UC Irvine
UC Irvine Law Review

Title
Generative AI and the Limits of Work Product Doctrine

Permalink
https://escholarship.org/uc/item/25z0n36h

Journal
UC Irvine Law Review , 16(2)

ISSN
2327-4514

Author
Singh, Sukhraj

Publication Date
2026-10-02

eScholarship.org Powered by the California Digital Library
University of California

Generative AI and the Limits of Work Product Doctrine

Sukhraj Singh*

* J.D. Candidate, U.C. Irvine, School of Law; B.A. Sociology, 2021, U.C. Irvine. All glory be to God. I would like to thank the U.C. Irvine Law Review team for all their help, and my family, without whom none of this would be possible.

Introduction 676
I. The Work Product Doctrine 677
A. Legal Standard 677
B. *Hickman v. Taylor* & Fed. R. Civ. P. 26(b)(3) 678
C. Purpose of the Work Product Doctrine 679
II. Fairness & Efficiency Dilemma 680
III. *Tremblay v. OpenAI, Inc.* 681
A. Case Background 681
B. Fairness Solution 683
Conclusion 684

INTRODUCTION

The rapid integration of generative artificial intelligence into the legal practice has certainly created new efficiencies along with new uncertainties, particularly around confidentiality, discovery, and the work product doctrine. Given the recent advancements of AI, more specifically Large Language Models (LLMs), it seems the integration of AI into legal practice is inevitable.[1] Due to large law firms already creating their own systems,[2] a natural consequence is that small firms, offices, or individuals that cannot afford to painstakingly create their own secure LLM[3] will rely on OpenAI to even the field.[4] Naturally, issues in regards to the work product doctrine will arise,[5] and judges must decide when the work product doctrine should or should not bend. As judges delve into the issues of work product doctrine, it is important to ground these arguments with the explicit purpose of the work product doctrine: fairness.

In fact, courts are already beginning to confront whether these AI-assisted outputs are protected under traditional doctrines. In *Tremblay v. OpenAI, Inc.*,[6] the district court faced this question directly when plaintiffs who used ChatGPT in pre-suit investigation sought to shield their prompts and outputs as attorney

1. Harry Surden, *ChatGPT, AI Large Language Models, and Law Symposium: The New AI: The Legal and Ethical Implications of ChatGPT and Other Emerging Technologies*, 92 FORDHAM L. REV. 1941, 1972 (2023).

2. Ella Sherman, *Firms Are Building Their Own LLMs to Stay Ahead of Innovation, Competition*, LAW.COM (Mar. 6, 2025), https://www.law.com/legaltechnews/2025/03/06/firms-are-building-their-own-llms-to-stay-ahead-of-innovation-competition-/?slreturn=2025040643426 [perma.cc/53V5-HRTB].

3. Ankwei Chen, *Roll Your Own Bot: Notes on Law Firms Building Their Own Private AI Systems*, LEE, TSAI & PARTNERS (July 24, 2024), https://www.leetsai.com/ltp-special-column/roll-your-own-bot-notes-on-law-firms-building-their-own-private-ai-systems [web.archive.org/web/20240815090130/https://www.leetsai.com/ltp-special-column/roll-your-own-bot-notes-on-law-firms-building-their-own-private-ai-systems].

4. *See* Surden, *supra* note 1, at 1074–76 (discussing the advantages in settlements of firms and repeat institutional players).

5. *Id.* at 1076 ("[L]itigants will seek the other side's machine outputs. What labels did you apply to the seed set? How much does your software say this case is worth? What legal arguments did your software say this case is worth?").

6. Tremblay v. OpenAI, Inc., No. 23-cv-03223-AMO, 2024 WL 3748003 (N.D. Cal. Aug. 8, 2024).

work product. Though a magistrate initially ordered disclosure, the district judge reversed, holding that prompts crafted by counsel reflected mental impressions and were therefore protected as opinion work product under Rule 26(b)(3).

Before discussing why the district judge's holding was correct, this Comment will review the work product doctrine and the judicial intent behind it. This Comment proceeds by discussing the fairness concerns posed by unequal access to private versus public AI tools, then discussing why the *Tremblay* holding was correct. Ultimately, extending the work product protection to AI use is essential to preserving fairness and strategic privacy in a modern legal landscape, and like other litigation tools revealing strategy or judgment, should be treated as opinion work product.

