

MARKET INSIGHTS UNITED STATES

NEGOTIATING BESPOKE AI CONTRACT TERMS

Vivek Rao, Kyle Glover & Matthew Stein
Pierce Atwood

This chapter forms part of:

ARTIFICIAL INTELLIGENCE

Law Over Borders Comparative Guide 2026

www.globallegalpost.com/lawoverborders

Beyond boilerplate: Negotiating bespoke AI contract terms

The shift to more balanced, use-case-driven negotiation

If you're old enough to remember the early 2020s boom that triggered the global proliferation of new (or at least, newly widespread) AI technologies across virtually every economic sector, you'll likely also recall the first wave of contract terms governing the provision and use of these AI technologies.

AI contract terms during that first wave — call them AI Terms 1.0 — typically fell into two predictable, distinct categories.

In one corner, AI technology providers pushed aggressive boilerplate marked by:

- Broad licenses to use customer data for secondary purposes, often without meaningful limits.
- Blanket “AS IS”-type disclaimers.
- Narrowly drawn indemnities covering the relatively safe areas of copyright and trade secret infringement in the development and offering of the vendor's software code, but staying well clear of indemnifying the customer for infringement risks relating to training data and output data.

In the other corner, sophisticated AI procurers with meaningful leverage demanded the opposite:

- Robust representations and warranties on accuracy, non-discrimination, and compliance.
- Extensive model documentation, monitoring obligations, and audit rights.
- Uncapped indemnities covering AI hallucinations, errors, and infringement risks relating to training data and output data, not just software code.

AI Terms 1.0 can and do stick in unbalanced negotiations — for example, where a startup AI technology provider is desperate to land a large health insurer as a reputation-building early adopter, or where a mid-sized business lacks the leverage to negotiate effectively with a market-leading AI technology provider.

But negotiations between parties of similar sophistication and leverage have entered what we might call the “AI Terms 2.0” era. The defining characteristic is not that one side capitulates, but that both sides reckon with market realities requiring more thoughtful, use-case-driven negotiations and targeted, risk-informed concessions on issues critical to the counterparty.

On the procurement side, cross-disciplinary teams assess how AI technology works and how it will be deployed, and then risk-informed lawyers prioritize and negotiate (or deprioritize or drop altogether) specific clauses accordingly.

On the provider side, vendors determine which commitments they can credibly make, survey their target markets to understand customer needs, and offer meaningful assurances, while avoiding overreach inconsistent with their risk appetite and commercialization model.

In sum, AI boilerplate is giving way to thoughtful, tailored terms aligned with the technology's actual use and risks.

Pre-negotiation AI contract risk assessment

If AI Terms 2.0 is defined by nuance, then the prerequisite to effective negotiation is understanding the specific risks at stake for a given AI system.

Practitioners on both sides should start with a focused, pre-negotiation AI risk assessment. This exercise can draw on the same disciplines as internal AI impact assessments — aligned with emerging regulatory expectations and frameworks such as ISO 42001 and the NIST AI RMF — but here the objective is narrower: to identify and prioritize risks that should be specifically allocated, mitigated, or addressed via the contract.

While there is no universally applicable approach, topics to consider include:

Who could be harmed? Identify potential recipients of harm caused by use or misuse of the AI system — end users, data subjects, the deploying organization, third parties whose intellectual property (IP) may be implicated.

What is the nature of harm? Map potential harms: for example, privacy breach, discriminatory output, violation of applicable laws or inability to respond adequately to regulatory inquiries, IP infringement, financial loss, end-customer dissatisfaction, reputational damage, or physical harm in safety-critical applications.

Who can prevent the harm? Determine which party is best positioned to prevent or mitigate each harm. The vendor typically has superior visibility into the model's capabilities, limitations, and training data; the customer typically controls the deployment environment, how the AI output is used and relied upon, and how to manage notices to or questions from affected end customers; but numerous variations exist.

What operational measures are required or desirable for assessing, preventing, or mitigating the harm? Consider the practical safeguards, controls, and processes that are relevant and feasible to assess, prevent, and mitigate the identified harms. These may include, for example, technical measures, such as human-in-the-loop review during model development and fine tuning, output filtering, and bias testing; administrative measures such as usage policies, training programs, and escalation protocols; and contractual mechanisms such as audit rights, performance metrics, and remediation obligations. Which controls make sense will vary significantly depending on the use case: a customer service chatbot may warrant different controls than a medical diagnostic tool or a financial underwriting system, and a low-risk AI system may warrant few, if any, of these measures.

The answers to all of these questions should directly shape each party's negotiation strategy — distinguishing the terms that warrant firm advocacy because they address material, unmitigated risks from those that can be narrowed or conceded because the underlying risk is minimal, already addressed through operational controls, or more effectively managed internally rather than shifted contractually to the counterparty.

