Texas Law Review Online
Volume 104 | 2026

Shadow Derivatives: The Quiet Propertization of AI Learning

*Mark C. McCabe**

Artificial intelligence systems generate durable competitive advantage not through any single output but through the learning that accumulates across training, fine-tuning, and deployment. Model weights, embeddings, and internal parameter updates encode this learning, yet they fall outside the scope of copyright's derivative works doctrine because they are functional, not expressive. Copyright law's exclusion of systems and processes from protection is not an oversight; it is a deliberate boundary. But that boundary has not prevented the allocation of exclusive control over AI learning. It has merely shifted the site of allocation from statute to contract.

This Essay identifies and names a phenomenon it calls contractual derivatives*: learning-based assets that are not statutory derivative works but are rendered proprietary through private ordering. Through service framing, narrow output definitions, residual rights clauses, improvement provisions, and use restrictions, AI vendors have constructed contractual architectures that replicate the exclusionary effects of intellectual property without invoking its doctrines. The Essay argues that these arrangements raise concerns at the intersection of copyright preemption, consent theory, and competition*

* Former Associate General Counsel & Chief IP Counsel, Genpact Ltd. (2017–2025). J.D., Temple University Beasley School of Law; B.A., Computer & Information Systems, Stockton University. Registered Patent Attorney, U.S. Patent and Trademark Office. The author spent twenty years negotiating enterprise technology and intellectual property transactions at Genpact, Accenture, and IBM, including AI-enabled SaaS, cloud infrastructure, and large-scale outsourcing agreements. The observations in Part I are informed by that experience. The author thanks Emma L. McCabe, whose careful reading of this Essay and ongoing collaboration on related work prompted the addition of the *Bartz v. Anthropic* discussion in Part I.

policy—particularly where they allocate control over subject matter that public law has deliberately left unowned.

The Essay proposes a taxonomy distinguishing statutory, contractual, and functional derivatives, and offers a three-part framework courts can apply to determine when contractual allocation of AI learning warrants heightened scrutiny rather than routine enforcement.

Table of Contents

Introduction

Artificial intelligence (AI) systems learn. In today's AI markets, durable advantage comes less from any single output than from the learning that accumulates through training, fine-tuning, and downstream feedback loops.[1] Each interaction, correction,

1. *See* Mark A. Lemley & Bryan Casey, *Fair Learning*, 99 TEXAS L. REV. 743, 745, 753 (2021) (noting that the "vast potential" of machine learning is tied to an "appetite for data" and that unlocking this potential requires "successive iterations" through an active learning process); Mkt. Structure and Antitrust Subcomm., *Report*, *in* STIGLER COMM. ON DIGITAL PLATFORMS, STIGLER CTR. FOR STUDY OF ECON. AND STATE, FINAL REPORT 23, 40–41 (2019) [hereinafter STIGLER CTR. REPORT], https://publicknowledge.org/wp-content/uploads/2021/11/Stigler-Committee-on-Digital-Platforms-Final-Report.pdf [https://perma.cc/VN7B-7MMF] (describing a "virtuous

and deployment contributes incrementally to improved performance, enabling systems to generalize, adapt, and optimize over time.[2] Yet the law has struggled to locate this learning within familiar intellectual property categories. Model weights, embeddings, and internal parameter updates are neither expressive works nor discrete inventions.[3] They sit uneasily between doctrinal regimes designed for books and machines, not for continuously adapting systems.

Copyright law's response to this problem is not ambiguity, but refusal. For more than a century, copyright has drawn a firm boundary between expression and function.[4] Systems, processes, and methods of operation are excluded from protection even when they are economically valuable.[5] That exclusion is not a gap to be filled. It is a policy choice designed to preserve competition, enable follow-on innovation, and prevent the

---

circle" where firms apply machine learning to accumulated data to improve products and entrench market position).

2. *See* INT'L ORG. SEC. COMM'N, THE USE OF ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING BY MARKET INTERMEDIARIES AND ASSET MANAGERS 10 (2019), https://www.iosco.org/library/pubdocs/pdf/IOSCOPD684.pdf [https://perma.cc/MMG8-DDRX] ("Unlike traditional algorithms, [machine learning] algorithms continually learn and develop over time."); *see also* Michael Chen, *What Is AI Model Training & Why Is It Important?*, ORACLE (Dec. 6, 2023), https://www.oracle.com/artificial-intelligence/ai-model-training/ [https://perma.cc/5S5H-V2HE].

3. *See* Alan Z. Rozenshtein, *There Is No General First Amendment Right to Distribute Machine-Learning Model Weights*, LAWFARE (Apr. 4, 2024, at 08:02 CT), https://www.lawfaremedia.org/article/there-is-no-general-first-amendment-right-to-distribute-machine-learning-model-weights [https://perma.cc/3HSX-54FA] ("Unlike source code, which humans use to express ideas to each other, model weights function solely as machine-readable instructions.").

4. *See* Baker v. Selden, 101 U.S. 99, 102–04 (1879).

5. *See* 17 U.S.C. § 102(b); Feist Publ'ns, Inc. v. Rural Tel. Serv. Co., 499 U.S. 340, 349–50 (1991).

privatization of functional knowledge.[6] Patent law, while capable of protecting some functional advances, is ill-suited to govern opaque, cumulative learning processes that arise through use rather than discrete inventive acts. The result is a category of value that public law deliberately leaves unowned.

This doctrinal settlement has not prevented the allocation of learning value in practice. Instead, it has shifted the site of allocation. AI vendors have increasingly turned to contract to govern learning indirectly, without asserting ownership over learning artifacts as intellectual property. Through standard-form service agreements, vendors define customer-owned outputs narrowly, treat learning as an internal system function, reserve residual rights in unallocated artifacts, and impose restrictions on reuse and competitive deployment.[7] These agreements allocate exclusive control over learning *ex ante*, before learning occurs, and do

6. *Feist*, 499 U.S. at 349–50; *see also* H.R. REP. NO. 94-1476, at 57 (1976) ("Section 102(b) is intended, among other things, to make clear that the expression adopted by the programmer is the copyrightable element in a computer program, and that the actual processes or methods embodied in the program are not within the scope of the copyright law.")

7. *See, e.g.*, *OpenAI Services Agreement,* OPENAI (Dec. 1, 2025), https://openai.com/policies/services-agreement/ [https://perma.cc/M8VV-XBK8] [hereinafter *OpenAI Services Agreement*]; *Terms of Use,* OPENAI (Jan. 1, 2026), https://openai.com/policies/row-terms-of-use/ [https://perma.cc/GQ96-74YC]; *Service Specific Terms,* GOOGLE CLOUD (Jan. 29, 2026), https://cloud.google.com/terms/service-terms [https://perma.cc/2RUG-T78D]; *Supplemental Terms of Use for Microsoft Azure Previews,* MICROSOFT AZURE (Nov. 2025), https://azure.microsoft.com/en-us/support/legal/preview-supplemental-terms [https://perma.cc/N3L6-DA2Z] [hereinafter *Azure Preview Supplemental Terms*]; *Anthropic on Vertex Commercial Terms of Service,* ANTHROPIC (Mar. 19, 2024), https://www-cdn.anthropic.com/471bd07290603ee509a5ea0d5ccf131ea5897232/anthropic-vertex-commercial-terms-march-2024.pdf [https://perma.cc/7L9E-NQF2].

so without confronting the doctrinal limits that copyright and patent law impose.

This Essay argues that these contractual mechanisms create shadow derivatives: assets that function like intellectual property but arise through private ordering rather than public law. The Essay refers to these assets as "contractual derivatives." They are not derivative works under copyright law.[8] They do not depend on expression, originality, or authorship.[9] Instead, they derive from contractual allocation of access and control over learning processes embedded within AI services.