I. THE WORK PRODUCT DOCTRINE

A. Legal Standard

On paper, the work product doctrine can appear simple enough but is often more complex in practice. As an overview, in suits where one party seeks another attorney's work product, first, the party seeking discovery must show that the subject documents or tangible things are relevant to the subject matter involved in the pending litigation and are not privileged.[7] Once such a showing has been made, the burden shifts to the party seeking protection to show that the requested materials were prepared in anticipation of litigation, or for trial by or for, the party or the party's attorney, consultant, surety, indemnitor, insurer, or agent.[8] If a court concludes that the items were prepared in anticipation of litigation, the burden shifts back to the requesting party to show (1) a substantial need for the materials in the preparation of the party's case; and (2) the inability without undue hardship of obtaining the substantial equivalent of the materials by other means.[9] Finally, even if substantial need and inability to obtain without substantial burden are demonstrated, the court must determine whether the materials are factual work product (which is usually subject to discovery), or opinion work product—mental impressions, opinions, and conclusions—which is rarely, if ever, subject to discovery.[10]

The scope of the work product doctrine protection has been recognized as "one of the 'most controversial and vexing problems' in the Federal Rules of Civil Procedure."[11] In fact, the work product protection is one of the most frequently litigated discovery issues,[12] despite guidance provided by *Hickman v. Taylor*[13] and Federal Rules of Civil Procedure 26(b)(3).[14]

---

7. Martin v. Monfort, Inc., 150 F.R.D. 172, 172–73 (D. Colo. 1993).
8. PepsiCo, Inc. v. Baird, Kurtz & Dobson LLP, 305 F.3d 813, 817 (8th Cir. 2002).
9. *See id.*
10. *Martin*, 150 F.R.D. at 173.
11. Jeff A. Anderson, Gina E. Cadieux, George Hays E. & Michael B. Hingerty, *Work Product Doctrine*, 68 CORNELL L. REV. 760, 762 (1983).
12. *Id.* at 763 (citing 4 JAMES WM. MOORE ET AL., MOORE'S FEDERAL PRACTICE ¶ 26.63 (2d ed. 1983)).
13. Hickman v. Taylor, 329 U.S. 495 (1947).
14. FED. R. CIV. P. 26(b)(3).

B. Hickman v. Taylor *& Fed. R. Civ. P. 26(b)(3)*

Work product protection is generally not absolute, unlike attorney-client privilege. Most work products receive only conditional protection, meaning the court may order disclosure if "good cause" is shown.[15] It is crucial, however, to note that when a work product reflects the attorney's mental impressions, conclusions, opinions, or legal theories, the protection is much greater than for other work products.[16] When determining whether to compel production of work product or not, courts have recognized work products as either "fact work product" or "opinion work product."[17]

The attorney work product doctrine protects trial-preparation materials that reveal an attorney's strategy, intended lines of proof, evaluation of strengths and weaknesses, and inferences drawn from interviews.[18] The work product protection is designed to preserve the privacy of attorneys' thought processes and to prevent parties from borrowing the wits of their adversaries.[19] Work product protects "documents and tangible things that are prepared in anticipation of litigation or for trial" by or on behalf of a party.[20] The scope of the work product doctrine, however, is broader than Fed. R. Civ. P. 26(b)(3) because it encompasses "intangible" as well as tangible work products.[21]

On paper, the distinction between the two is simple enough. Fact work product consists of documents prepared by an attorney that do not contain the attorney's mental impressions, conclusions, or opinions of the attorney.[22] Fact work product can be discovered upon a showing of both "a substantial need and an inability to secure the substantial equivalent of the materials by alternate means without undue hardship."[23] On the other hand, opinion work product is one where an attorney's mental impressions, legal strategy, intended lines of proof, evaluation of the strengths and weaknesses of the attorney's case, and inferences drawn from interviews of witnesses.[24] Opinion work products have been repeatedly recognized as "virtually undiscoverable."[25]

As one can imagine, in practice there can be difficulty in distinguishing between the two. "[T]he task of drawing a line between what is fact and what is

---

15. *Hickman*, 392 U.S. at 512.

16. *See* Upjohn Co. v. United States, 449 U.S. 383, 400 (1981).

17. United States v. Clemens, 793 F. Supp. 2d 236, 239–40 (D.D.C. 2011); *In re* Doe, 662 F.2d 1073, 1076 (4th Cir. 1981).

18. *See* FED. R. CIV. P. 26(b)(3); *see Hickman*, 392 U.S. at 510–11.

