



CONTRACTUAL SILENCE AND GAP-FILLING IN AI-GENERATED AGREEMENTS: WHO BEARS THE RISK?

RAGHUNATHAN Priya¹

National Law School of India University, Bangalore, India

Email: praghunathan@gmail.com

FISCHER Lukas²

Bucerius Law School, Hamburg, Germany

Email: lukasfishaa@gmail.com

ABSTRACT

Large language models are increasingly used not merely to review contracts but to draft and negotiate them, generating terms with no identifiable human author for a given clause. This article asks how the common law's gap-filling doctrines — implied terms, contra proferentem, and mutual mistake — should adapt when the “drafter” whose intent traditionally anchors these doctrines is an artificial system rather than a party or its lawyer. We argue that existing doctrine can accommodate AI-generated ambiguity without wholesale reform, but only if courts treat the deploying party, rather than the AI system itself, as the relevant drafter for purposes of contra proferentem, while treating genuine AI-generated inconsistency between the parties' understandings as a species of mutual mistake rather than unilateral mistake. We further argue, drawing on efficient risk allocation principles from law and economics, that liability for AI-drafting errors should presumptively fall on the party best positioned to verify the output — typically the party who selected and deployed the tool — subject to contractual reallocation. The article closes with a proposed model clause allocating AI-drafting risk, informed by comparative practice in Indian, German, and English contract law.

Introduction

Contract drafting has historically been an act attributable to an identifiable human agent: a solicitor, an in-house counsel, a contracting party acting pro se. Doctrines that resolve ambiguity — most prominently contra proferentem, which construes ambiguous terms against the party who drafted them — depend on this attribution.¹ The increasing use of generative AI tools to draft entire agreements, or discrete clauses within otherwise human-drafted agreements, disrupts this premise. Where a clause is generated by a large language model prompted by one party's counsel, refined through several rounds of AI-assisted

¹ Kevin M Teeven, ‘A History of the Anglo-American Common Law of Contract’ (Greenwood Press 1990) 234–36.

negotiation, and finally accepted without material human edit by both sides, it becomes genuinely unclear who “drafted” the resulting ambiguity for doctrinal purposes.

This article proceeds in five parts. Part II describes the contractual gap-filling doctrines most likely to be implicated by AI-assisted drafting: implied terms, *contra proferentem*, and mistake. Part III develops the central doctrinal argument — that attribution of AI-generated text should track deployment and reliance rather than authorship in any literal sense. Part IV applies a law-and-economics lens to ask which party is best placed, at lowest cost, to verify AI output, and argues that this economic analysis should inform default risk allocation in the absence of express contractual terms. Part V offers a comparative view across Indian, German, and English contract law, noting convergence on outcome despite doctrinal divergence in method. Part VI proposes a model clause and concludes.

II. Gap-Filling Doctrines and Their Premises

Implied terms. Courts imply terms into contracts either as a matter of law (default rules applicable to a class of contract) or as a matter of fact (terms necessary to give business efficacy to the particular agreement, per the classic formulation in *The Moorcock*).² The business efficacy test asks what the parties would have agreed had they turned their minds to the gap in question — a counterfactual inquiry that presupposes parties capable of turning their minds to anything at all. Where the gap arises because an AI drafting tool omitted a term neither party’s human negotiators ever considered, the counterfactual inquiry becomes strained: the “parties” in the relevant sense never approached the question, and the tool that did approach it (in the loose sense of generating text) has no intent to impute.

Contra proferentem. The rule construing ambiguity against the drafter rests on an incentive rationale: a party controlling the language of a term is best placed to draft clearly, and should bear the cost of failing to do so.³ English courts have narrowed the doctrine’s application in commercial contracts between parties of

² *The Moorcock* (1889) 14 PD 64 (CA) 68 (Bowen LJ).

³ Catherine Mitchell, *Interpretation of Contracts* (2nd edn, Routledge 2019) 112–15.

equal bargaining power, treating it as a rule of last resort after purposive construction has been exhausted.⁴ Its incentive rationale nonetheless remains relevant to AI-assisted drafting: whichever party selected, prompted, and deployed the AI tool controlled the process that generated the language, even if no human at that party read every generated word before signature.

Mistake. Common mistake voids a contract where both parties share a fundamental misapprehension about the subject matter, per *Bell v Lever Brothers* and its stricter modern application in *Great Peace Shipping*.⁵ Unilateral mistake, by contrast, generally does not vitiate consent unless the non-mistaken party knew or ought to have known of the other's error.⁶ AI-generated ambiguity sits awkwardly between these categories: it is not obviously "shared" in the sense the mistake doctrine contemplates, since neither party may have formed any belief about the disputed term at all, and it is not obviously unilateral either, since neither party's error is more culpable than the other's in any traditional sense.