Framing the problem in this way clarifies several ongoing debates. Much of the current literature on AI and intellectual property focuses on training data, infringement, and output ownership.[10] Those questions matter, but they miss a deeper shift. The most consequential legal work in AI markets is no longer being done by intellectual property doctrine; it is being done by contract.[11] Through boilerplate agreements, vendors allocate the future value of learning in ways that copyright law would not permit if presented as claims of ownership.

The timing is consequential. Courts are actively resolving questions about the ownership and allocation of learning

8. *See* 17 U.S.C. § 101 (defining derivative work).

9. *See id.*

10. *See, e.g.*, Pamela Samuelson, *Generative AI Meets Copyright*, 381 SCIENCE 158, 158–61 (2023); Matthew Sag, *The New Legal Landscape for Text Mining and Machine Learning*, 66 J. COPYRIGHT SOC'Y U.S. 291, 308–15 (2019).

11. *Cf.* David Nimmer, Elliot Brown, & Gary N. Frischling, *The Metamorphosis of Contract into Expand*, 87 CALIF. L. REV. 17, 22–30 (1999) (discussing the relationship between copyright law and contract law generally); JULIE E. COHEN, BETWEEN TRUTH AND POWER: THE LEGAL CONSTRUCTIONS OF INFORMATIONAL CAPITALISM 44–46 (2019) (describing contract's role in constructing informational capitalism's property-like controls).

artifacts.[12] One District Court judge in Delaware has already rejected fair use for AI training on copyrighted material;[13] the GitHub Copilot litigation[14] and the New York Times v. OpenAI dispute[15] remain pending on appeal and in discovery, respectively. Enterprise AI adoption is accelerating, with billions of dollars in value flowing through contractual provisions that courts have not yet examined.[16] The doctrinal assumptions courts adopt in early cases will shape AI markets for years to come.

---

12. *See, e.g.*, Bartz v. Anthropic PBC, 787 F. Supp. 3d 1007, 1034 (N.D. Cal. 2025) (holding AI model training constitutes fair use but denying summary judgment on piracy-based library claims); Kadrey v. Meta Platforms, Inc., 788 F. Supp. 3d 1026, 1060 (N.D. Cal. 2025) (granting summary judgment on fair use grounds based on absence of market harm evidence). The *Bartz* and *Kadrey* decisions are the first to apply the four-factor fair use analysis to the specific intermediate copying mechanics of generative AI training. *See Bartz*, 787 F. Supp. 3d at 1020; *Kadrey*, 788 F. Supp. 3d at 1059–60. Both found for the AI developer on the training question, though on different analytical grounds and with different implications for the market harm analysis under Factor 4. That *Bartz* is on track to settle before trial, and that *Kadrey* left the market harm question largely open, means the circuit-level resolution of these issues remains forthcoming. *See* Dave Hansen, *The Bartz v. Anthropic Settlement: Understanding America's Largest Copyright Settlement*, KLUWER COPYRIGHT BLOG (Nov. 10, 2025) https://legalblogs.wolterskluwer.com/copyright-blog/the-bartz-v-anthropic-settlement-understanding-americas-largest-copyright-settlement/ [https://perma.cc/QMB6-3ADB] (discussing the preliminary *Bartz* settlement).

13. Thomson Reuters Enter. Ctr. v. Ross Intel. Inc., 765 F. Supp. 3d. 382, 401 (D. Del. 2025) (granting partial summary judgment and rejecting fair use defense for AI training on copyrighted headnotes), *interlocutory appeal granted and proceedings stayed*, No. 1:20-cv-613-SB, 2025 WL 1488015 at *4 (D. Del. May 23, 2025).

14. Doe 1 v. GitHub, Inc., 672 F. Supp. 3d. 837 (N.D. Cal. 2023), *appeal docketed*, No. 24-7700, (9th Cir. argued Feb. 11, 2026).

15. N.Y. Times Co. v. Microsoft Corp., 777 F. Supp.3d 283, 328–29 (S.D.N.Y. 2025) (denying motions to dismiss core copyright claims and proceeding to discovery).

16. *See* STANFORD INST. FOR HUMAN-CENTERED AI, ARTIFICIAL INTELLIGENCE INDEX REPORT 2025 at 17 (2025), https://hai.stanford.edu/assets/files/hai_ai_index_report_2025.pdf

This shift has practical and doctrinal consequences. The allocation of learning in AI markets has moved away from public-law doctrines governing intellectual property and toward private ordering through contract, allowing exclusive control over functional value that public law has not recognized.[17] It shapes market structure by concentrating learning advantages in vendors, increases switching costs, and affects competition in learning-driven markets. It also raises foundational questions about the relationship between contract and intellectual property. If private agreements can routinely allocate exclusive control over learning that public law has excluded from protection, the expressive limits of copyright risk becoming hollow in practice.

The claim advanced here is not that AI contracts are illegitimate or that freedom of contract should yield to public policy wholesale. The claim is narrower and more precise: where contracts allocate exclusive control over learning that copyright law has deliberately excluded from protection, courts should hesitate before enforcing those allocations reflexively. Heightened scrutiny is warranted not because contracts are suspect, but because the allocations at issue substitute for intellectual property rights that the law has refused to recognize.

The Essay proceeds in four Parts. Part I explains why AI training artifacts are not statutory derivative works and why that exclusion reflects copyright's boundary-maintenance function. Part II examines how AI contracts allocate control over learning through service framing, output minimalism, residual rights, and use restrictions. Part III analyzes the economic and doctrinal consequences of contractual derivatives, including the limits of contract autonomy, preemption concerns, unjust enrichment, and competition concerns. Part IV offers a taxonomy of

[https://perma.cc/2CQM-CSMM] (reporting $252.3 billion in total corporate AI investment in 2024).

17. *See* COHEN, *supra* note 11, at 44–46 (describing contract's displacement of public-law information controls).

derivative value and a framework for courts evaluating AI contracts that allocate learning.

I. The Doctrinal Gap: Why AI Learning Is Not a Statutory Derivative

Copyright law protects expression, not function.[18] That principle has done more work in shaping the scope of copyright than any other.[19] It explains why systems, processes, and methods of operation remain outside copyright protection even when they are embodied in expressive works and even when they generate substantial economic value. Understanding why AI learning artifacts fall outside the category of statutory derivative works therefore requires understanding the role that exclusion plays in copyright doctrine.

A. Derivative Works as Boundary Maintenance

The Copyright Act defines a derivative work as one that recasts, transforms, or adapts a preexisting work.[20] That definition presupposes that the new work incorporates protectable expression from the old.[21] From its earliest cases, the Supreme Court has resisted efforts to extend this concept to functional systems. In *Baker v. Selden*,[22] the Court held that copyright in a book describing a bookkeeping system did not confer exclusive rights in

18. *See generally* 17 U.S.C. § 102(b).

19. *See, e.g.*, Baker v. Selden, 101 U.S. 99, 102–04 (1879) (holding that the copyright in question protected publication of the book, not the practice of the methods the book described); Pamela Samuelson, *Why Copyright Law Excludes Systems and Processes from the Scope of Its Protection*, 85 TEXAS L. REV. 1921, 1924–31 (2007) (discussing the history of *Baker*).

20. 17 U.S.C. § 101.

21. *See* MELVILLE B. NIMMER & DAVID NIMMER, NIMMER ON COPYRIGHT § 3.01 (rev. ed. 2025) ("A work is not derivative unless it has *substantially* copied from the prior work.").

22. 101 U.S. 99 (1879).

the system itself.[23] The expressive description was protected; the functional method was not.