19. *Hickman*, 392 U.S. at 510–11.

20. FED R. CIV. P. 26(b)(3)(A).

21. *See Hickman*, 392 U.S. at 511.

22. *See In re* Grand Jury Proceedings #5 Empanelled Jan. 28, 2004, 401 F.3d 247 (4th Cir. 2005).

23. *In re* Doe, 662 F.2d at 1080.

24. Sporck v. Peil, 759 F.2d 312, 316 (3d Cir. 1985); Williamson v. Moore, 221 F.3d 1177, 1182 (11th Cir. 2000) (citing *Hickman*, 329 U.S. 495).

25. Dir., Off. of Thrift Supervision v. Vinson & Elkins, LLP, 124 F.3d 1304, 1307 (D.C. Cir. 1997); United States v. Deloitte LLP, 610 F.3d 129, 135 (D.C. Cir. 2010) (citing *Dir., Off. of Thrift Supervision*); Eastman v. Thompson, 636 F. Supp. 3d 1078, 1087 (C.D. Cal. 2022); *see* Republic of Ecuador v. Mackay, 742 F.3d 860, 869 (9th Cir. 2014).

opinion can at times be frustrating and perplexing."[26] To add to the confusion, courts have held that "an attorney can take certain actions that will, in essence, transform 'opinion' work product into 'fact' work product."[27]

The work product doctrine is not absolute; there are certainly ways to overcome the protection. When a party requests work product from an attorney, we look to Fed. R. Civ. P. 26(b)(3) in order to determine whether or not a party may discover documents and tangible things of another party prepared in anticipation of litigation or for trial.[28] A party must show (1) a "substantial need" of the materials to prepare the case and (2) inability to obtain a substantial equivalent without "undue hardship."[29]

Also, a party can waive the work product protection when "protected materials are disclosed in a way that 'substantially increases the opportunity for potential adversaries to obtain the information.'"[30] It is important to note for our discussion that courts have found that sharing material with vendors, contractors, or litigation support firms under confidentiality agreements does not waive work product protection.[31]

C. Purpose of the Work Product Doctrine

It is at this point, one may begin wondering what the point of such a cumbersome doctrine is, and the answer to this lies in *Hickman*. The *Hickman* Court noted that a lawyer be allowed a certain degree of privacy in order to "promote justice and to protect their clients' interests."[32] The Court stated if material used to prepare for trial were open to opposing counsel, then lawyers would begin refraining from writing things down.[33] This would have a demoralizing effect on the legal profession, leading to inefficient and *unfair* practices.[34] The Court in *U.S. v. Nobles* further notes that "[T]he [work product] doctrine is an intensely practical one, grounded in the realities of litigation in our adversary system."[35] As we can see, the purpose of the doctrine is crucial in our application of the work product doctrine to OpenAI. The concepts of fairness and efficiency are stressed throughout the early discussion of the work product doctrine. With this in mind, let us assess a dilemma that may arise, and then assess not only how the holding of *Tremblay* is the correct holding but also how the holding closely adheres to the purposes of the work product doctrine.

---

26. Fla. House of Representatives v. U.S. Dep't of Com., 961 F.2d 941, 947 (11th Cir. 1992).
27. United States v. Clemens, 793 F. Supp. 2d 236, 253 (D.D.C. 2011).
28. Upjohn Co. v. United States, 449 U.S. 383, 400 (1981).
29. *Id.*
30. Stern v. O'Quinn, 253 F.R.D. 663, 681 (S.D. Fla. 2008) (citing Niagara Mohawk Power Corp. v. Stone & Webster Engineering Corp., 125 F.R.D. 578, 587 (N.D.N.Y. 1989)).
31. *See Deloitte LLP*, 610 F.3d at 129.
32. *Hickman*, 329 U.S. at 510–11.
33. *Id.* at 511.
34. *See id.* (emphasis added).
35. United States v. Nobles, 422 U.S. 225, 238 (1975).

II. FAIRNESS & EFFICIENCY DILEMMA

Given the prevalence of AI in the legal practice,[36] courts have the power to drastically sway the scales of balance. The legal profession's, and more specifically large law firms', embrace of LLMs has created a dichotomy in how attorneys access AI. Many leading law firms are developing or deploying *private* AI systems, proprietary LLMs, or licensed AI platforms, behind their own firewalls.[37] In early 2023, the global firm Allen & Overy announced it would roll out a GPT-4-based legal assistant to 3,500 of its lawyers for tasks like document drafting and research.[38] Other major firms have followed suit: DLA Piper, Orrick, and Fisher Phillips are among those integrating GPT-4 legal tools into their practice, while firms like Holland & Knight and Baker McKenzie have begun building in-house AI applications tailored to their needs.[39] These private LLM deployments allow big-firm attorneys to leverage AI's efficiency gains within a closed environment. And this is where the issue arises.

