stuartteicher esq
I teach lawyers

CLE Performer & Ethics Educator
stuart.tpg@gmail.com
www.stuartteicher.com
732-522-0371

Hook, Line, and Stinker: Phishing, Deepfakes and Other Cybersecurity Concerns for Lawyers Written Materials

I. Cybersecurity, Deepfakes, Phishing, and Supervision Under Rule 5.3

Rule 5.3 - Responsibilities Regarding Nonlawyer Assistance

Rule 5.3 establishes clear obligations for lawyers who employ or work with nonlawyer assistants. The Rule provides that a lawyer having direct supervisory authority over a nonlawyer must make reasonable efforts to ensure that the person's conduct is compatible with the professional obligations of the lawyer. Partners and lawyers with comparable managerial authority in a law firm must make reasonable efforts to ensure that the firm has measures in effect giving reasonable assurance that the nonlawyer's conduct is compatible with the professional obligations of the lawyer.

The Evolving Threat Landscape

Modern phishing attacks have moved far beyond the easily identifiable emails of the past. Today's sophisticated social engineering schemes specifically target law firms because they handle valuable client information, maintain trust accounts with significant funds, and serve as gatekeepers to confidential business transactions. Phishing attacks against law firms increased significantly over the past several years.

These attacks frequently target nonlawyer staff members who may lack the same level of security awareness as attorneys. Paralegals, legal secretaries, administrative assistants, and IT personnel all represent potential entry points for cybercriminals. A successful phishing attack on any staff member can compromise the entire firm's network and client data.

Deepfake Technology and Impersonation Risks

Deepfake technology represents a particularly insidious evolution in social engineering attacks. Using artificial intelligence, malicious actors can now create convincing video or audio impersonations of senior partners, clients, or other trusted individuals. In 2024, a multinational company reported that a staff member transferred $25 million from the company's account after participating in a video conference call with a deepfake impersonations.

The technology required to create convincing deepfakes has become increasingly accessible. Software that once required significant technical expertise can now be deployed by individuals with minimal training. Public videos from law firm websites, conference presentations, and social media provide ample source material for creating convincing impersonations of lawyers and their clients.

Supervision Obligations in the Context of Cybersecurity

Rule 5.3 requires more than passive oversight. The obligation to "make reasonable efforts" means that lawyers must actively implement systems and training to protect against cybersecurity threats that could compromise client information or firm resources. This includes ensuring that nonlawyer assistants can recognize and respond appropriately to phishing attempts and social engineering schemes.

Reasonable supervision in the cybersecurity context includes several key components. First, firms must provide regular training to all staff members on current threat vectors. Training conducted once during onboarding proves insufficient given the rapidly evolving nature of cybersecurity threats. Quarterly or semi-annual refresher training better satisfies the reasonable efforts standard.

Second, firms must implement technical controls that reduce the risk that a single successful phishing attack can compromise the entire system. Multi-factor authentication, email filtering systems, and restrictions on staff access to sensitive data all serve as reasonable measures under Rule 5.3. A lawyer cannot delegate cybersecurity entirely to an IT department and claim to have satisfied supervisory obligations when that delegation leaves obvious vulnerabilities. Lawyers need to retain responsibility and accountability.

Third, firms must establish clear protocols for verifying unusual requests, particularly those involving fund transfers or the disclosure of confidential information. The deepfake scenario described above succeeded because the staff member had no established protocol for independently verifying instructions to transfer funds, even when those instructions came through unusual channels.

Practical Implementation

A reasonable supervision system under Rule 5.3 should include specific procedures for staff to follow when they receive requests that involve client confidences, firm resources, or sensitive data. These procedures might include requirements that staff members independently verify requests through a separate communication channel before acting on them. For example, if a paralegal receives an email that appears to come from a partner requesting the immediate transfer of client trust funds, the protocol should require the paralegal to call the partner using a known phone number rather than replying to the email or using contact information contained in the message.

Firms should also implement technical measures that flag potential phishing attempts. Email systems can be configured to warn users when messages originate from outside the organization but display names that match internal staff members. This simple measure can prevent the scenario where an attacker creates an email account that displays the name "Sarah Johnson, Partner" but uses an external email address.