The goal is not to achieve theoretically ideal protections across every conceivable issue, but to secure meaningful safeguards where they matter most and to make reasonable compromises where insisting on perfect terms would stall negotiations without meaningfully reducing risk.

Illustrative examples of AI Terms 2.0

When both parties let go of AI Terms 1.0 and engage in risk-informed negotiations, productive compromises can emerge. Below is a non-exhaustive list of key topics in AI contracts, how aggressive vendors and customers typically address them, and how nuanced approaches can bridge the gap.

Ownership of IP:

- Vendors typically seek to retain ownership of the underlying AI model and all improvements, enhancements, and derivative works arising from the customer's use, including fine-tuned models and custom configurations.
- Customers with leverage demand ownership, or at minimum a perpetual royalty-free license, over custom-trained models, fine-tuned layers, and outputs generated using their proprietary data.
- A pragmatic middle ground might distinguish between the vendor's pre-existing IP, which remains with the vendor, and customer-specific fine-tuning, configurations, and outputs, over which the customer receives broad perpetual license rights or joint ownership with delineated fields of exclusive use.

Secondary data use rights:

- Vendors typically seek sweeping secondary use rights with respect to customer inputs and outputs.
- Customers, thoughtfully or reflexively, balk at such terms.
- In more nuanced negotiations, where securing secondary data use rights is important to the vendor's commercialization model and cannot simply be dropped, the vendor should explain how the customer may also ultimately benefit from granting such rights. The vendor can then seek, and the customer is more likely to give, a tailored authorization that defines the customer data approved for secondary use, the specific purposes for which such use is authorized, and any additional conditions, which may include commitments not to disclose the relevant customer data to third parties and to take responsibility for processing and protecting such data in compliance with applicable laws (including, where applicable, the vendor confirming it is the "data controller" with respect to such processing).

Model documentation:

- Vendor standard terms typically offer little more than high-level marketing documentation, user guides, and application programming interface (API) references.
- Sophisticated enterprise customers, particularly in regulated sectors, often demand comprehensive model documentation packages.
- A negotiated compromise may involve the vendor delivering initial documentation covering training data categories (without disclosing the underlying datasets themselves), summary-level architecture descriptions, bias testing results relevant to the customer's use case, and key performance metrics, with ongoing obligations to update such documentation upon material model changes and to provide supplementary information reasonably requested by the customer's compliance or risk functions under appropriate confidentiality protections.

AI output quality, accuracy, and fairness:

- Vendors prefer to disclaim warranties regarding accuracy, completeness, or fitness for purpose, positioning AI as a probabilistic tool requiring independent validation.

- Customers insist on express warranties that the model will be fit for purpose and delivered in a professional and diligent manner (sometimes blending product liability and service liability theories), meet documented accuracy thresholds, produce non-discriminatory outputs, and be suitable for their use case.
- A nuanced approach might establish mutually agreed performance metrics tied to the specific use case, with the vendor warranting benchmarks at delivery or during acceptance testing, while the customer, depending on the use case, accepts responsibility for validating outputs before making consequential decisions.

Regulatory compliance:

- Vendors prefer to allocate compliance risk to customers, only warranting compliance with laws applicable to their own operations.
- Customers seek broad assurances that their use of AI outputs will comply with all applicable laws.
- The appropriate allocation depends on the technology: general-purpose AI typically justifies placing compliance risk on the customer, while narrow AI for regulated sectors (e.g., insurance underwriting, resume screening) typically justifies more shared responsibility. A balanced approach has the vendor warrant compliance with laws applicable to technology providers and commit to reasonable cooperation with the customer's compliance efforts, while the customer retains primary responsibility for its regulatory obligations. The parties should also go beyond the surface concept of "responsibility for compliance" and dig deeper into the specific operational activities necessary to achieve compliance, with accompanying contract terms aligned to a shared understanding of which party is best positioned to carry out each of those activities. And the parties should establish a framework for monitoring regulatory changes, with notice obligations and rights to modify or suspend service to maintain compliance.

Infringement risk:

- Vendors typically limit IP indemnification to claims that their proprietary code infringes third-party copyrights or trade secrets, while carving out infringement arising from training data, customer inputs, or AI-generated outputs – precisely the most novel and likely-to-be-litigated risks with AI.
- Customers push for broad indemnification covering all IP claims, including training data provenance and output similarity to copyrighted works.
- A middle ground may involve the vendor providing its standard code-level indemnity plus representations that it has implemented reasonable procedures to vet training data for IP compliance, with a limited indemnity or shared-risk framework for training data claims. Output-related risk might be addressed through a combination of the vendor's obligation to implement guardrails and filters, the customer's obligation to screen outputs before external use, and a negotiated allocation of liability for residual claims that survive both parties' reasonable precautions. And all of this should align with the vendor's core sales pitch (or, where applicable, core disclaimer) to customers on how it trains its models and the purposes for which customers can expect to use the output.