When an AI system is developed and controlled internally, any outputs or queries remain within the firm's protected domain, much like other internal research files. In such settings, there is little question that the attorney's use of the AI or the attorney's account falls under work product protection. These AI systems are simply another tool under the firm's direct control, analogous to an internal research database or software. No opponent would contend that a law firm must disclose its *private* LLM's settings or reveal the inputs its attorneys posed, just as one would not expect to compel production of a firm's internal memorandum bank or litigation management software. In short, large firms' use of private AI does not generally raise novel discovery issues because the AI operates as an extension of the firm's own knowledge infrastructure.

By contrast, attorneys without the resources of big law firms would have to rely on public generative AI tools (such as OpenAI's ChatGPT) to obtain similar benefits in research and drafting.[40] Solo practitioners, small law firms, and one-time

---

36. *See* Majorite Richter, *How AI Is Transforming the Legal Profession*, THOMSON REUTERS: BLOG (Aug. 18, 2025), https://legal.thomsonreuters.com/blog/how-ai-is-transforming-the-legal-profession/ [perma.cc/RHN4-7YZG].

37. Sara Merken, *Legal AI Race Draws More Investors as Law Firms Line Up*, REUTERS (Apr. 26, 2023), https://www.reuters.com/legal/legal-ai-race-draws-more-investors-law-firms-line-up-2023-04-26/ [perma.cc/4Q2L-MHLV] (reporting that Allen & Overy has adopted the Harvey AI platform firm-wide, with 3,500 attorneys and staff slated to use it for generating legal documents and conducting research tasks).

38. *Id.*; *see also* Mohar Chatterjee, *How Lawyers Use AI*, POLITICO (May 23, 2023), https://www.politico.com/newsletters/digital-future-daily/2023/05/23/how-lawyers-use-ai-00098417 [web.archive.org/web/20230731230905/https://www.politico.com/newsletters/digital-future-daily/2023/05/23/how-lawyers-use-ai-00098417] (noting Allen & Overy's partnership with an AI startup to integrate GPT-based tools into the firm's practice).

39. *See* Merken, *supra* note 37 (noting that DLA Piper and other large firms agreed to use Casetext's CoCounsel (GPT-4) tool, and that firms such as Orrick and Fisher Phillips have also begun using CoCounsel; further noting Holland & Knight's development of a custom AI for contract review and Baker McKenzie's pilot programs with LLM technology).

40. *See* Laura Lorek, *AI Legal Innovations, The Benefits and Drawbacks of Chat-GPT and Generative AI in the Legal Industry Artificial Intelligence and the Practice of Law*, 50 OHIO N.U. L. REV.,

or infrequent litigants face a potential double bind: if their use of an external AI tool is deemed discoverable by an opposing party, they risk exposing their litigation strategy. If, however, they were to refrain from using such tools out of caution, they operate at an efficiency disadvantage.

The policy implications here are significant. Treating an attorney's OpenAI or other public AI accounts and account activity as broadly discoverable would create a serious competitive imbalance in litigation. Big law firms can afford dedicated AI systems whose workings remain shielded while smaller actors would be uniquely vulnerable to intrusive discovery into their trial preparation. It would be fundamentally unfair for courts to compel a solo attorney to turn over their OpenAI records, or outright forgo the usage of AI, when no one would suggest compelling a firm to hand over the keys to its privately developed AI or research software. Such a disparity would undermine the level playing field that the work product doctrine endeavors to maintain.

Commentators have accordingly warned that overbroad discovery or disclosure requirements regarding AI usage could chill lawyers' willingness to use these valuable tools, disproportionately harming those who need them most for affordable legal services and access to justice. If plaintiffs and small-firm attorneys must assume that anything they ask a public OpenAI tool, or even their entire account, might later be turned over to an opposing party, many will simply forgo using the technology. The end result would be to the advantage of those litigants who can internally absorb AI's functions, while denying opposing parties of the efficiency gains that AI can provide. Such an outcome is directly at odds with the work product doctrine's fairness-centric purpose. The doctrine was never intended to allow savvy litigants to exploit differences in technology or resources; rather, its aim was to ensure each side can prepare its case with minimal fear that its investigative efforts will be laid bare to the other side.