III. Attribution Without Authorship: Reframing the Drafter Inquiry

We propose that courts confronting AI-generated ambiguity should ask not "which party's AI system generated this specific string of text" but "which party selected, deployed, and had the opportunity to review the tool's output before it became part of the operative agreement." This reframing tracks the underlying incentive rationale of *contra proferentem* more faithfully than a literal authorship inquiry would, for three reasons.

First, a literal authorship inquiry is frequently indeterminate in practice. Modern contract-negotiation workflows often involve iterative exchange of AI-modified drafts, such that by the time of signature neither party's counsel can reliably reconstruct which party's tool introduced a given clause.⁷ A doctrine that depends on resolving this question will simply fail to resolve most disputes, generating high-variance litigation outcomes turning on discovery of prompt logs and model version histories.

⁴ *K/S Victoria Street v House of Fraser (Stores Management) Ltd* [2011] EWCA Civ 904, [2012] Ch 497 [68] (Rimer LJ).

⁵ *Bell v Lever Brothers Ltd* [1932] AC 161 (HL); *Great Peace Shipping Ltd v Tsavliris Salvage (International) Ltd* [2002] EWCA Civ 1407, [2003] QB 679.

⁶ *Smith v Hughes* (1871) LR 6 QB 597; *Hartog v Colin & Shields* [1939] 3 All ER 566 (KB).

⁷ See generally Karen Yeung and Martin Lodge (eds), *Algorithmic Regulation* (Oxford University Press 2019) ch 6.

Second, the deployment-based inquiry preserves the doctrine's incentive structure. If liability for ambiguity tracks deployment rather than literal generation, parties are incentivised to review AI output before relying on it — precisely the behaviour *contra proferentem* was designed to encourage among human drafters who might otherwise produce careless boilerplate.⁸ A party that deploys an AI tool and signs the resulting agreement without meaningful review has, in every functionally relevant sense, adopted the tool's output as its own drafting, whatever the tool's internal architecture.

Third, this approach avoids awarding either party a strategic incentive to obscure their use of AI tools, an incentive a literal-authorship rule would create by rewarding parties who can plausibly claim the other side's tool “really” generated the disputed term.

Where both parties jointly used a shared AI negotiation platform — an increasingly common arrangement in cross-border commercial transactions — the deployment-based inquiry collapses, since both parties equally selected and had opportunity to review the tool. We argue that in this joint-deployment scenario, resulting ambiguity should default to the mutual mistake framework discussed below, rather than *contra proferentem*, since no single party can be said to have controlled the drafting process in the sense the doctrine presupposes.

IV. Efficient Risk Allocation: A Law and Economics Perspective

Where deployment cannot be cleanly attributed to a single party, or where the parties have contracted around *contra proferentem* (as sophisticated commercial parties frequently do through boilerplate disclaiming its application⁹), a residual question remains: who should bear the cost of AI-drafting errors as a matter of default rule?

The Coasean starting point is that, in a world without transaction costs, the initial allocation of liability would not affect efficiency, since parties would simply bargain to the efficient outcome regardless of the default rule.¹⁰ Contract drafting is not costless, however, and default rules matter precisely because

⁸ John D McCamus, *The Law of Contracts* (3rd edn, Irwin Law 2020) 787–90.

⁹ Edwin Peel, *Treitel on the Law of Contract* (15th edn, Sweet & Maxwell 2020) para 7-016.

¹⁰ RH Coase, ‘The Problem of Social Cost’ (1960) 3 *Journal of Law and Economics* 1, 15–19.

renegotiation after signature is expensive or, once a dispute has crystallised, impossible. The relevant question therefore becomes which party is the cheapest cost avoider — the party who could have prevented the error at the lowest cost.¹¹

We argue that, absent contrary indication, the party who selected the AI tool and had access to its output before execution is typically the cheapest cost avoider, for the simple reason that they alone had the opportunity to review the specific text before signature. This is not an argument that AI-drafting risk should always fall on the deploying party regardless of circumstance — where the counterparty possessed superior technical knowledge of the tool’s known failure modes and withheld that knowledge, ordinary principles of misrepresentation or unconscionability should apply.¹² It is instead an argument for a sensible default: absent express allocation, the deploying party bears the risk of its own tool’s errors, and both parties bear the risk of a jointly deployed tool’s errors under the mutual mistake framework.