Congress codified this principle in Section 102(b), which excludes ideas, procedures, processes, systems, and methods of operation from copyright protection.[24] This exclusion performs a boundary-maintenance function. It prevents copyright from becoming a general law of economic control over useful knowledge.[25] By withholding protection from functional systems, copyright preserves space for competition and reuse even when those systems are valuable.

Derivative works doctrine operates within this framework. It protects expressive transformations of protected works, not the functional mechanisms that enable those transformations.[26] The doctrine's limits are not accidental. They are structural.

B. *Software and the Expression–Function Divide*

Software forced courts to confront the expression–function distinction in new ways. Software is written in expressive code, yet it performs functional tasks. Courts responded by reaffirming the boundary rather than erasing it. In cases such as *Computer Associates v. Altai*[27] and *Lotus v. Borland*,[28] courts

23. *See id.* at 102–04.

24. 17 U.S.C. § 102(b).

25. *See, e.g.*, Samuelson, *supra* note 19, at 1925–31 (explaining how Selden could only economically control sales of the copyrighted book, not the unpatented method it described).

26. *See* Comput. Assocs. Int'l, Inc. v. Altai, Inc., 982 F.2d 693, 706 (2d Cir. 1992) (explaining that "functional" and "utilitarian" aspects of a work must be "filtered" out before determining the scope of protection for derivative or subsequent works); *see also* Sega Enters. Ltd. v. Accolade, Inc., 977 F.2d 1510, 1524 (9th Cir. 1992) (noting that copyright does not grant a monopoly over "functional requirements" even when they are part of a copyrighted work).

27. 982 F.2d 693 (2d Cir. 1992).

28. Lotus Dev. Corp. v. Borland Int'l, Inc., 49 F.3d 807 (1st Cir. 1995), *aff'd* by an equally divided Court, 516 U.S. 233 (1996).

distinguished protectable expression from unprotectable methods of operation, interfaces, and functional structures.[29]

In *Altai*, the Second Circuit established the abstraction-filtration-comparison test, requiring courts to identify levels of abstraction and filter out unprotectable elements—including ideas, processes, facts, public domain information, and elements dictated by efficiency—before comparing what remains.[30] Applied rigorously, this framework recognizes that much of what makes software valuable lies outside copyright's domain.[31] In *Lotus*, the First Circuit held that a spreadsheet's menu command hierarchy was an unprotectable method of operation, emphasizing that protecting those commands would grant copyright over the method of operation itself rather than over any particular expression of it.[32] Judge Boudin's influential concurrence warned against extending copyright to cover the basic method by which a program operates.[33]

Most recently, the Supreme Court's decision in *Google LLC v. Oracle America*[34] reinforced the functional limits of copyright protection.[35] The Court found fair use in Google's reimplementation of Java API declarations, emphasizing the functional nature of the copied elements and the importance of interoperability to competitive software markets.[36] The majority stressed that the declaring code at issue served organizational and functional

29. *See Lotus*, 49 F.3d at 815; *Altai*, 982 F.2d at 704–12; *See also Lotus*, 516 U.S. at 235–36.

30. *Altai*, 982 F.2d at 706–07; *see also* Pamela Samuelson et al., *A Manifesto Concerning the Legal Protection of Computer Programs*, 94 COLUM. L. REV. 2308, 2365–75 (1994).

31. *See Altai*, 982 F.2d at 708–10; *see also* Samuelson *supra* note 19, at 1924–31.

32. *Lotus*, 49 F.3d at 815–19.

33. *See id.* at 819–21 (Boudin, J., concurring).

34. 141 S. Ct. 1183 (2021).

35. *Id.* at 1208.

36. *Id.* at 1202–03.

purposes rather than expressing creative content.[37] Together with *Altai* and *Lotus*, *Oracle* confirms that copyright's boundary-maintenance function is sharpest precisely where functional value is greatest.

The lesson of these cases is not that software lacks copyright protection; it is that copyright protection does not extend to the functional architecture that enables a program to operate. AI systems present a similar but more complex challenge. Like software, they are implemented through expressive artifacts such as code and training data. But the value they generate lies increasingly in internal representations that enable function rather than in expressive outputs.

C. *Why AI Learning Artifacts Fall Outside Derivative Works Doctrine*

Model weights, embeddings, and internal parameter updates do not recast expressive content in any recognizable sense.[38] They are internal system states that enable prediction and generation.[39] While they may be causally influenced by training data, they do not embody expressive elements of that data. They are not readable, interpretable, or consumable as expression.

Treating learning artifacts as derivative works would stretch the doctrine beyond its foundations and convert causal influence

37. *See id.* at 1192, 1198 (comparing the code to other organizational systems like the Dewey Decimal System and travel guides and contrasting it from books, films, and expressive literary works).

38. *See* Lemley & Casey, *supra* note 1, at 748–52 (discussing use of copyrighted materials to train machine learning tools).

39. *See Weight,* GOOGLE: MACHINE LEARNING GLOSSARY, https://developers.google.com/machine-learning/glossary#weight [https://perma.cc/M238-LTV5]; *Parameter,* GOOGLE: MACHINE LEARNING GLOSSARY, https://developers.google.com/machine-learning/glossary#parameter [https://perma.cc/3CUQ-YYX4]; *Embedding Vector,* GOOGLE: MACHINE LEARNING GLOSSARY, https://developers.google.com/machine-learning/glossary#embedding-vector [https://perma.cc/A7Q2-EFHJ].

into expressive incorporation, undermining Section 102(b)'s exclusion of systems and processes.[40] Copyright's refusal to protect functional knowledge is not a bug. It is a feature.

Patent law does not meaningfully fill this gap. While some AI-related inventions may be patentable, the cumulative, opaque, and continuously evolving nature of learning artifacts makes them poor candidates for patent protection.[41] Much of what matters in AI learning arises through use rather than discrete inventive acts. Public law has therefore declined to allocate exclusive rights in learning itself.

Recent judicial authority confirms this conclusion from the fair use side. In *Bartz v. Anthropic*,[42] the first federal court order to apply the four-factor fair use test squarely to the intermediate copying process required for generative AI training, Judge Alsup held that using copyrighted works to train large language models is "exceedingly transformative" because training does not reproduce or supplant those works but generates something functionally distinct from them.[43] The court's reasoning tracks the doctrinal logic developed here: learning artifacts do not embody expressive content from training data, and copyright law does not extend to the methods, patterns, and statistical relationships encoded through training. Critically, the court distinguished this

40. *See* James Grimmelmann, *There's No Such Thing as a Computer-Authored Work*, 39 COLUM. J.L. & ARTS 403, 410–12 (2016) (comparing the allocations of copyright between computer programmer and computer user to other "thing-maker and thing-user" relationships and explaining the difficulty of allocating rights to multiple contributors over time).

41. *See* Sag, *supra* note 10, at 308–15 (discussing the limitations of existing IP frameworks for machine learning).

42. 787 F. Supp. 3d 1007 (N.D. Cal. 2025).

43. *Id.* at 1019–20; *see also Kadrey* v. Meta Platforms, Inc., 788 F. Supp. 3d 1026, 1044, 1054–55 (N.D. Cal. 2025). The two decisions reflect an emerging tension between Factor 1 (transformative purpose) and Factor 4 (market effect) that remains unresolved at the circuit level. *See supra* note 12.

conclusion from the question of how training copies were obtained—holding that the transformative character of the training use was independent of whether source material came from legitimate or pirated sources.[44] The fair use holding thus confirms that learning artifacts occupy the space copyright has deliberately left vacant. It does not settle how that space is governed once copyright withdraws. That is the question this Essay addresses.

