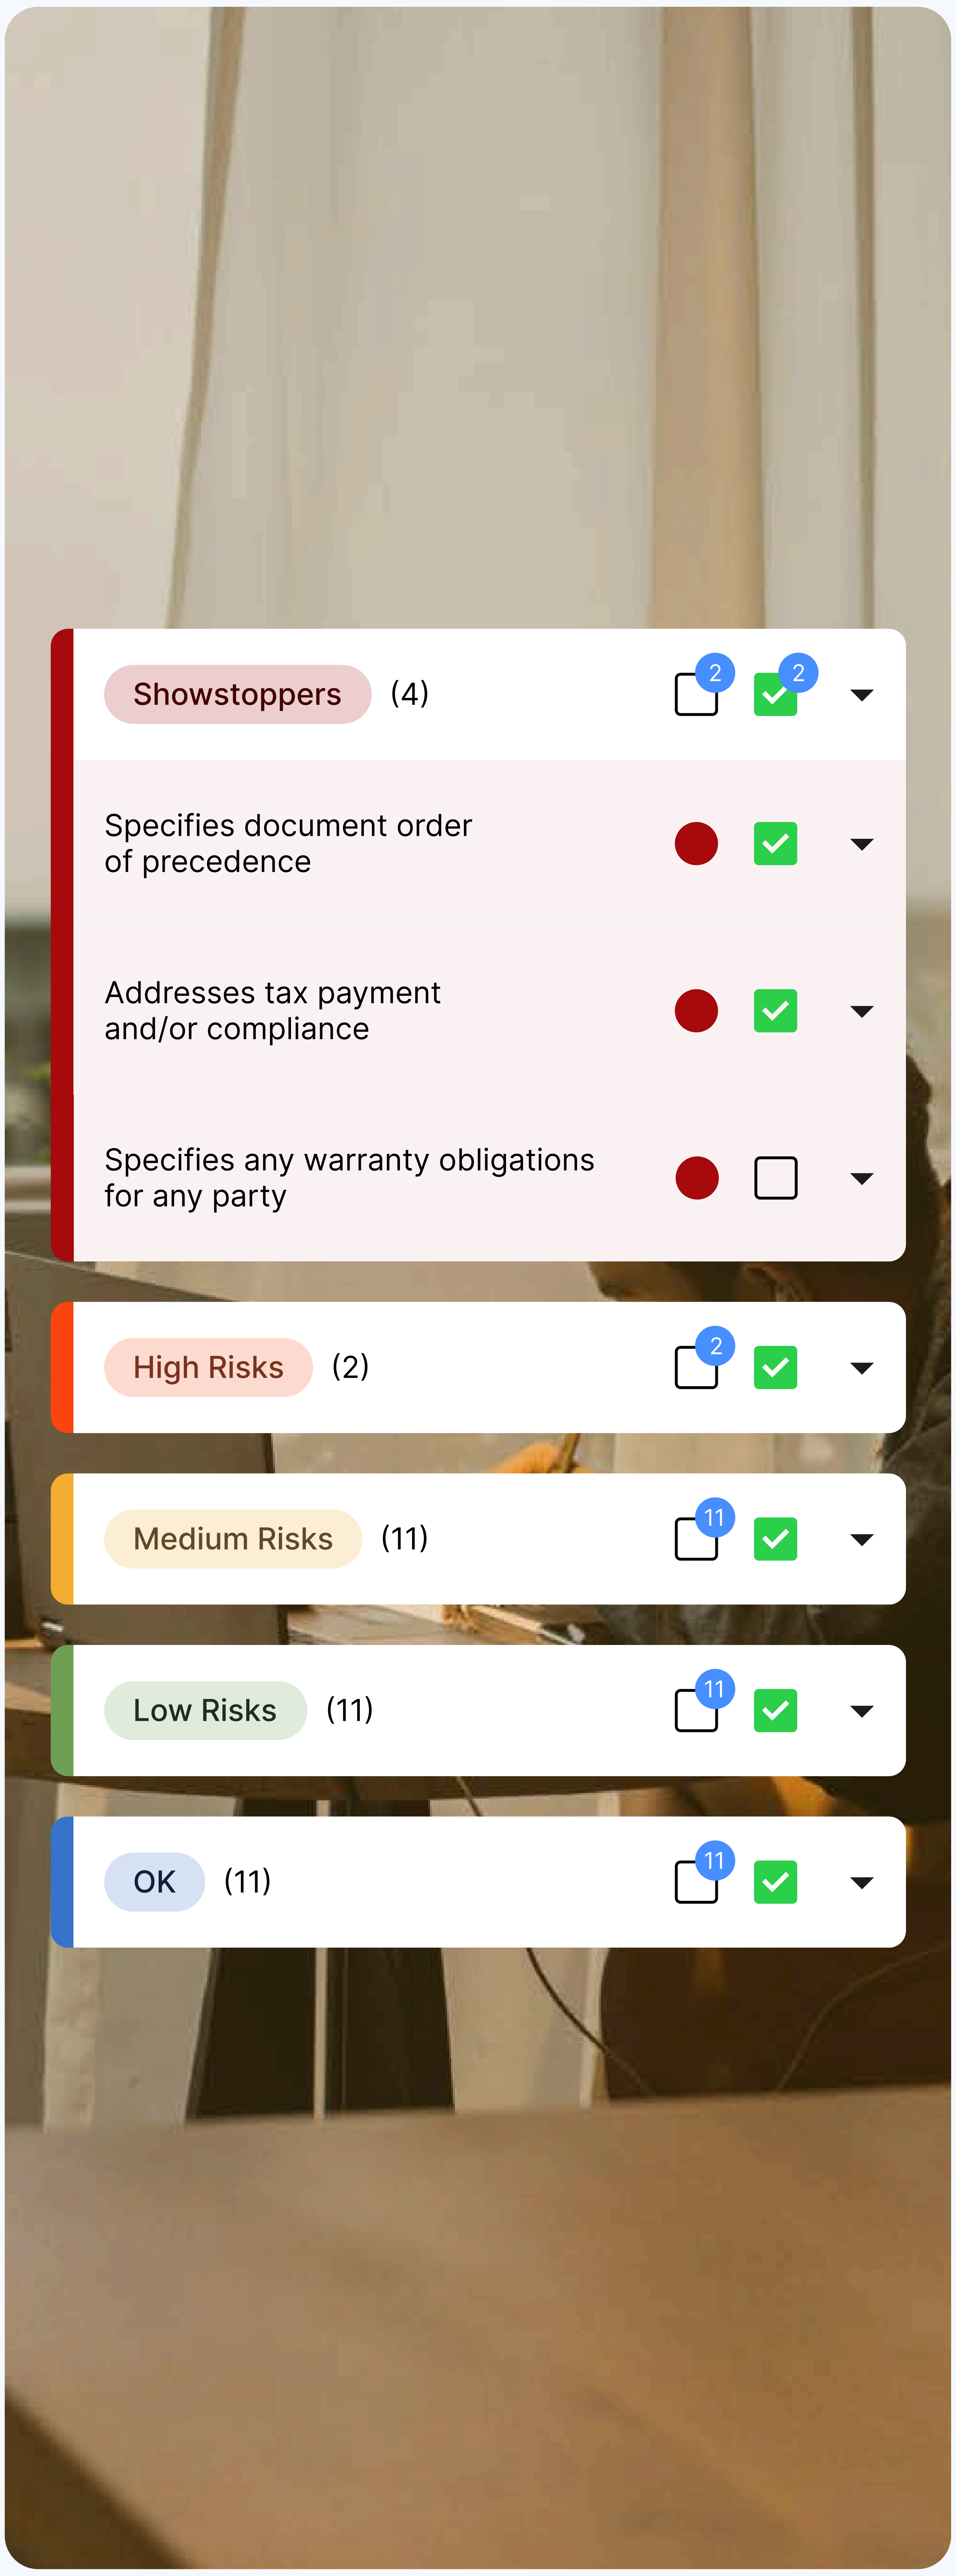

AI Clauses in Commercial Contracts: **Provisions Still Under Contention**

1. Accuracy Warranties

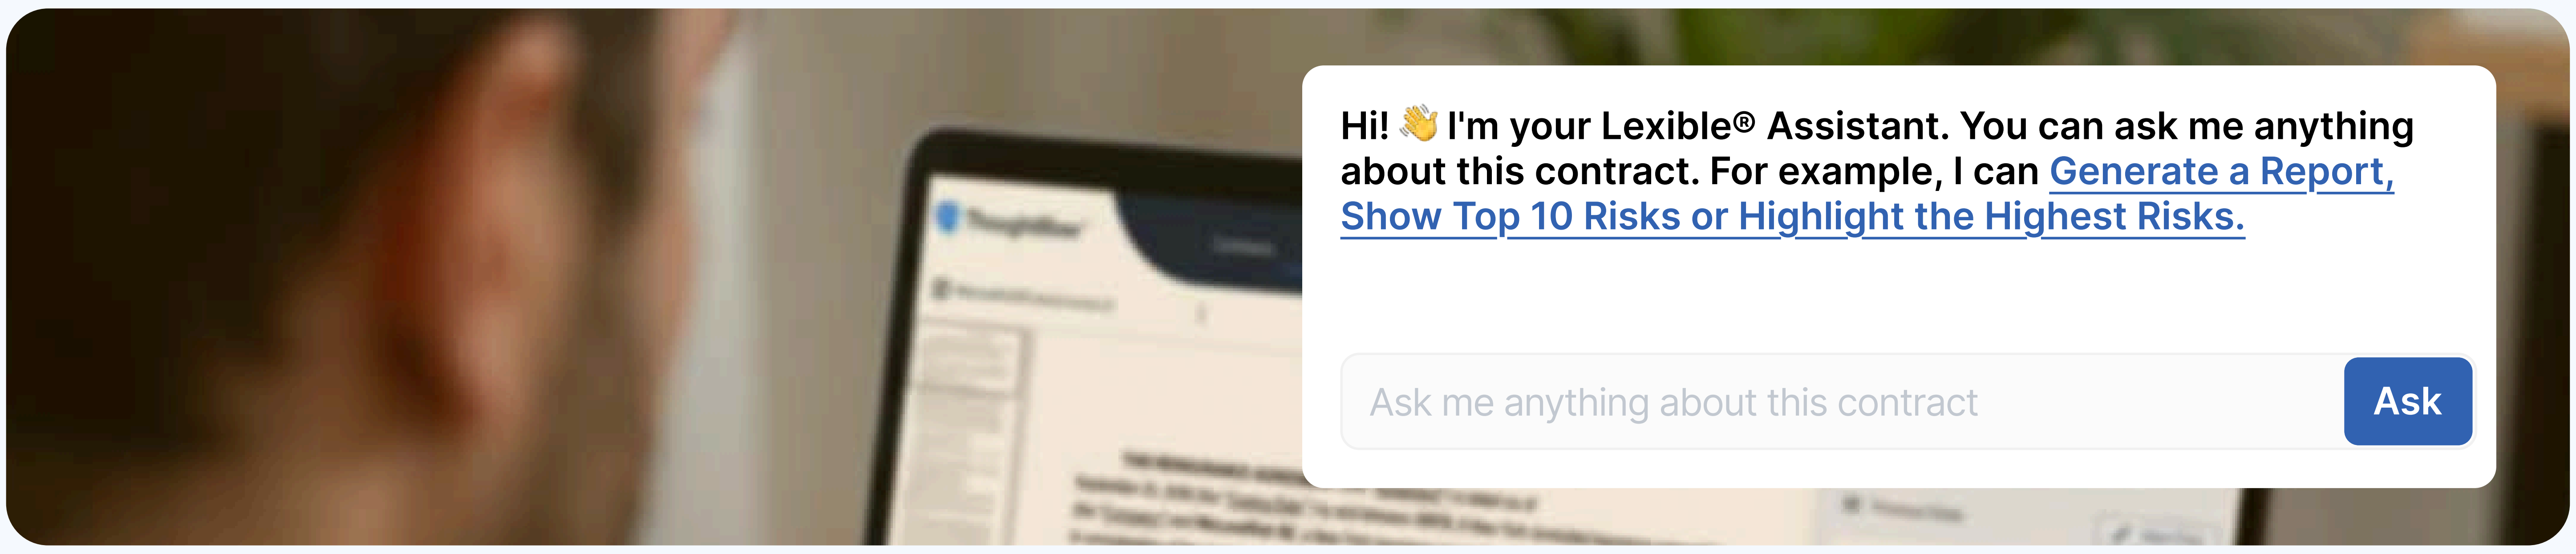

The flashpoint: Vendors typically refuse to warrant that AI outputs will be accurate, arguing that AI is inherently probabilistic. Buyers pushing back have a fair point: if you're selling a system based on its performance, that promise should mean something.

Vendors cannot always avoid taking responsibility for the performance of their AI systems. For vendors of AI systems classified as 'high-risk' under the EU AI Act, the law imposes specific obligations to design and maintain systems that achieve appropriate levels of accuracy, robustness and cybersecurity throughout their lifecycle.

Vendors are also required to define and document the expected level of accuracy for the system and implement measures to help ensure that performance remains consistent when the system is used as intended.

While these regulatory obligations do not amount to a contractual warranty in favour of a customer, they do demonstrate that, in certain circumstances, the law expects vendors to stand behind the performance of their AI systems.

What suppliers offer (vendor-friendly):

Supplier makes no warranty regarding the accuracy, completeness, or reliability of any AI-generated outputs. Customer acknowledges that AI systems may produce erroneous results and agrees to independently verify all outputs before relying on them.

What buyers want (buyer-friendly):

Supplier warrants that AI-generated outputs will meet the accuracy standards and performance metrics set forth in Exhibit A. If outputs fail to meet such standards, Supplier shall re-perform the work at no cost or refund fees paid.

Possible compromise:

Supplier warrants that AI systems will perform substantially in accordance with the specifications and performance benchmarks documented in Exhibit A, measured over a statistically significant sample size.

If performance falls materially below documented benchmarks (e.g., >10% deviation), Customer may request remediation or, if remediation is not achieved within 30 days, terminate the affected portion of the Services for refund.