The reasonable efforts standard under Rule 5.3 scales with firm size and resources. A solo practitioner working with a single legal assistant faces different practical requirements than a firm with hundreds of nonlawyer employees. However, the fundamental obligation remains the same: lawyers must take affirmative steps to ensure their nonlawyer assistants maintain security practices compatible with the lawyer's professional obligations.

Violations and Disciplinary Consequences

Failures in supervision that lead to data breaches or unauthorized disclosures can result in disciplinary action. Disciplinary authorities could find Rule 5.3 violations where lawyers fail to implement basic security measures and nonlawyer staff members subsequently fall victim to phishing schemes that compromises client information. The fact that the breach might occur through a staff member rather than a lawyer directly does not insulate the lawyer from responsibility.

Supervisors bear responsibility not only for implementing systems but also for monitoring compliance with those systems. A firm policy requiring multi-factor authentication provides no protection if the lawyer never verifies that staff members actually use the system correctly. Regular security audits, even informal ones, help demonstrate the reasonable efforts required by Rule 5.3.

II. Facial Recognition, Imposters, and the Competence Requirement of Rule 1.1

Rule 1.1 - Competence

Rule 1.1 requires that a lawyer provide competent representation to a client, which demands the legal knowledge, skill, thoroughness, and preparation reasonably necessary for the representation. Comment 8 to the Rule states that maintaining competence requires keeping abreast of changes in the law and its practice, including the benefits and risks associated with relevant technology.

New York RPC 1.1 includes this technology competence requirement in Comment 8, matching the ABA Model Rule. New Jersey RPC 1.1 does not include an explicit comment regarding technology competence. However, the Supreme Court of New Jersey has indicated through its recent CLE requirement focused on technology-related subjects that it considers technology competence an aspect of overall professional competence even without a specific comment to the Rule.

The Imposter Problem in Legal Practice

Identity verification has become a critical competence issue for lawyers as cybercriminals have developed increasingly sophisticated methods of impersonating clients, opposing counsel, and other parties to legal proceedings. The traditional reliance on email addresses, phone numbers, and even video calls no longer provides adequate assurance that the person communicating with the lawyer is who they claim to be.

Imposter schemes targeting lawyers typically follow several patterns. In client impersonation schemes, criminals pose as prospective clients seeking legal services. They engage the lawyer in what appears to be legitimate representation, often involving real estate transactions, estate planning, or business formations. The imposter provides false identification documents and induces the lawyer to receive and forward funds that are actually proceeds of fraud or money laundering schemes.

In business email compromise schemes, attackers gain access to email accounts and monitor communications between lawyers and clients. They then insert themselves into ongoing transactions at critical moments, sending fraudulent wire transfer instructions that appear to come from legitimate parties. These schemes have cost lawyers and their clients hundreds of millions of dollars.

Facial Recognition Technology and Its Limitations

Some lawyers have turned to facial recognition software and video conferencing as a solution to imposter problems. The logic appears sound: requiring a video call with a client provides greater assurance of identity than email or phone contact alone. However, facial

recognition technology carries significant limitations and risks that lawyers must understand to maintain competence under Rule 1.1.

Facial recognition systems can be defeated through various means. Deepfake technology can create video presentations of individuals that appear authentic even to sophisticated facial recognition software. Attackers can also use photographs or videos of real individuals, presenting them to cameras in ways that fool both human observers and recognition software.

More problematically, facial recognition technology may create a false sense of security. Lawyers who believe they have verified client identity through a video call may relax other verification procedures that would catch imposter schemes. This creates what cybersecurity experts call "single point of failure" vulnerability. When all security depends on one measure, defeating that measure compromises the entire system.

Competent Identity Verification Practices

The competence requirement of Rule 1.1 demands that lawyers understand both the capabilities and limitations of identity verification technologies. This understanding must inform the development of verification procedures that appropriately balance security needs with practical considerations.

A competent approach to identity verification uses multiple independent factors. For example, a lawyer representing a client in a real estate transaction might require the client to provide government-issued identification, verify the client's identity through a video call, and independently confirm the client's phone number and address through public records or third-party databases. No single measure provides certainty, but multiple independent verification points significantly reduce imposter risk.