D. The Unallocated Remainder

The absence of statutory protection for AI learning artifacts creates a doctrinal gap. That gap does not imply that learning is valueless or ungoverned. It implies that public law has chosen not to allocate exclusive rights in it. The critical question is how that value is allocated in practice.

The gap is not filled by trade secret law, though the analogy is instructive. Trade secret protection can cover functional, nonexpressive information, but it requires that information derive independent economic value from its secrecy and that the holder take reasonable steps to maintain that secrecy.[45] Contractual derivatives do not depend on secrecy in this sense. The vendor's control over learning artifacts arises from the structure of the agreement, from service framing, residual rights, and use

44. *Bartz*, 787 F. Supp. 3d at 1022. The court drew a sharp distinction between the training use itself, which it found fair, and Anthropic's retention of a permanent digital library of pirated copies for general purposes, which it found indefensible. *Id.* at 1022, 1029. The latter use failed each factor because it was not transformative, served no independent purpose, and involved deliberate appropriation of works Anthropic could have acquired lawfully. *Id.* at 1029–31. The distinction is important for the present analysis: the court's fair use holding for training is not a general license to copy without payment; it is a finding that the functional use of works in training does not extend copyright's reach to the resulting artifacts.

45. 18 U.S.C. § 1839(3) (defining trade secret); *see also* COHEN, *supra* note 11, at 88–92 (discussing the relationship between contractual and property-based information controls).

restrictions, not from the information's secret character. Model weights may or may not independently qualify as trade secrets, but the contractual allocation of learning operates regardless of whether secrecy conditions are satisfied. The distinction matters because trade secret law has its own built-in limitations, including independent discovery, reverse engineering, and the requirement of secrecy, that contractual derivatives bypass.[46]

As the next Part shows, AI vendors have answered the allocation question through contract. Rather than asserting ownership over learning as intellectual property, they have designed agreements that allocate control indirectly, by structuring access, reserving residual rights, and restricting reuse.[47] To understand contractual derivatives, one must first understand why copyright law has declined to recognize learning as a derivative work.

II. The Contractual Workaround: How AI Contracts Allocate Learning

Copyright law's refusal to protect systems and processes leaves learning unowned as a matter of public law. That refusal, however, has not prevented the allocation of learning value in practice. Instead, it has shifted the site of allocation from statute to contract. AI vendors do not attempt to characterize learning artifacts as expressive works or patentable inventions. They do something more effective: They design contractual architectures

46. *See* Kewanee Oil Co. v. Bicron Corp., 416 U.S. 470, 476 (1974) ("A trade secret law, however, does not offer protection against discovery by fair and honest means, such as by independent invention, accidental disclosure, or by so-called reverse engineering, that is by starting with the known product and working backward to divine the process which aided in its development or manufacture.")

47. *See, e.g.*, *OpenAI Services Agreement*, *supra* note 7; Google Cloud, *Service Specific Terms*, *supra* note 7.

that allocate exclusive control over learning indirectly, without ever asserting ownership in the traditional sense.[48]

This Part explains how that allocation occurs. The central insight is that no single clause does the work alone. Control over learning emerges from the interaction of service framing, narrow output definitions, residual rights, improvement clauses, and use restrictions. Together, these provisions create durable exclusion over learning while avoiding the doctrinal limits that would apply if learning were framed as intellectual property.

A. *Why Contract, Not Intellectual Property, Allocates Learning*

Contract is uniquely suited to govern learning because it operates *ex ante* and without categorical constraints.[49] Intellectual property doctrines impose threshold requirements. Copyright demands expression. Patent law demands novelty, non-obviousness, disclosure, and fixation.[50] Learning artifacts satisfy none of these requirements in stable form. They are functional, cumulative, and continuously evolving. Attempting to force them into existing IP categories would require doctrinal distortion.

Contract, by contrast, allows parties to allocate future, undefined value before it exists.[51] Parties routinely agree on how value

---

48. *See* COHEN, *supra* note 11, at 88–92; *see also OpenAI Services Agreement*, *supra* note 7.

49. *See* 17 U.S.C. §§ 101, 102(a)–(b); 35 U.S.C. §§ 101–03, 112(a); Baker v. Selden, 101 U.S. 99, 102–05 (1879).

50. UNITED STATES PATENT AND TRADEMARK OFFICE, PATENT ELIGIBLE SUBJECT MATTER: REPORT ON VIEWS AND RECOMMENDATIONS FROM THE PUBLIC 1 (2017) ("the invention must meet other patentability requirements, including novelty, non-obviousness, written description, and enablement."); *see also* 35 U.S.C. §§ 102–103.

51. *See* Ronald J. Gilson, Charles F. Sabal, & Robert E. Scott, *Contracting for Innovation: Vertical Disintegration and Interfirm Collaboration*, 109 COLUM. L. REV. 431, 437–40 (2009); Thomas W. Merrill & Henry E. Smith, *Optimal Standardization in the Law of Property: The Numerus Clausus Principle*, 110 YALE L.J. 1, 3–4, 24–30 (2000).

will be allocated once it emerges, even when that value cannot be precisely described at the time of contracting. In AI markets, this temporal feature is decisive. Learning artifacts do not exist when agreements are signed. Contract nonetheless determines where they will reside once they come into being.

Equally important, contract enforcement does not require courts to decide whether learning qualifies for protection. Courts enforce agreements governing access and use, not ownership in the abstract.[52] This allows vendors to achieve exclusionary effects without triggering copyright's boundary-maintenance doctrines. Control over learning is achieved through conditions on access, not claims of authorship.

B. Service Framing and the Black-Box Allocation of Learning

The foundational move in AI contracting is the characterization of AI systems as services rather than goods, deliverables, or jointly developed assets.[53] Customers purchase access to functionality, not transfer of technology. The underlying system remains provider-controlled at all times.

Service framing performs quiet but decisive legal work. In a service relationship, use does not confer ownership.[54] Courts have long rejected the idea that accessing a service grants proprietary interests in the infrastructure through which the service operates.[55] This principle, well established in SaaS and cloud computing, carries over seamlessly to AI systems.

By framing AI as a service, vendors foreclose alternative ownership narratives. If an AI system were characterized as a

52. *See* Samuelson, *supra* note 19, at 1924–27.

53. *See, e.g.*, *OpenAI Services Agreement*, *supra* note 7; GOOGLE CLOUD, *Service Specific Terms*, *supra* note 7; *Azure Preview Supplemental Terms*, *supra* note 7.

54. *See* Marquette Univ. v. Kuali, Inc., 584 F. Supp. 3d 720, 727 (E.D. Wis. 2022).

55. *See* Vernor v. Autodesk, Inc., 621 F.3d 1102, 1110–12 (9th Cir. 2010).

customized model or a jointly developed asset, customer participation in training could plausibly support claims of shared ownership or derivative interest. Service framing prevents those arguments from arising. Learning remains inside the provider's black box, legally and technically inaccessible.

This framing aligns with technical opacity.[56] Customers cannot observe internal parameter updates or isolate their contribution to system improvement. Legal service framing and technical opacity reinforce one another. Together, they ensure that learning remains vendor-controlled by default.

C. Output Minimalism and the Exclusion of Learning

Within the service framework, AI contracts define customer-owned output narrowly.[57] Output typically refers to the immediate content returned in response to a prompt: generated text, images, or structured responses. Nothing more.

This minimalism is not incidental. AI systems do not derive their value from isolated outputs. They derive value from accumulated learning across interactions.[58] By defining output narrowly, contracts exclude the most valuable learning artifacts from customer ownership by definition.

The classification matters because ownership follows definition. If learning artifacts were treated as outputs, customers could plausibly claim rights in them. Questions of authorship, derivation, and contribution would follow. Narrow output definitions prevent those questions from arising. Learning is classified as process, not product.