III. *TREMBLAY V. OPENAI, INC.*

A. Case Background

In June 2023, a group of authors led by Paul Tremblay filed a putative class action in the Northern District of California against OpenAI, alleging that OpenAI's ChatGPT was trained on their copyrighted books without permission.[41] Plaintiffs are authors of books who alleged that "their books were used to train OpenAI language models that operate the artificial intelligence ('AI') software

492, 513 (2024) ("Generative AI has many benefits in the legal industry, including increased efficiency, greater accuracy, innovations, and more.").

41. Tremblay v. OpenAI, Inc., No. 23-cv-03223-AMO (N.D. Cal. filed June 28, 2023). *Tremblay* has now been consolidated with Silverman v. OpenAI, Inc., No. 4:23-cv-03416-KAW (N.D. Cal. filed July 7, 2023), and Chabon v. OpenAI, Inc., No. 4:23-cv-04625-AMO (N.D. Cal. filed Sept. 8, 2023), into *In re* OpenAI ChatGPT Litigation, No. 23-cv-03223. *See* Pretrial Order No. 1, Tremblay v. OpenAI, Inc., No. 3:23-cv-03223-AMO (N.D. Cal. Nov. 9, 2023).

ChatGPT."[42] More specifically, OpenAI copied Plaintiffs' copyrighted books and used them in its training dataset (a set of examples used to train machine learning models which enables them to learn patterns and make predictions).[43] OpenAI moved to dismiss most of the claims, and the district court granted the motion in part while allowing the core copyright claim to proceed.[44]

The importance of this case for our discussion is what occurred during discovery. There are two of OpenAI's requests that are particularly noteworthy for our discussion. One of the things OpenAI asked for was the OpenAI account information for individuals who used ChatGPT to investigate Plaintiffs' claim.[45] The second thing OpenAI asked for was the prompts and outputs for Plaintiffs' testing of ChatGPT in connection with their pre-suit ChatGPT testing.[46] This included prompts and outputs that did not reproduce or summarize Plaintiffs' works or otherwise support Plaintiffs' claims, along with documentation of Plaintiffs' testing process.[47] OpenAI argued that Plaintiffs waived work product protection over the account and prompt information, claiming, among other reasons, the totality of the data would be needed for OpenAI to subject Plaintiffs' claims to meaningful adversarial testing.[48] As for the specific account information, OpenAI claimed the "custom instructions" feature allows users to add specific preferences for ChatGPT to consider when it generates responses.[49]

Plaintiffs asserted this requested material was shielded from discovery as it was an attorney work product.[50] Plaintiffs stated they agreed to submit prompts and outputs that support their claim, but the material not disclosed divulged "thoughts and analysis conducted by Plaintiffs' lawyers in preparation for litigation" and thus were shielded.[51]

The court ruled that Plaintiffs' prompts, outputs, and settings from their pre-suit testing were discoverable and granted OpenAI's motion to compel discovery.[52] The court held that the materials were work products because they are "more in the nature of bare facts," and by placing the test results in their complaint, Plaintiffs waived the ability to assert work product doctrine.[53] The court also noted that OpenAI could not generate reliable test results "without knowing the account settings used by Plaintiffs."[54]

---

42. Tremblay v. OpenAI, Inc., 716 F. Supp. 3d 772, 775 (N.D. Cal. 2024).
43. *Id.* at 776.
44. *Id.* at 783.
45. Tremblay v. OpenAI, Inc., No. 23-cv-03223-AMO, 2024 WL 3159292, at *1 (N.D. Cal. June 24, 2024).
46. *Id.*
47. *Id.*
48. *Id.* at *2.
49. *Id.*
50. *Id.*
51. *Id.*
52. *Id.* at *4.
53. *Id.* at *3.
54. *Id.*

Plaintiffs objected to the magistrate's ruling and moved for relief pursuant to Rule 72(a).[55] The district court disagreed with the characterization of the requested material as work product because "the ChatGPT prompts were queries crafted by counsel and contain counsel's mental impressions and opinions about how to interrogate ChatGPT."[56] The district court would go on to note that in regards to waiving opinion work product "only where 'mental impressions are at issue in a case and the need for the material is compelling.'"[57] Neither the order nor OpenAI showed that the mental impressions were at issue or that OpenAI had a compelling need for the material.[58] Thus, because the requested material was opinion work product, the district court granted Plaintiffs' motion for relief and denied OpenAI's request to compel production.[59] With the case and holding discussed and contextualized, let us now look at the prevalence of AI in the law and the issues that will arise in regards to the work product doctrine and see how this holding can be adopted to solve that problem.