Lawyers must also recognize situations that present heightened imposter risk. Representations that involve fund transfers to third parties, particularly when those transfers occur quickly or to international recipients, demand enhanced scrutiny. Unsolicited inquiries from prospective clients who found the lawyer through internet searches rather than personal referrals warrant additional verification steps. Clients who resist normal verification procedures or create urgency around transactions should trigger skepticism rather than accommodation.

Geographic and Practice Area Considerations

Lawyers practicing in areas with high rates of cyber fraud must implement more rigorous identity verification procedures to satisfy Rule 1.1's competence standard. Real estate attorneys, estate planning practitioners, and business lawyers handling transactions all face significant imposter risks that require specific protective measures.

Similarly, lawyers who represent clients remotely or who maintain virtual law practices face unique verification challenges. The absence of in-person meetings removes a traditional safeguard against imposter schemes. Competent practice in this context requires developing alternative verification methods that provide similar assurance without face-to-face contact.

Continuing Education and Technology Competence

Maintaining competence under Rule 1.1 requires ongoing education about evolving cybersecurity threats and available protective technologies. The technology landscape changes rapidly. Security measures that provided adequate protection two years ago may prove insufficient today. Lawyers cannot rely on verification procedures implemented years ago without periodically reviewing and updating those procedures in light of new threats and available technologies.

III. Casual Communication, Rule 1.4, and Increased Cybersecurity Risks

Rule 1.4 - Communication

Rule 1.4 requires lawyers to keep clients reasonably informed about the status of their matters and to promptly comply with reasonable requests for information. The Rule also mandates that lawyers explain matters to the extent reasonably necessary to permit clients to make informed decisions regarding the representation.

The Evolution of Legal Communication Norms

Legal communication has undergone significant transformation over the past decade. Where formal written correspondence once dominated lawyer-client interactions, text messages, instant messaging applications, and informal email have become commonplace. Younger lawyers and clients particularly favor these more casual communication channels, viewing them as more efficient and accessible than traditional methods.

This shift toward informal communication creates a paradox. While these channels may facilitate more frequent communication between lawyers and clients, potentially improving compliance with Rule 1.4's requirement to keep clients reasonably informed, they simultaneously increase cybersecurity risks in ways that may ultimately undermine the communication objectives the Rule serves.

Security Vulnerabilities in Casual Communication Channels

Text messages and instant messaging applications present several cybersecurity concerns that lawyers must understand to maintain both effective communication and client confidentiality. Many popular messaging platforms do not provide end-to-end encryption, meaning that messages travel across networks in forms that third parties can intercept and read. Even platforms that offer encryption may not enable it by default, leaving security dependent on user configuration choices that lawyers and clients may not make correctly.

SMS text messages, despite their ubiquity, provide minimal security. These messages traverse multiple networks between sender and recipient, creating numerous points where interception can occur. Cybercriminals can exploit vulnerabilities in cellular networks to intercept text messages. Foreign intelligence services routinely monitor text message traffic. Yet lawyers frequently use text messages to communicate information about client matters, including sensitive case developments and strategic decisions.

The casual nature of informal communication channels creates additional security risks beyond the inherent technical vulnerabilities. Lawyers and clients often communicate through these channels using personal devices that may lack adequate security

protections. Personal smartphones frequently contain a mix of professional and personal applications, expanding the attack surface that cybercriminals can exploit. A security breach in a personal gaming application or social media account can provide access to legal communications if they exist on the same device.

Phishing Through Informal Channels

The casual communication style prevalent in text messages and instant messaging makes these channels particularly vulnerable to phishing and social engineering attacks. Email phishing schemes have become relatively easy to identify because they often violate the formal communication norms that lawyers expect in professional email. Obvious grammatical errors, strange formatting, and inappropriate informality all serve as warning signs in email communications.

Text message phishing, sometimes called "smishing," exploits the opposite dynamic. Because text messages normally use abbreviated language and informal tone, phishing messages blend seamlessly with legitimate communications. A text message reading "Need case update ASAP" could come from a client, a partner, or a cybercriminal who has spoofed the sender's phone number. The informal communication style provides no basis for distinguishing legitimate from fraudulent messages.