This default has the further virtue of creating a market incentive for AI contract-drafting vendors to compete on output verifiability — for example, by providing confidence scores or clause-level provenance logs — since deploying parties bearing default liability will demand such features from their vendors.¹³

V. A Comparative View: India, Germany, England

India. The Indian Contract Act, 1872 codifies mistake and consent doctrines in sections 20 to 22, with section 20 requiring that both parties be mistaken as to a matter of fact essential to the agreement for the contract to be void.¹⁴ Indian courts have applied a functional, substance-over-form approach to contractual interpretation that would likely accommodate the deployment-based attribution proposed in Part III without requiring legislative amendment, particularly given the flexibility Indian courts have

¹¹ Guido Calabresi, *The Costs of Accidents: A Legal and Economic Analysis* (Yale University Press 1970) 135–40, adapting the cheapest-cost-avoider concept from tort to contractual risk allocation contexts.

¹² Cf *Derry v Peek* (1889) 14 App Cas 337 (HL) on the scienter requirement for deceit, and its relationship to the lower threshold for negligent misrepresentation under *Hedley Byrne & Co Ltd v Heller & Partners Ltd* [1964] AC 465 (HL).

¹³ Cf discussion of algorithmic auditing requirements in Sandra Wachter, Brent Mittelstadt and Chris Russell, ‘Counterfactual Explanations Without Opening the Black Box’ (2018) 31 *Harvard Journal of Law & Technology* 841, 862–65.

¹⁴ Indian Contract Act 1872, s 20.

shown in applying the business efficacy test under section 29 read with established common law principles.¹⁵

Germany. German law approaches the problem through the doctrine of *Willenserklärung* (declaration of intent) under § 133 and § 157 BGB, which require interpretation according to the true intent of the parties rather than the literal meaning of a declaration, tempered by the requirements of good faith (*Treu und Glauben*).¹⁶ Where an AI-generated term diverges from what a party actually intended to declare, German doctrine offers a route to rectification through § 119 BGB's provisions on error in declaration (*Erklärungsirrtum*), potentially avoiding the deployment-attribution question altogether by focusing on subjective intent rather than drafting responsibility.¹⁷ This represents a genuine point of doctrinal divergence from the English and Indian common law approach, though the practical outcomes may converge where courts apply the good-faith requirement to allocate the cost of an AI tool's error to the party who deployed it without adequate verification.

England. As set out in Part II, English courts have narrowed contra proferentem's application among commercial parties, preferring purposive interpretation under the framework established in *Investors Compensation Scheme v West Bromwich Building Society* and refined in *Wood v Capita Insurance Services*.¹⁸ This purposive approach, asking what a reasonable person with the parties' background knowledge would have understood the term to mean, may in practice absorb much of the deployment-based inquiry proposed here, since a court asking what a reasonable commercial party would have understood is implicitly asking which party's process generated the relevant expectation.

The comparative picture suggests convergence on substance — cost typically falls on the party best positioned to have caught the error — even where the doctrinal routes differ substantially between codified civil law and purposive common law interpretation.

¹⁵ Avtar Singh, *Contract and Specific Relief* (12th edn, Eastern Book Company 2017) 201–06.

¹⁶ Bürgerliches Gesetzbuch (BGB) §§ 133, 157.

¹⁷ BGB § 119(1); Dieter Medicus and Jens Petersen, *Bürgerliches Recht* (27th edn, Vahlen 2019) paras 143–48.

¹⁸ *Investors Compensation Scheme Ltd v West Bromwich Building Society* [1998] 1 WLR 896 (HL) 912–13 (Lord Hoffmann); *Wood v Capita Insurance Services Ltd* [2017] UKSC 24, [2017] AC 1173 [10]–[15] (Lord Hodge).

VI. A Model Clause and Conclusion

Sophisticated commercial parties need not wait for doctrine to settle before addressing AI-drafting risk expressly. We propose the following model clause, offered here for illustrative rather than jurisdiction-specific drafting purposes:

“Where any provision of this Agreement was generated, in whole or material part, by an artificial intelligence drafting tool selected or deployed by a Party, that Party shall be treated as the drafter of such provision for the purposes of any rule of construction resolving ambiguity against the drafter, unless the tool was jointly selected and deployed by both Parties, in which case any resulting ambiguity shall be resolved by reference to the Parties’ documented pre-contractual negotiations.”

This clause does not purport to resolve every dispute an AI-drafted ambiguity might generate, but it removes the threshold indeterminacy identified in Part III by fixing attribution to deployment rather than requiring reconstruction of a generative process that is frequently undiscoverable after the fact.

The broader argument of this article is that existing contract doctrine is more adaptable to AI-generated ambiguity than the novelty of the technology might suggest, provided courts are willing to reframe the drafter inquiry around deployment and review opportunity rather than literal authorship. This reframing preserves the incentive structures underlying contra proferentem and mistake doctrine while avoiding the practical indeterminacy that a literal-authorship approach would generate. Legislative intervention may eventually prove necessary as AI-negotiation platforms become more deeply embedded in commercial practice, but the doctrinal resources to manage the interim period already exist within the common law and its civilian counterparts alike.