This ties warranties to documented, measurable performance rather than absolute accuracy, which is more defensible for suppliers while still giving buyers recourse.

2. Liability for AI Hallucinations

Where a system generates convincing but completely false outputs and someone relies on them, who is responsible? Courts are starting to answer this question, and the answers aren't always what vendors hope.

For example, for vendors of AI systems classified as high-risk, hallucinations and other output errors should be appropriately mitigated. A failure to implement adequate safeguards may result in non-compliance with Article 15 (Accuracy, Robustness and Cybersecurity) and, where sufficient human oversight is not in place, Article 14 (Human Oversight) of the EU AI Act. Such non-compliance can expose vendors to significant regulatory fines.

In addition, the EU AI Act imposes transparency obligations on certain AI systems. In particular, users may need to be informed when they are interacting with AI, and AI-generated or AI-manipulated content may need to be clearly labelled, unless an exception applies, such as where the use of AI is obvious from the context, the AI performs a purely assistive function, or the content has been subject to meaningful human editorial control.

What suppliers want (vendor-friendly):

Supplier shall not be liable for any damages arising from Customer's or any third party's reliance on AI-generated outputs. Customer is solely responsible for verifying the accuracy of all outputs before use.

What buyers want (buyer-friendly):

Supplier shall be liable for all direct and consequential damages arising from materially inaccurate AI outputs, including but not limited to legal costs, regulatory fines, and reputational harm.

It's improbable for a Supplier to accept all liability given that the input may be one of many.

Possible compromise:

Supplier shall be liable for direct damages arising from AI outputs that are materially inaccurate due to Supplier's failure to implement reasonable quality controls or human oversight as specified in this Agreement, up to the liability cap set forth in Section X.

Supplier shall not be liable for consequential damages or for inaccuracies in outputs where Customer failed to perform the verification procedures specified in the Documentation.

This creates shared responsibility: suppliers are accountable for implementing reasonable safeguards, but buyers must also verify outputs in high-stakes situations.

Emerging case law to reference: In Air Canada's chatbot case (2024), the tribunal held that businesses cannot disclaim liability for what their AI systems tell customers. Point to this when pushing back on blanket disclaimers.

3. AI Training Data and Copyright Indemnification

Several major AI developers are facing litigation over whether their models were trained on copyrighted material without authorization. Sophisticated buyers are asking for indemnities covering not just use of outputs, but the provenance of training data itself.

What buyers want (buyer-friendly):

Supplier warrants that: (i) its AI models were trained only on data for which Supplier holds all necessary rights and licenses; and (ii) Customer's use of AI-generated outputs will not infringe any third-party intellectual property rights.

Supplier shall indemnify, defend, and hold Customer harmless from any claims, damages, or costs arising from breaches of these warranties, including claims related to the training data used to develop Supplier's AI models.

What suppliers resist: Providing indemnities for training data provenance when the legal landscape is unsettled and they may have relied on fair use or other defences.

Possible compromise:

Supplier represents that it has taken commercially reasonable steps to ensure its AI models were trained in compliance with applicable intellectual property laws and contractual obligations.

Supplier shall indemnify Customer for third-party claims arising from Customer's use of AI outputs, but such indemnity shall not extend to claims based solely on the composition of Supplier's training data, provided Supplier has complied with applicable law.

If a final court ruling determines that Supplier's training data violates third-party rights, Supplier shall at its option: (i) obtain necessary licences; (ii) modify the AI system to remove infringing elements; or (iii) refund fees paid.

This acknowledges legal uncertainty while still providing buyers with recourse if training data issues materialise.

Your Next Move

AI contract language is moving faster than most legal teams can comfortably track. What was a buyer's wish list twelve months ago is now baseline, and positions are still shifting as case law catches up.

A few principles hold across every clause covered in this guide:

- Define what you are talking about. Broad, undefined references to "AI" or "data" create the loopholes you will later regret. Make definitions as precise as permissible.
- Tie obligations to your standards. Reference measurable specifications that are truly important to your business. Knowing where to push and where to compromise is what separates a workable agreement from a stuck negotiation.
- Build in a path for change. Regulation, technology, best practice, and case law are all moving. Notice periods and review cycles should give you room to respond.
- Keep an open dialogue. Treat vendor negotiations like partnerships. Both parties are exploring new terrain. An open dialogue can help you evolve and get ahead of needs together.

Rapidly shifting baselines are now the norm. Clauses that feel hard-fought today will be standard within a year or two, and new flashpoints will take their place. Treat this guide as a working reference, not a final position.

See how ThoughtRiver can help

ThoughtRiver's AI-powered contract review surfaces risk, flags non-standard terms, redlines agreements, and aligns drafts in minutes.

Contact ThoughtRiver for a demo →

and see how we can empower your team with fast, consistent, and dependable contract review with enterprise-quality accuracy.