This vulnerability extends to impersonation schemes where attackers pose as lawyers to communicate with clients. A cybercriminal monitoring case filings might identify parties to litigation and send text messages posing as their attorney. The informal nature of text communication provides cover for these schemes. Clients expect their lawyers to text briefly and informally, so an imposter can adopt that style without raising suspicion.

Rule 1.4 Compliance in the Context of Insecure Channels

The intersection of Rule 1.4's communication requirements and cybersecurity concerns creates a tension that lawyers must navigate carefully. Rule 1.4 mandates reasonably prompt communication with clients, but using insecure channels to achieve that promptness may violate other ethical obligations, particularly the duty to maintain client confidentiality under Rule 1.6.

Lawyers satisfy Rule 1.4 by establishing clear communication protocols with clients at the outset of representation. These protocols should address both the channels that will be used for communication and the types of information that can be shared through each channel. For example, a lawyer might agree to use text messages to confirm meeting times and send brief case status updates, while reserving detailed legal analysis and confidential client information for more secure communication methods.

The reasonableness standard embedded in Rule 1.4 requires lawyers to consider both client preferences and security requirements when selecting communication channels. Some clients may strongly prefer text message communication and resist using more secure alternatives. However, client preference does not eliminate the lawyer's professional judgment obligation. A lawyer must explain the security risks associated with different communication channels and make reasonable recommendations, even if the client ultimately chooses a less secure option.

Verification and Authentication in Casual Communication

The informal nature of modern communication channels demands enhanced authentication procedures to ensure that lawyers are actually communicating with clients rather than imposters. Before acting on instructions received through text message or instant messaging, lawyers should verify the communication through an independent channel. This practice protects both lawyer and client while maintaining compliance with Rule 1.4's communication requirements.

Authentication becomes particularly critical when clients send instructions through text messages or other informal channels that could have significant consequences. A text message appearing to come from a client requesting a settlement offer or agreeing to a plea bargain should trigger verification through a phone call or secure email before the lawyer acts on the instruction. While this additional step may delay communication slightly, it prevents potentially catastrophic errors that could result from acting on fraudulent instructions.

Educating Clients About Secure Communication

Rule 1.4's requirement to explain matters to clients necessarily includes explaining the security implications of different communication methods. Lawyers cannot assume that clients understand the risks associated with text messages, instant messaging, or other informal communication channels. Part of reasonably informing clients under Rule 1.4 involves explaining these risks and helping clients make informed decisions about communication methods.

This educational obligation does not require lawyers to deliver lectures on network security protocols. Rather, lawyers should provide practical guidance that helps clients understand risks in accessible terms. For example, a lawyer might explain that text messages work like postcards rather than sealed letters—anyone who handles them along the way can read them. This explanation helps clients make informed decisions about what information they share through different channels.

Firm Policies and Communication Standards

Consistent with the rules on supervision (Rule 5.1(a) for lawyers and Rule 5.3(a) for nonlawyers) Law firms should establish clear policies regarding communication channels and the types of information that can be shared through each channel. These policies help ensure consistency across the firm and provide guidance to lawyers who may otherwise make ad hoc decisions about communication security. A well-designed policy balances the efficiency benefits of informal communication channels with the security requirements necessary to protect client confidences.

Policies should address not only lawyer-client communication but also communications among firm members. Internal communications about client matters shared through insecure channels create the same risks as direct communications with clients. A text message between associates discussing case strategy presents similar vulnerabilities to a text message with the client.

IV. The Metaverse, Spatial Audio, and Confidentiality Under Rule 1.6

Rule 1.6 - Confidentiality of Information

Rule 1.6 prohibits lawyers from revealing information relating to the representation of a client unless the client gives informed consent, the disclosure is impliedly authorized to carry out the representation, or the disclosure falls within specific exceptions enumerated in the Rule. This confidentiality obligation extends beyond the attorney-client privilege and protects all information relating to representation regardless of its source.

If you happen to be in New York, you should note that the above statement still applies, but there is one thing to note. New York's 1.6(a) includes a definition that you don't see on other state codes: ""Confidential information" consists of information gained during or relating to the representation of a client, whatever its source, that is (a) protected by the attorney-client privilege, (b) likely to be embarrassing or detrimental to the client if disclosed, or (c) information that the client has requested be kept confidential ...