56. *See* Jenna Burrell, *How the Machine 'Thinks': Understanding Opacity in Machine Learning Algorithms*, 3 BIG DATA & SOC'Y 1, 3–5 (2016).

57. *See, e.g.*, *OpenAI Services Agreement*, *supra* note 7.

58. *See* Lemley & Casey, *supra* note 1, at 771–75.

D. *Residual Rights and Allocation by Omission*

AI contracts almost invariably include residual rights clauses providing that all rights not expressly granted are reserved to the provider.[59] These clauses are deceptively powerful. Because learning artifacts are never expressly allocated, residual rights ensure that control over them remains with the vendor by default.

This strategy operates through omission rather than assertion. Vendors do not need to claim ownership over weights or embeddings.[60] By declining to mention them, they ensure that no rights pass to the customer. Contractual silence therefore operates in tandem with technical opacity. What cannot be seen cannot be negotiated. What cannot be negotiated remains reserved.

E. *Improvement Clauses and the Erasure of Contribution*

Many AI agreements include provisions allocating ownership of improvements, updates, or modifications to the service.[61] These clauses typically apply regardless of source. Whether an improvement arises from provider research, aggregate user interaction, or customer-specific feedback, it is treated as part of the provider's service.

Improvement clauses erase contribution as a legally relevant category.[62] Learning is not attributed to any customer, dataset, or use case. It is absorbed into the service as a whole. Here, contract substitutes categorical allocation for doctrinal analysis. It does not ask who contributed to learning. It declares where learning belongs.

---

59. *See, e.g.*, *OpenAI Services Agreement*, *supra* note 7; ANTHROPIC, *supra* note 7.

60. *See* COHEN, *supra* note 11, at 88–92.

61. *See, e.g.*, *OpenAI Services Agreement*, *supra* note 7; GOOGLE CLOUD, *supra* note 7.

62. *See generally* Gilson et al., *supra* note 51, at 442–44.

F. *Use Restrictions as Barriers to Learning Diffusion*

Definitional and allocation strategies are reinforced by behavioral restrictions. AI contracts commonly prohibit reverse engineering, model extraction, and the use of outputs to train competing systems.[63] These restrictions prevent customers from externalizing learning even indirectly.

These clauses are not ancillary. They are essential. Without them, customers could attempt to reconstruct learning through observation, benchmarking, or reuse of outputs. Use restrictions close that avenue.[64] Learning remains locked inside the vendor's platform. In effect, these restrictions create legal barriers to learning diffusion. They ensure that learning accumulates centrally rather than spreading through use and competition. This is not a technological necessity. It is a legal choice.

G. *Standardization and the Limits of Enterprise Negotiation*

It is tempting to assume that sophisticated enterprise customers can negotiate around these provisions. In practice, they rarely do.[65] While pricing, service levels, and indemnities may be negotiable, ownership allocations and use restrictions typically are not.

This reflects repeat-player dynamics. Vendors operate at scale and depend on centralized learning to maintain competitive advantage.[66] Custom allocations would undermine that model. As a result, even enterprise customers encounter standardized terms with limited variation on learning control. The result is not coercion in the traditional sense. It is structural

63. *See, e.g.*, *OpenAI Services Agreement*, *supra* note 7; *Azure Preview Supplemental Terms*, *supra* note 7.

64. *See* COHEN, *supra* note 11, at 88–92.

65. *See* MARGARET JANE RADIN, BOILERPLATE: THE FINE PRINT, VANISHING RIGHTS, AND THE RULE OF LAW 40–42 (2013).

66. *See* COHEN, *supra* note 11, at 42–45, 85; Gilson et al., *supra* note 51, at 437–40.

asymmetry. Customers may negotiate the price of access, but not the allocation of learning generated through that access.

H. Contractual Derivatives Defined

Taken together, these mechanisms create contractual derivatives: learning-based assets that are not statutory derivative works but are rendered proprietary through private ordering.[67] They replicate the exclusionary effects of intellectual property without invoking its doctrines.

Naming this phenomenon matters. Without a concept of contractual derivatives, disputes over AI contracts appear fragmented. With it, the pattern becomes visible. Contract is doing the work that intellectual property law has declined to do. The next Part examines why that substitution matters, and where the limits of contractual allocation should lie.

III. Why Contractual Derivatives Matter: Consequences and Limits

Contractual derivatives do more than allocate value between vendors and customers. They reshape market structure, alter competitive dynamics, and test the limits of doctrinal accommodation between contract and intellectual property. This Part explains why the contractual allocation of learning is not merely a private ordering choice, but a legal development with systemic consequences. It then confronts the strongest defense of these arrangements: the claim that freedom of contract, especially among sophisticated parties, should resolve the matter.

A. Learning as the Primary Source of Competitive Advantage

In traditional software markets, competitive advantage often turned on features, pricing, or distribution. In AI markets, it turns on learning.[68] The ability to aggregate, retain, and redeploy

---

67. *See* COHEN, *supra* note 11, at 44–46.

68. *See generally* Lemley & Casey, *supra* note 1, at 752–56; STIGLER CTR. REPORT, *supra* note 1, at 40–42, 46–48.

learning across deployments determines performance, scalability, and cost. Systems trained on broader interaction sets improve faster, generalize better, and dominate markets characterized by feedback effects.[69]

Contractual derivatives concentrate this learning advantage in vendors. Each customer interaction contributes incrementally to system improvement, but contractual allocation ensures that those increments accrete centrally rather than remaining with or returning to the customer. Over time, this produces asymmetric accumulation. Vendors learn from everyone. Customers learn from no one but the vendor.

This asymmetry persists even where customers receive improved service in return. Access to improved functionality is not equivalent to control over learning. The former is conditional and revocable. The latter is durable and exclusionary. Contractual derivatives therefore create a one-way ratchet. Learning flows inward, never outward.

B. Lock-In, Switching Costs, and Market Entrenchment

The accumulation of learning through contract has predictable effects on market structure.[70] As vendors aggregate learning, switching costs increase. Customers who move to competing systems do not take their learning with them. They return to baseline performance, while the incumbent retains the benefit of prior interactions.

This dynamic amplifies lock-in beyond what traditional SaaS contracts produce. In SaaS, switching costs arise from data migration, retraining, and integration. In AI services, they also arise

---

69. *See generally* Lemley & Casey, *supra* note 1, at 765–68 (describing feedback effects in machine learning systems); STIGLER CTR. REPORT, *supra* note 1, at 40–43.

70. *See* Carl Shapiro & Hal R. Varian, INFORMATION RULES: A STRATEGIC GUIDE TO THE NETWORK ECONOMY 103–107 (1998).

from the loss of accumulated learning.[71] That loss is often invisible at the time of contracting but becomes salient only after extended use. The result is path dependence. Early contractual choices shape long-term competitive outcomes, even when customers later regret those choices.[72] Contractual derivatives thus operate as market-shaping instruments, not merely bilateral allocations.

C. *The Contract Autonomy Defense*

The strongest defense of contractual derivatives appeals to freedom of contract. AI customers, particularly enterprises, are sophisticated actors. They negotiate complex agreements, evaluate tradeoffs, and allocate risk routinely.[73] If such parties agree to allocate learning to vendors, courts should respect that choice.

This argument has force. Contract autonomy is a cornerstone of commercial law. Courts are rightly reluctant to rewrite bargains or protect parties from their own decisions. In many contexts, boilerplate terms are enforced even when they allocate value asymmetrically.[74]

But contractual derivatives strain this logic. Autonomy presupposes meaningful choice. Meaningful choice presupposes information. In AI markets, neither condition is reliably satisfied with respect to learning allocation.