B. Fairness Solution

To avoid unfairness and promote efficiency, courts should follow the *Tremblay* court's lead and find AI-assisted work to constitute opinion work product. One of the ways courts confronting the discoverability of AI-assisted work products could analogize to existing protections for attorney-controlled tools and research materials. An instructive comparison can be made with litigation databases and other knowledge-management systems attorneys use. In *Jaffe v. Household International*, for example, in-house counsel maintained a private electronic database tracking pending lawsuits, which included attorney notes and evaluations of each case. The court held that this database, created "to assist [the company's] counsel in understanding, managing and providing legal advice" on litigation, was protected by attorney-client privilege and work product doctrine, rejecting attempts to label it a mere business record.[60] That decision underscores that when attorneys gather information and analyze cases using an internal platform, the compiled result is shielded if it embodies counsel's impressions or was prepared for litigation purposes.

The fundamental principle here is that the *process* by which an attorney prepares their case is not ordinarily subject to an opponent's scrutiny, and that should not change with AI. We do not require attorneys to turn over their Westlaw

---

55. Tremblay v. OpenAI, Inc., No. 23-cv-03223-AMO, 2024 WL 3748003, at *1 (N.D. Cal. Aug. 8, 2024).

56. *Id.* at *2 (citing Republic of Ecuador v. Mackay, 742 F.3d 860, 869 (9th Cir. 2014)).

57. *Id.* (citing United States v. Sanmina Corp., 968 F.3d 1107, 1124–25 (9th Cir. 2020)).

58. *Id.* at *3.

59. *See id.*

60. Lawrence E. Jaffe Pension Plan v. Household Int'l, Inc., 237 F.R.D. 176, 179–85 (N.D. Ill. July 6, 2006) (finding an internal "litigation database" containing attorneys' comments, impressions, and legal strategies was prepared for the purpose of managing pending litigation and thus protected by privilege and work product, noting the database had never been shared outside the legal team). The court rejected the argument that use of the database for business auditing purposes stripped it of protection, emphasizing that its primary purpose was to assist counsel in advising on and defending litigation.

or Lexis query histories,[61] their personal notes on which leads they found fruitful, nor the identities of people they chose *not* to depose, unless a rare showing of necessity or waiver is made. Requiring disclosure of an attorney's ChatGPT prompts or outputs would be analogous to invading the attorney's private research file or forcing the production of a proprietary search algorithm's result. So, unless the content of those AI interactions itself is directly at issue in the litigation, as stated in *Tremblay*, such demands should be met with the same skepticism courts have shown toward fishing expeditions into counsel's thought processes.

The work product doctrine, as numerous courts have affirmed, is meant to keep attorneys' preparation privileged absent extraordinary justification.[62] Extending the work product protection to AI-assisted work products not only fits comfortably within existing doctrine but also avoids an unreasonable scenario in which technological innovation becomes a loophole to exploit adversaries' diligence. In practice, treating counsel's AI interactions as protected work products will encourage lawyers of all stripes to responsibly utilize these advanced tools to serve their clients, without fear that doing so will later hand an opponent the roadmap of their case.

CONCLUSION

The emergence of generative artificial intelligence as a tool in legal practice requires courts to reaffirm and carefully adapt longstanding doctrines of attorney work product. The foundational purpose of the work product doctrine, ensuring fairness and efficiency, remains critically important in an era where attorneys increasingly interact with AI platforms to conduct pre-suit investigations, assess claims, and develop legal theories. To compel disclosure of such materials, particularly from resource-constrained plaintiffs or small firms using public LLMs like ChatGPT, would not only erode doctrinal protections but also exacerbate existing inequities in the legal system. The court in *Tremblay* ruled correctly and in line with the purposes of the work product doctrine and should serve as a guiding influence on other cases moving forward.

61. *See* Brooks v. Thomson Reuters Corp., No. 21-cv-01418-EMC, 2022 WL 16698686, at *3 (N.D. Cal. Nov. 3, 2022).

62. Hickman v. Taylor, 329 U.S. 495, 510–511 (1947).