Understanding the Metaverse

The metaverse represents an evolution of internet technology that creates persistent virtual environments where users interact through digital avatars. Unlike traditional video conferencing, which presents flat images of participants in separate windows, metaverse platforms create spatial environments where avatars occupy positions relative to each other in three-dimensional space. Users navigate these environments, moving their avatars closer to or farther from other participants, creating experiences that more closely approximate physical presence than conventional online interactions.

Several companies have developed metaverse platforms aimed at professional use, including virtual offices, conference rooms, and courtrooms. These platforms promise to improve remote collaboration by creating more natural interaction dynamics than traditional video conferencing. Legal organizations have begun experimenting with metaverse technologies for depositions, client meetings, mediations, and even trials.

The appeal of metaverse platforms for legal work stems from their ability to simulate important aspects of in-person interaction. Traditional video conferencing forces all participants into identical relationships with each other—everyone appears in similar-sized windows with no spatial differentiation. The metaverse allows for sidebar conversations, private discussions, and the formation of sub-groups within larger meetings, mimicking dynamics that occur naturally in physical settings.

Spatial Audio and Its Risks

Spatial audio represents one of the key technologies that makes metaverse platforms feel more natural than conventional video conferencing. Rather than having all participants' voices arrive at the same volume regardless of their position, spatial audio adjusts volume and directionality based on the relative positions of avatars in the virtual environment. When your avatar stands next to another participant's avatar, their voice sounds close and clear. As you move your avatar away, their voice becomes quieter and eventually inaudible.

This spatial audio functionality creates significant confidentiality risks for lawyers using metaverse platforms. In physical meetings, lawyers can step into a hallway or conference room to have confidential discussions with clients. The physical separation ensures that others cannot overhear the conversation. Metaverse platforms attempt to replicate this dynamic through spatial audio, allowing lawyers to move their avatars away from the main group to have private discussions.

However, spatial audio in virtual environments provides much weaker confidentiality protection than physical separation. The software determines who can hear what based on avatar positions, but this determination depends entirely on proper platform configuration and functioning. Unlike physical walls that reliably block sound transmission, virtual separation offers no such guarantees.

Technical Vulnerabilities in Metaverse Platforms

Metaverse platforms face numerous technical vulnerabilities that can compromise attorney-client confidentiality. First, the platforms themselves may contain security flaws that allow unauthorized access to communications. A participant might exploit software vulnerabilities to hear conversations that should be inaudible based on avatar positioning. Because metaverse technology remains relatively new, these platforms have not undergone the same extensive security testing and hardening that more established communication technologies have received.

Second, metaverse platforms typically record much more data about user interactions than traditional video conferencing systems. They track avatar positions, movements, gestures, and the formation of conversational groups. This metadata, combined with audio recordings of all conversations within the platform, creates detailed records of attorney-client interactions. If this data is inadequately protected, it could be accessed by unauthorized parties, violating Rule 1.6's confidentiality requirements.

Third, the complexity of metaverse platforms creates numerous opportunities for user error that could compromise confidentiality. A lawyer who believes they have moved their avatar into a private space may have actually remained within hearing range of other participants

due to misjudging distances in the virtual environment. Unlike physical meetings, where participants can rely on innate spatial awareness developed over a lifetime, metaverse interactions require users to learn artificial spatial rules that may not match their intuitions.

Platform Control and Third-Party Access

Metaverse platforms, like all cloud-based services, involve storing and transmitting client information through systems controlled by third parties. The platform provider has access to all communications that occur within its system. This creates Rule 1.6 concerns because lawyers must ensure that third-party service providers maintain confidentiality standards consistent with professional obligations.

Comment 18 to ABA Model Rule 1.6 addresses the use of technology in providing legal services and acknowledges that lawyers may use services that store or transmit client information. The comment requires lawyers to make reasonable efforts to prevent inadvertent or unauthorized disclosure of client information, including selecting and using appropriate technology for the representation.

Applying this standard to metaverse platforms requires lawyers to investigate the security practices of platform providers. Do they encrypt communications? How long do they retain recordings? Who has access to stored data? Can the provider disclose data to government agencies without client consent? These questions must be answered satisfactorily before lawyers can ethically use metaverse platforms for confidential client communications.