71. *See generally* Lemley & Casey, *supra* note 1, at 765–68; Michael L. Katz & Carl Shapiro, *Network Externalities, Competition, and Compatibility*, 75 AM. ECON. REV. 424, 425–30 (1985).

72. *See* W. Brian Arthur, *Self-Reinforcing Mechanisms in Economics*, *in* INCREASING RETURNS AND PATH DEPENDENCE IN THE ECONOMY 111, 116–20 (W. Brian Arthur ed., 1994).

73. *Cf.* ProCD, Inc. v. Zeidenberg, 86 F.3d 1447, 1453–55 (7th Cir. 1996) (discussing negotiations with sophisticated enterprises).

74. *See* RADIN, *supra* note 65, at 33–34, 38–40.

D. Why Consent Is Incomplete in AI Learning Allocation

Learning artifacts are not observable. Customers cannot inspect internal parameter updates, isolate their contribution to system improvement, nor quantify the value of learning generated through their use.[75] Even sophisticated parties lack the information necessary to assess what they are allocating.

Learning value is also temporally misaligned with consent.[76] Its significance emerges over time, often after contracts are renewed or terminated. Customers may consent to service terms without appreciating that their interactions will generate durable assets used to improve the vendor's system across customers and markets.

Finally, learning allocation is rarely negotiated. It appears as background architecture rather than as a bargained-for term.[77] Customers may negotiate price, service levels, or indemnities. They almost never negotiate residual rights in learning. This is not because learning is unimportant, but because it is difficult to see and harder to value *ex ante*.

These features distinguish contractual derivatives from ordinary allocations of risk or ownership. Customers may knowingly

---

75. *See* Burrell, *supra* note 56, at 3–5.

76. *Cf.* George A. Akerlof, *The Market for "Lemons": Quality Uncertainty and the Market Mechanism*, 84 Q.J. ECON. 488, 488–90, 492–95, 497–98 (1970) (establishing the foundational problem of information asymmetry); Oliver E. Williamson, *The Economic Institutions of Capitalism* 44–46, 66–67 (1985) (describing "contractual incompleteness" as the result of bounded rationality and the inability to specify complex future contingencies); Oliver Hart, *Incomplete Contracts and Public Ownership: Remarks, and an Application to Public-Private Partnerships*, ECON. J., March 2003, at 691, 670 (explaining that "residual rights" to value typically fall to the party with the most bargaining power when a contract fails to explicitly allocate a specific asset).

77. *Cf.* NANCY S. KIM, WRAP CONTRACTS: FOUNDATIONS AND RAMIFICATIONS 87–112 (2013).

accept unfavorable pricing. They do not knowingly accept the future competitive use of learning they cannot observe.

E. Preemption, Policy, and Boundary Erosion

These consent failures intersect with copyright policy. Federal copyright law reflects a deliberate choice to exclude systems and processes from protection.[78] Allowing private agreements to recreate exclusive control over those systems risks undermining that choice.

Courts have long distinguished between contracts that govern access to copyrighted works and contracts that create rights equivalent to copyright. The former are generally enforceable. The latter may be preempted where they interfere with the balance Congress struck between protection and free use.[79]

The circuit split on this question is instructive. In *ProCD, Inc. v. Zeidenberg*,[80] the Seventh Circuit held that contracts restricting use of uncopyrightable databases do not conflict with federal copyright policy, reasoning that contracts create obligations only between parties rather than rights against the world.[81] In *Vault Corp. v. Quaid Software Ltd.*,[82] the Fifth Circuit took a contrary view, invalidating arrangements creating copyright-equivalent rights in functional elements.[83] The Federal Circuit in *Bowers v. Baystate Technologies*[84] further complicated the picture by enforcing a shrinkwrap license prohibiting reverse engineering, over a vigorous dissent arguing that such enforcement

78. 17 U.S.C. § 117; *see also* Vault Corp. v. Quaid Software Ltd., 847 F.2d 255, 269–70 (5th Cir. 1988).

79. *See* Mark A. Lemley, *Beyond Preemption: The Law and Policy of Intellectual Property Licensing*, 87 CALIF. L. REV. 111, 143–47, 151–52 (1999).

80. 86 F.3d 1447 (7th Cir. 1996).

81. *Id.* at 1453–55.

82. 847 F.2d 255 (5th Cir. 1988).

83. *Id.* at 268–70.

84. 320 F.3d 1317 (Fed. Cir. 2003).

nullified the fair use defense Congress had provided.[85] This split leaves the enforceability of contractual derivatives uncertain but signals that courts recognize the tension between private ordering and federal copyright policy.[86]

Contractual derivatives sit uneasily between these categories. Vendors do not claim copyright in learning artifacts. Yet they use contract to achieve similar exclusionary effects. Customers are prevented from reusing, extracting, or redeploying learning even though that learning is not protected expression. This does not require categorical preemption. Contract has a legitimate role in allocating access and confidentiality. But where contractual enforcement recreates exclusion over functional knowledge that copyright law deliberately leaves unprotected, heightened scrutiny is warranted.[87] Otherwise, copyright's boundary-maintenance function survives only in theory.

F. *Unjust Enrichment and the Limits of Consideration*

Customers contribute data, validation, domain expertise, and behavioral feedback that materially improve system performance. Yet contractual derivatives treat learning artifacts as vendor property irrespective of contribution. Courts recognize that enrichment may be unjust even where contracts exist, particularly where one party captures value enhanced by another's contribution without commensurate exchange.[88]

---

85. *Id.* at 1325–26.

86. *See* Nimmer et al., *supra* note 11, at 27–35; Maureen A. O'Rourke, *Drawing the Boundary Between Copyright and Contract*, 45 DUKE L.J. 479, 548-55 (1995).

87. *See* Lemley, *supra* note 79, at 152–55.

88. *See* Schock v. Nash, 732 A.2d 217, 232–33 (Del. 1999) (noting that enrichment is "unjust" where it would be "inequitable" for a recipient to retain a benefit without payment, regardless of the parties' initial intent); Bd. of Trade of Chi. v. Dow Jones & Co., 456 N.E.2d 84, 89–90 (Ill. 1983) (finding that free-riding on the investment of another to capture enhanced commercial value is an actionable misappropriation); *see also generally* Ayelet Gordon-Tapiero & Yotam Kaplan,

Vendors respond that service access constitutes adequate consideration. But learning benefits flow outward—enhancing the vendor's platform across all customers and markets—in ways that far exceed the value delivered to any individual customer.[89] Whether one-sided allocation of contingent future value satisfies consideration requirements, or whether it should be subject to greater scrutiny where the contributing party cannot observe or value what it is giving up, warrants closer examination.[90]

G. *Competition and Innovation Effects*

Beyond bilateral fairness, contractual derivatives affect innovation.[91] Centralized learning advantages favor incumbents and raise barriers to entry. New entrants face not only capital and talent constraints, but learning deficits that cannot be overcome quickly.

This dynamic interacts with network effects and scale economies.[92] Contractual derivatives accelerate winner-take-most outcomes by ensuring that learning accrues where it compounds fastest.[93] While such outcomes may be efficient in some contexts, they also raise competition policy concerns, particularly where contractual design rather than technical necessity drives accumulation.

Generative AI Training as Unjust Enrichment, 86 OHIO ST. L.J. 285 (2025) (applying unjust enrichment to AI developers' capture of creators' labor; analogous dynamics apply to vendors' capture of customer-generated learning).

89. *See* Salome Viljoen, *A Relational Theory of Data Governance*, 131 YALE L.J. 573, 615–25 (2021).