Disclaimers and liability caps:

- Vendors favor broad “AS IS” disclaimers, low liability caps (often 12 months’ fees), and blanket consequential damages exclusions.
- Customers seek narrower disclaimers, express warranties, higher caps for AI-specific harms (data breaches, regulatory penalties, IP infringement), and carve-outs for foreseeable losses.
- A negotiated outcome might preserve a general liability limitation at a reasonable multiple of annual fees, with super-caps or uncapped carve-outs for high-risk categories (confidentiality breaches, willful misconduct, IP indemnification, data protection), while the vendor provides targeted warranties tied to agreed performance standards. The parties might also separate categories of liability that a traditional SaaS framework might treat as monolithic (e.g., infringement liability) into subcategories meriting different treatment.

Model drift, monitoring, audit, and change management:

- Vendors prefer full discretion over model updates with no notification, monitoring, or audit obligations.
- Customers demand ongoing monitoring commitments, regular performance measurement and reporting, drift and bias detection and reporting, broad audit rights, and formal change management requiring advance notice and consent before material updates.
- A practical compromise might involve the vendor monitoring against agreed metrics and providing periodic performance reports, implementing change management with advance notice and a testing window before production deployment, and granting audit rights through an independent third-party auditor under confidentiality protections, scoped to contractual compliance rather than wholesale inspection of proprietary technology.

Conclusion

AI contracting has moved from an all-or-nothing phase to a more pragmatic, risk-informed approach. Both vendors and customers now enter negotiations with targeted objectives: vendors clarify what risks they can shoulder and prepare to offer limited but meaningful warranties or indemnities, while customers define the safeguards they need for their data and use case. By focusing on the actual AI application and negotiating tailored terms, parties can unlock AI’s benefits while keeping risks in check.

For additional reading on addressing AI risks through contracting, see ‘Machines Make Mistakes Too: Planning for AI Liability in Contracting’ (<https://scholarlycommons.law.case.edu/jolti/vol15/iss2/5/>) by Pierce Atwood attorneys Mark Sayre and Kyle Glover, published by Case Western Reserve Journal of Law, Technology & the Internet (volume 15, issue 2).

Continued on the following page.

AUTHOR BIOGRAPHIES



Vivek Rao

Vivek Rao leads Pierce Atwood's Privacy and Data Security practice and is a member of the firm's Technology Transactions and Outsourcing practice. Vivek emphasizes client-tailored, practical approaches to helping businesses maximize the tremendous commercial opportunities and evolving legal challenges of the technology economy. His clients range from Fortune 500 companies in highly regulated industries to family-owned businesses, to sophisticated startups commercializing innovative software and data-driven services. Vivek focuses his practice on technology transactions, data privacy compliance, consumer-facing intellectual property terms, brand protection, marketing management, and technology mergers and acquisitions risk counseling.



Kyle Glover

Kyle Glover leads Pierce Atwood's Technology Transactions and Outsourcing practice. Kyle assists clients with a wide range of matters related to the negotiation of complex technology and intellectual property transactions, the prosecution and protection of trademarks, and compliance with consumer protection, data privacy, and security requirements. In addition to his subject-matter and legal expertise, Kyle advises clients on negotiation strategy throughout the deal and settlement life cycle, including advising on negotiation set-up and structure, process, and substantive contract issues. He also regularly represents clients in negotiations with counterparties to the deal or settlement, where he employs cutting-edge negotiation techniques and best practices to resolve conflict, clarify concepts, and create and claim value for the client.



Matthew Stein

Matt Stein leads Pierce Atwood's Business Group and is a member of the firm's Technology Transactions and Outsourcing practice. From startups to Fortune 500 companies, Matt Stein helps clients navigate, negotiate, and close technology and intellectual property transactions involving software, Software as a Service (SaaS), AI and models, data processing, managed services, and cloud computing. Matt understands the real-world challenges and opportunities presented by these emerging technologies and provides clients with thoughtful, pragmatic advice to get complex deals done efficiently and responsibly. Matt focuses his practice on negotiating all manner of commercial contracts, advising clients with respect to privacy and information security related matters, AI risk management and compliance, and managing, defending, licensing, and enforcing domestic and international trademark portfolios.

FIRM PROFILE

Pierce Atwood

Pierce Atwood, a highly regarded full-service law firm based in New England, has 150 attorneys who serve regional, national, and international clients. Pierce Atwood provides legal and business counsel and advocacy services to its clients, which range from regional and local enterprises, utilities, financiers, middle-market companies, entrepreneurs, and individuals to Fortune 500 companies, multinational corporations, and state and foreign governments.