Informed Consent and Metaverse Use

Before using metaverse platforms for client communications, lawyers should obtain informed consent from clients. This consent requirement flows from Rule 1.6's provision allowing disclosure of client information when the client gives informed consent. While using a particular communication platform might not seem like "disclosure" in the traditional sense, it does involve transmitting client information to third parties who control the platform.

The term “informed consent” is defined in Rule 1.0(e). It requires that clients understand the material risks and available alternatives. Lawyers must explain that metaverse platforms involve transmitting communications through servers controlled by the platform provider, that these communications may be recorded and stored, and that the technology remains relatively new with potentially undiscovered security vulnerabilities. Clients should also understand alternative communication methods that may offer greater security.

Practical Limitations on Metaverse Use

Given the confidentiality risks associated with metaverse platforms, lawyers should carefully limit their use of these technologies for client communications. Preliminary client meetings, general case discussions, and collaborative work sessions may be appropriate for metaverse platforms when properly secured. However, discussions involving highly sensitive information, privileged communications, or strategic deliberations warrant more secure communication methods.

Lawyers who do use metaverse platforms should implement additional safeguards to protect confidentiality. Recording settings should be carefully controlled to prevent unauthorized recording of confidential discussions. Participants should be verified before meetings begin to prevent unauthorized individuals from accessing attorney-client communications. Regular security assessments of the platform and its provider should be conducted to identify emerging vulnerabilities.

The Future of Virtual Legal Practice

As metaverse technology matures and security practices improve, these platforms may become more suitable for confidential legal communications. However, lawyers must avoid assuming that technological advancement automatically solves confidentiality concerns. Each new generation of communication technology brings new vulnerabilities alongside its benefits. Maintaining compliance with Rule 1.6 requires lawyers to continually assess the risks associated with communication technologies and adjust their practices accordingly.

V. Generative AI

No program on technology these days is complete if we don't talk about generative AI. We will be discussing the issue of hallucinations and transparency in the spoken program. However, for the purposes of these materials I wanted to go a little overboard and provide you with a larger amount of substantive information. We will cover some, but not all of the concepts in this next section.

It seems like every day there is another opinion coming out that addresses a lawyer's use of generative AI in the practice. If you think that meant it would be tough to keep track of the ethical standards that apply, you'd be partially right. The volume of guidance being thrown at us seems to be increasing in pace. Below I'm going to synthesize the guidance from many of those opinions. One thing that you'll likely realize is that it isn't quite as daunting as it seems. That's because a lot of these opinions are all saying the same thing. So, let's look at it all...

Note: Bloomberg seems to be compiling a running list of all opinions being issued across the country. You could find that list here:

https://www.bloomberglaw.com/external/document/X2JK49QC000000/legal-profession-comparison-table-state-legal-ethics-guidance-on

A. Preliminary Guidance

Let's start of considering some basic ideas that were set forth by my home state, New Jersey. It's important to consider these ideas before we get into the ethics guidance because they establish some fundamental concerns when using generative AI and all technology in the practice. The numbered items below are from the report issued by the Task Force on Artificial Intelligence (AI) and the Law. The comments in bullets after each numbered item are my comments.[1]

1. Legal professionals must understand that education, knowledge and guidance are necessary to operate AI tools safely and ethically in a legal setting.

- We're going to see how important competence is in all of this. Of course, it's not exactly earth shattering, conceptually. But it's something that's been focused on since the first computer entered the legal practice.

2. When assessing AI tools and services, it is crucial to categorize them according to their intended users and recipients. Tools designed for the public, as opposed to legal professionals, should not be used for tasks considered "the practice of law."

[1] All references to information from New Jersey, or a reference to a NJ opinion are actually from this document: Task Force on Artificial Intelligence (AI) and the Law: Report, Requests, Recommendations, and Findings May 2024.

- Whether something is being done in the course of the practice is an issue that comes up in the ethics world a lot. I think, however, that it's a miustake to rely too much on this distinction when it comes to genAI and technology. There really is no more separation between our private lives and our professional lives.

3. When evaluating AI tools and services, it is essential to identify and document how data, especially client data, is transmitted, used and stored by the AI to ensure its confidentiality. This information should guide the assessment of whether a particular AI tool is suitable for its intended use.