90. *See* Brett M. Frischmann & Mark A. Lemley, *Spillovers*, 107 COLUM. L. REV. 257, 280–84 (2007).

91. *See* Herbert Hovenkamp, *Antitrust and Platform Monopoly*, 130 YALE L.J. 1952, 1980–88 (2021).

92. STIGLER CTR. REPORT, *supra* note 1, at 40–41.

93. *See* Jean-Charles Rochet & Jean Tirole, *Platform Competition in Two-Sided Markets*, 1 J. EUR. ECON. ASS'N 990, 1002-1005 (2003).

Antitrust law offers tools to address these concerns, but those tools typically operate *ex post* and at scale. Contractual scrutiny offers an *ex ante* complement. By examining how learning is allocated at the point of agreement, courts can address competitive effects before they harden.

H. Heightened Scrutiny, Not Categorical Invalidation

None of this requires courts to invalidate AI contracts wholesale. Nor does it justify treating contractual derivatives as per se unlawful. Contract remains an essential governance tool in complex markets.[94]

What it does require is a shift in posture. Where contracts allocate exclusive control over learning that public law has deliberately excluded from protection, courts should ask harder questions:[95] Was the allocation observable and reasonably valued at the time of consent? Does enforcement undermine the policy reasons for excluding systems and processes from copyright? Are use restrictions functioning as de facto intellectual property rights?

Heightened scrutiny respects contract autonomy while acknowledging its limits. It allows courts to distinguish between routine service terms and provisions that function as substitutes for intellectual property rights. The next Part develops a

94. *See* ProCD, Inc. v. Zeidenberg, 86 F.3d 1447, 1454 (7th Cir. 1996) (affirming that private contractual ordering remains a primary tool for governing information use where intellectual property rights are thin or non-existent).

95. *See* Amit Elazari, *Unconscionability 2.0 and the IP Boilerplate*, 34 BERKELEY TECH. L.J. 567, 582–83 (2019) (calling for a "revised doctrine of unconscionability" specifically for contracts that function as de facto intellectual property regimes); *see also* Restatement (Second) of Contracts § 178 (Am. L. Inst. 1981) (providing that a promise is unenforceable on grounds of public policy if the interest in enforcement is outweighed by a public policy against such enforcement—such as the policy underlying the § 102(b) exclusion).

taxonomy of derivative value and offers a framework for courts applying this scrutiny in practice.

IV. A Taxonomy of Derivative Value and a Framework for Courts

The preceding Parts have shown that AI contracts allocate exclusive control over learning artifacts that fall outside statutory intellectual property. This Part supplies the analytic vocabulary courts need to recognize what is happening and to respond without distorting existing doctrine. It does so in two steps. First, it distinguishes among three forms of derivative value that appear in AI markets. Second, it offers a framework for evaluating when contractual allocation of learning warrants heightened scrutiny rather than routine enforcement.

A. Three Forms of Derivative Value

Disputes over AI contracts often collapse distinct phenomena into a single question of "ownership." That collapse obscures what is actually being allocated. The taxonomy offered here separates three forms of derivative value, each governed by different legal constraints.[96]

1. Statutory Derivatives

Statutory derivatives are the familiar category governed by copyright law.[97] They arise when a protected work is recast, transformed, or adapted in a manner that incorporates protected expression. Ownership and infringement turn on statutory criteria such as originality, fixation, and substantial similarity.

96. *See generally* Paul Goldstein, *Derivative Rights and Derivative Works in Copyright*, 30 J. COPYRIGHT SOC'Y U.S.A. 209, 216–18 (1983).; O'Rourke, *supra* note 86, at 532–48.

97. 17 U.S.C. § 106 ("[T]he owner of copyright under this title has the exclusive rights to . . . prepare derivative works based upon the copyrighted work.")

In the AI context, statutory derivatives may arise at the output level, where generated content bears a sufficiently close relationship to protected expression.[98] Whether particular outputs actually constitute infringement remains contested. What matters here is what statutory derivatives do *not* include. They do not include internal learning artifacts such as model weights, embeddings, or internal parameter updates that enable function rather than expression.

2. Contractual Derivatives

Contractual derivatives occupy the gap left by statutory derivatives.[99] They are assets rendered proprietary not by statute but by agreement. Their defining feature is not expressive transformation, but contractual allocation: contracts allocate exclusive control over learning artifacts that are not protected expression.

This category explains why ownership language is often absent from AI contracts. Vendors do not need to claim ownership in the traditional sense. By structuring access and restricting use, they achieve the same exclusionary effects.[100] Contractual derivatives therefore replicate the economic function of intellectual property without invoking its legal tests.

Contractual derivatives are not inherently illegitimate. Contracts routinely allocate value that public law leaves unowned, as in trade secrets or confidential know-how.[101] What distinguishes contractual derivatives in AI markets is their scale, opacity, and interaction with copyright's deliberate exclusion of functional systems.

98. *See* Samuelson, *supra* note 10, at 158–61.

99. *See* O'Rourke, *supra* note 86, at 548–55.

100. *See supra* Part III.

101. *See* Robert P. Merges, *Contracting into Liability Rules: Intellectual Property Rights and Collective Rights Organizations*, 84 CALIF. L. REV. 1293, 1295–1300 (1996).

3. Functional Derivatives

Functional derivatives describe value that arises from use, learning, and adaptation but is not allocated exclusively by statute or contract.[102] They include improvements users internalize, efficiencies gained through experience, and knowledge that diffuses through observation and competition.

In many markets, functional derivatives spread naturally. Copyright facilitates this diffusion by excluding systems and processes from protection. Learning moves through markets as experience, not property.

AI contracts disrupt this pattern. Learning that would otherwise diffuse as a functional derivative is captured centrally as a contractual derivative.[103] This transformation does not occur through technological necessity alone. It occurs through legal design. The taxonomy makes visible what is displaced when learning is contractually enclosed.

B. Why the Taxonomy Matters

The taxonomy clarifies why existing doctrinal tools often feel misapplied in AI contract disputes.[104] Asking whether learning artifacts are statutory derivatives answers the wrong question. They are not. Asking whether contracts "own" learning is similarly unhelpful. Ownership is rarely claimed.

The relevant inquiry is whether contracts should be permitted to allocate exclusive control over learning in ways that undermine the policy choices that justified its exclusion from statutory protection in the first place. By naming contractual derivatives as a distinct category, the taxonomy allows courts to recognize that something derivative-like is occurring without distorting copyright doctrine.

102. *See* Lemley & Casey, *supra* note 1, at 760–65.

103. *See* Viljoen, *supra* note 89, at 625–30.

104. *See* Lemley, *supra* note 79, at 152–55.

The taxonomy also explains why disputes over AI contracts often feel incoherent. Without it, courts oscillate between enforcing boilerplate terms and expressing unease about their effects. With it, courts can identify which category of value is being allocated and apply the appropriate level of scrutiny.

C. A Framework for Judicial Evaluation

The taxonomy supports a practical framework for courts faced with disputes over AI contracts. Rather than beginning with ownership, courts should focus on allocation and effect. Three questions are central.

First, does the contract allocate exclusive control over learning artifacts or improvements that are functionally excluded from statutory intellectual property protection?[105] If the answer is no, routine enforcement is appropriate. If the answer is yes, further inquiry is warranted.

Second, was the allocation observable and reasonably valued at the time of consent?[106] Courts should consider whether learning was disclosed, described, or meaningfully negotiable. Silence, technical opacity, and standardized terms weigh against routine enforcement where learning constitutes the primary source of value.

Third, does enforcement undermine the boundary-maintenance function of copyright law? Where contractual enforcement recreates exclusion over systems or processes that copyright law deliberately leaves unprotected, heightened scrutiny is appropriate. The question is not whether contract is disfavored, but whether its enforcement substitutes for intellectual property rights the law has refused to grant.

