Use of AI Tools Impacts Confidentiality, Privilege & Attorney Work-Product: Concerns for Attorneys In the Wake of US v. Heppner

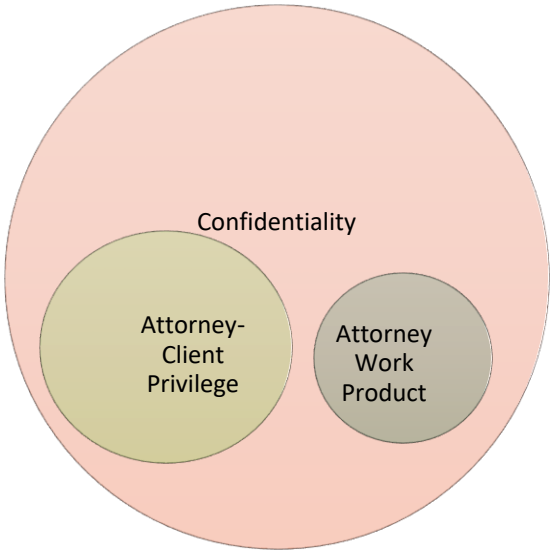

As illustrated in the Venn diagram, confidentiality is a broad ethical duty owed by lawyers and law firms to all clients not to disclose the subject matter of representation nor the identity of the client to any third parties, without that client's informed consent. The duty of confidentiality is outlined at ABA Model Rule 1.6. Confidentiality encompasses two narrower and more specific classes of verbal and written communications legally protected against disclosure in discovery, in court or before a tribunal, namely those communications and documents protected by attorney-client privilege, a statutory right belonging to the client, who may either invoke the privilege or waive it, and those deemed attorney work-product, prepared by or for the attorney in anticipation of litigation. In New York state, for example, privileged attorney-client communications are defined broadly at NY CPLR 4503 and include legal advice and preparation for administrative actions and hearings as well as trials, so it clearly covers most immigration matters, and CPLR 4548 extends that privilege to electronic communications; NY CPLR 3101(c)

extends absolute privilege against disclosure to materials created by an attorney that reveal their legal research, analysis, strategy, or theory of the case. All privileged attorney-client communications are confidential, but not all confidential communications are privileged.

In immigration law, where dual representation is the norm rather than the exception, we are often confronted with questions of who holds the privilege, and what kinds of documents prepared in an immigration context are privileged, which may depend in turn on state law. The application of state law is complicated in this federal practice area by questions of where is the employer, where is the beneficiary, and where is the law firm, so which state's law governs?

This landscape changed dramatically in February 2026 with the US v. Heppner decision (1:25-cr-00503-JSR), SDNY, Feb. 17, 2026. A criminal defendant charged with wire & securities fraud, false statements to auditors, and falsification of records in connection with a scheme to defraud investors fed questions about his fact pattern into Claude, a public AI platform, before conveying those facts and the resulting AI analysis to his attorney. Judge Rakoff found that the public AI platform was a third party, and that as a non-human, it was incapable of fiduciary duty. Both the defendant's fact pattern and the AI analysis of it were not, and could not be, covered by attorney-client privilege based on his communicating them to counsel after the fact, because he had already disclosed them to a third party (Claude, the AI platform), thus waiving confidentiality. Subsequent use of Claude's analysis by counsel in trial preparation could not retroactively turn the AI documents into protected attorney work-product. The court found that privilege cannot attach retroactively. The court also held that the attorney work product doctrine could not attach to AI documents because they had been created independently by the defendant, not prepared at the direction of counsel. So, to recap:

- A public AI platform is a third party with no fiduciary capacity: disclosure waives confidentiality & privilege
- Privilege cannot attach retroactively to anything already disclosed to 3rd parties
- Material already shared with a 3rd party does not become protected work product when an attorney uses it in preparing for a legal matter

Lessons moving forward: if a client, or a prospective client, feeds the fact pattern of a case into a public AI platform, that information is no longer confidential, and the client cannot assert privilege with respect to that fact pattern. A much more alarming supervisory and liability problem arises if a paralegal or junior associate in a law firm feeds a client's fact pattern into a public AI platform.

Best practice going forward is that for attorney-client privilege to survive in cases where any LLM or generative AI tools are used in legal research, analysis, writing or any steps in case preparation, all such AI use must be directed by counsel under the terms of engagement, using only enterprise-licensed AI tools operating in a closed ecosystem, with an express expectation of confidentiality, and supporting documentation of the contractual limits on the AI tool's training use, retention, disclosure, access and security controls.

Any attorney who has had a consultation in the past year is well aware that this concern may arise before the law firm is retained, as prospective clients may use public AI tools to frame their questions for counsel prior to initial consultation, and to craft detailed follow-up questions after consultation. To that end, I offer the following warning to prospective clients:

"Thank you for putting your trust in [Law Firm]. While our initial consultation is privileged and confidential, it is important for you to be aware that those protections are fragile, and can easily be destroyed by feeding your fact pattern, our legal advice, or any portions or combination thereof into a public AI platform. While the consultation fee covers one meeting of up to an hour and one email for follow-up questions, we urge you not to use any public AI platform (ChatGPT, Gemini, Claude, etc.) to draft your initial or follow-up questions, as doing so may waive both attorney-client privilege and confidentiality."