- Confidentiality is the grandaddy of all ethical issues. That's true with genAI and it's true with all technologies. What's interesting, however, is that it's becoming increasingly more difficult for lawyers to actually comply with the rules. That's because there is so much we can't control in this new world of technology. For instance, the realities are that tech company contractors are reviewing user content. Also, security breaches are commonplace with all technologies. So, it is possible for lawyers to make reasonable efforts to protect client data? Do we need to reconsider what reasonable means? We are sitting in the midst of a paradigm shift and these issues are currently unresolved.

4. All law firms should adopt an organizational AI policy with a risk assessment framework (sample template provided in Appendix 2).

- It's about being proactive. This is huge, as we'll see below.

5. The Rules Governing the Courts of New Jersey, the New Jersey Rules of Evidence and the Rules of Professional Conduct are sufficiently flexible to address considerations relating to AI.

B. The ethics issues

1. Competence, including verifying the accuracy. I think 1.3 should be here as well.

I'm not even sure if it's necessary to repeat, again, how lawyers have a duty to understand technology under Rule 1.1.

> Rule 1.1. Competence.
> A lawyer shall provide competent representation to a client. Competent representation requires the legal knowledge, skill, thoroughness and preparation reasonably necessary for the representation.
>
> [8] Maintaining competence. — To maintain the requisite knowledge and skill, a lawyer should keep abreast of changes in the law and its practice, including the benefits and risks associated with relevant technology, engage in continuing study and education and comply with all continuing legal education requirements to which the lawyer is subject.

But remember that in the context it gets a little more broad. For instance, it is imperative to review your local court rules to ensure that you know the requirements when using genAI. That is part of being competence. States like Kentucky mentioned that explicitly in its Ethics Opinion KBA E-457 issued March 15, 2024.

Also remember that lawyers need to double check all sources when using genAI for research. We all know the famous Avianca case where the lawyer submitted fake cases that were made up by genAI. In his case, he didn't double check the case law. Doing so is an ethical mandate under both competence and Diligence, Rule 1.3

2. Confidentiality, including former clients and 1.9, and 1.8(b)

The use of generative AI by lawyers introduces several confidentiality concerns that must be carefully managed to ensure compliance with the rules of professional conduct.

> Confidentiality and Rule 1.6

One of the primary concerns is the potential breach of client confidentiality, as outlined in Rule 1.6, which mandates that lawyers must not reveal information relating to the representation of a client unless the client gives informed consent, the disclosure is impliedly authorized to carry out the representation, or the disclosure is permitted by specific exceptions in the rules. When using generative AI, there is a risk that sensitive information could be inadvertently exposed to unauthorized parties, especially if the AI system is provided by a third-party vendor or operates on a platform that may not have adequate security measures in place.

Confidentiality and Former Clients (Rule 1.9)

Additionally, Rule 1.9 addresses duties to former clients, including the prohibition against using information related to the representation of a former client to their disadvantage, except as the rules would permit or require with respect to a client, or when the information has become generally known. The use of generative AI could inadvertently lead to the misuse of such information if the AI system retains data from previous interactions and applies it in a manner that compromises the confidentiality of former clients.

Use of Information to Disadvantage a Client (Rule 1.8(b))

Rule 1.8(b) further emphasizes that a lawyer must not use information relating to the representation of a client to the disadvantage of the client unless the client gives informed consent. Generative AI systems that analyze large datasets or produce output based on patterns identified in confidential client information could potentially be used in ways that might inadvertently harm the client's interests.

Also, remember that anything you put into these publicly available gAI systems like Chat GPT will be used as training data and could be spit out in response to other people's queries in the future. Confidentiality, therefore, mandates that we don't put client information into our queries. Plus, remember not to give too many details, or that might also run afoul of Comment [4] to Rule 1.6:

Rule 1.6, Comment [4]

Paragraph (a) prohibits a lawyer from revealing information relating to the representation of a client. This prohibition also applies to disclosures by a lawyer that do not in themselves reveal protected information but could reasonably lead to the discovery of such information by a third person. A lawyer's use of a hypothetical to discuss issues relating to the representation is permissible so long as there is no reasonable likelihood that the listener will be able to ascertain the identity of the client or the situation involved.