---

105. *See* O'Rourke, *supra* note 86, at 548–55.

106. *Cf.* Inge Graef et al., *Data Portability and Data Control*, 19 GERMAN L.J. 1359, 1375–85 (2018).

The framework can be illustrated concretely. Consider a bank that licenses an AI fraud detection platform. The platform improves substantially through fine-tuning on the bank's proprietary transaction data, fraud labels, and validation feedback. False positives decline, detection accuracy increases, and the model adapts to the bank's specific risk profile. The contract assigns ownership of all "derivatives and improvements," including trained models, internal parameter updates, and generalized learnings to the vendor. After termination, the vendor deploys an improved fraud detection model to the bank's competitors, incorporating insights derived from the bank's proprietary data.[107]

A court applying the framework would proceed as follows. First, the court would ask whether the claimed "derivatives," such as learning artifacts, incorporate protected expression in recognizable form. They do not. They are functional artifacts that encode statistical patterns, not expressive adaptations. The claim therefore does not sound in copyright, whatever terminology the contract uses. Second, the court would evaluate the contractual allocation under contract law. Was the allocation of learning observable and reasonably valued at the time of consent? Given the opacity of model training and the temporal misalignment between consent and value realization, heightened scrutiny is warranted. Third, the court would ask whether enforcement undermines copyright's boundary-maintenance function by recreating exclusive control over functional knowledge that copyright law deliberately leaves unprotected. Where the vendor's competitive use of customer-derived learning extends beyond the contracting relationship and into the customer's

107. This hypothetical is stylized but reflects contractual patterns the author has observed in enterprise AI agreements for financial services. *See supra* author information note (describing the author's experience negotiating AI-enabled service agreements); *supra* Part II (analyzing the contractual mechanisms that produce these allocations).

competitive market, enforcement risks precisely the kind of functional propertization that copyright law has resisted since *Baker v. Selden*.

This framework does not dictate outcomes; it structures analysis. In some cases, enforcement will be appropriate. In others, courts may construe clauses narrowly, refuse to enforce use restrictions that operate as de facto intellectual property, or require clearer disclosure of learning allocation.

D. Addressing Predictable Objections

One objection is that heightened scrutiny invites uncertainty. Parties benefit from predictable enforcement of contracts.[108] The response is that predictability already coexists with scrutiny in many areas of law. Courts routinely evaluate unconscionability, public policy, and preemption without destabilizing contract law. The framework proposed here is narrower and more targeted.

Another objection is that customers benefit from centralized learning through improved services. That is true. But benefit does not imply consent to all downstream uses of learning, nor does it justify ignoring structural asymmetries in allocation. Improved service access is not equivalent to control over learning.

Finally, some will argue that market forces suffice. Customers who dislike learning allocations can choose other vendors.[109] This assumes competitive conditions that contractual derivatives themselves may undermine. Where learning-driven lock-in shapes markets, choice is constrained by the very contracts under examination.

E. From Taxonomy to Judgment

The taxonomy does not mandate a particular remedial approach; it informs judgment. Courts may respond by interpreting ambiguous clauses narrowly, declining to enforce restrictions

108. *See* Shapiro & Varian, *supra* note 70, at 173–79.

109. *See* Hovenkamp, *supra* note 91, at 1985–90.

that function as intellectual property by another name, or demanding clearer disclosure where learning allocation is central to the bargain.[110]

What matters is recognition. Without recognizing contractual derivatives as a distinct phenomenon, courts risk enforcing private arrangements that quietly reshape the public law of intellectual property. With that recognition, courts can preserve both freedom of contract and the limits that intellectual property law has long enforced.

Conclusion

Artificial intelligence markets are organized around learning. That learning does not merely improve performance in the abstract. It accumulates value, shapes competition, and determines who controls the future trajectory of AI systems. Yet public law has declined to allocate exclusive rights in that learning. Copyright excludes systems and processes by design, and patent law captures only a narrow subset of functional advances.[111] The result is not a doctrinal oversight, but a deliberate refusal to privatize functional knowledge.

This Essay has shown how contract fills that space. AI vendors do not claim ownership over learning as intellectual property. Instead, they allocate exclusive control over learning indirectly. Operating together, these mechanisms produce what this Essay has called *contractual derivatives*: learning-based assets that function like intellectual property but arise through private ordering rather than statute.

Recognizing contractual derivatives clarifies why existing debates about AI and intellectual property often feel incomplete. Focusing on training data, infringement, or output ownership

110. *See* Vault Corp. v. Quaid Software Ltd, 847 F.2d 255, 268–70 (5th Cir. 1988).

111. *See* Baker v. Selden, 101 U.S. 99, 102–03 (1879); 17 U.S.C. § 102(b).

leaves untouched the most consequential allocation of value in AI markets. That allocation occurs at the level of learning. It is governed not by copyright doctrine, but by contract design.

The Essay does not argue that contractual derivatives are inherently illegitimate. Contract plays a vital role in coordinating complex systems and allocating risk, and customers may benefit from centralized learning through improved services. The claim advanced here is narrower. Where contracts allocate exclusive control over learning that public law has deliberately excluded from protection, courts should hesitate before enforcing those allocations reflexively. Heightened scrutiny is warranted not because contracts are suspect, but because the allocations at issue substitute for intellectual property rights the law has refused to recognize.

The taxonomy developed here offers a way to operationalize that scrutiny. By distinguishing among statutory, contractual, and functional derivatives, courts can identify what kind of value is being allocated and how. That clarity matters. It allows judges to ask the right questions: Does a contract allocate exclusive control over learning that copyright law excludes from protection? Was that allocation observable and reasonably valued at the time of consent? Does enforcement undermine the policy reasons for excluding systems and processes from copyright in the first place?

These questions should structure analysis. In some cases, enforcement will be appropriate. In others, courts may construe clauses narrowly, decline to enforce use restrictions that operate as de facto intellectual property, or require clearer disclosure where learning allocation is central to the bargain. The point is not to mandate results, but to ensure that contractual enforcement does not quietly erode the limits that intellectual property law has long enforced.

The implications extend beyond adjudication. For transactional lawyers and in-house counsel, learning allocation should no longer be treated as background noise. Silence is not neutral; it is itself an allocation. Counsel negotiating AI agreements should understand that control over learning may matter more than ownership of outputs and should address that allocation explicitly. For vendors, greater transparency around learning allocation may reduce litigation risk and build trust. For customers, awareness of contractual derivatives may inform procurement decisions and long-term strategy.[112]

More broadly, the analysis offered here suggests a shift in how intellectual property should be understood in data-driven markets.[113] As value migrates from expression to function, pressure to recreate exclusivity through private ordering will intensify. Contract will continue to do work that public law declines to do. The challenge is to ensure that this work does not quietly reconstruct intellectual property rights the law has deliberately refused to grant.[114]

Learning is not owned by default. When contracts allocate it anyway, the law should notice.

The emergence of judicial authority confirming that copyright does not reach AI learning artifacts—most directly in *Bartz v. Anthropic*[115]—makes this observation more urgent, not less. The doctrinal floor has now been drawn. What fills the space above it is contract. Whether courts will recognize what contract is doing in that space is the open question this Essay has tried to frame.

---

112. *See generally* K. Sabeel Rahman, *The New Utilities*, 39 CARDOZO L. REV. 1621, 1653–65 (2018).

113. *See* Hovenkamp, *supra* note 91, at 1985–90.

114. *See* Frischmann & Lemley, *supra* note 90, at 299–305.

115. *See supra* note 43 and accompanying discussion.