3. Reasonable fees and costs, Rule 1.5

The use of generative AI by lawyers also raises concerns regarding the reasonableness of fees and costs, as outlined in Rule 1.5, which requires that a lawyer's fees be reasonable and that fees and expenses for legal services be communicated to the client. When incorporating AI into legal practice, lawyers must navigate the complexities of determining and justifying the costs associated with AI tools. For instance, if the implementation of

generative AI significantly reduces the time and effort required to perform certain legal tasks, lawyers must ensure that their fees reflect these efficiencies and do not result in overcharging the client.

Moreover, if the costs of using AI systems, such as subscription fees or maintenance costs, are passed on to the client, this must be transparently communicated and agreed upon in advance. According to Rule 1.5(b), any changes in the basis or rate of the fee or expenses must be promptly communicated to the client. Failure to appropriately manage and disclose these costs could lead to disputes over billing and potentially undermine the client's trust in the lawyer's services. Therefore, it is crucial for lawyers to exercise due diligence in evaluating the impact of AI on their billing practices and to ensure that their fees remain fair and reasonable in light of the efficiencies gained through AI technology.

4. Prevention of misconduct, 8.4 and accuracy and truthfulness. Candor to the tribunal. Rule 3.3

Generative AI, while offering substantial benefits in terms of efficiency and capability, brings with it concerns related to the prevention of misconduct in legal practice. Rule 8.4 deals with maintaining professional integrity and outlines what constitutes professional misconduct. Lawyers must be mindful of these provisions when integrating generative AI into their practice to avoid potential ethical pitfalls.

One significant concern with using generative AI is the risk of inadvertently disseminating inaccurate or misleading information. Generative AI systems, despite their advanced capabilities, can produce errors or biased outputs. Lawyers are accountable for ensuring the accuracy and truthfulness of the information they present, as required by Rule 3.3, which mandates candor toward the tribunal. Any misrepresentation, even unintentional, that arises from over-reliance on AI could be considered professional misconduct under Rule 8.4.

5. Transparency, 1.4 and 1.2

Transparency is a critical concern for lawyers using generative AI, particularly under Rules 1.2 and 1.4. Rule 1.2 requires a lawyer to abide by a client's decisions regarding the objectives of representation and to consult with the client about how they are pursued. This rule emphasizes the importance of clear communication and ensuring client autonomy. When incorporating generative AI into legal practice, lawyers must ensure that clients are informed about how AI tools may affect their legal matters, especially if the use of AI could influence significant decisions or strategies.

Rule 1.4 mandates that lawyers keep clients reasonably informed about the status of their matters and promptly comply with reasonable requests for information. This rule underscores the necessity of transparency in the lawyer-client relationship. If AI-generated content or tools are used in a way that impacts the client's case, it is imperative to disclose this usage, particularly if it affects the fee structure or the scope of work being performed.

But is there a blanket, mandatory obligation to disclose the use of AI in the practice? It seems not. At least, that's what Kentucky said. The Kentucky Bar Association addressed the issue of transparency in its Ethics Opinion KBA E-457 issued March 15, 2024 as follows:

> Does an attorney have an ethical duty to disclose to the client that AI is being used with respect to legal matters entrusted to the attorney by the client?
>
> No, there is no ethical duty to disclose the rote use of AI generated research for a client's matter unless the work is being outsourced to a third party; the client is being charged for the cost of AI; and/or the disclosure of AI generated research is required by Court Rules.

New Jersey chimed in as well.

> Those RPCs do not impose an affirmative obligation on lawyers to tell clients every time that they use AI. However, if a client asks if the lawyer is using AI, or if the client cannot make an informed decision about the representation without knowing that the lawyer is using AI, then the lawyer has an obligation to inform the client of the lawyer's use of AI.[2]

6. Supervision, both 5.1, and 5.3, including racial bias

Supervision, encapsulated in Rules 5.1 and 5.3, is a critical concern when lawyers incorporate generative AI into their practices.

Rule 5.1 mandates that partners or supervisory lawyers ensure that all lawyers in their firm conform to the professional standards set by the rules. This responsibility extends to the implementation and oversight of generative AI tools, ensuring they are used ethically and effectively. Supervisors must establish policies and procedures to mitigate risks associated with AI, such as errors or biases in AI-generated outputs.

[2] https://www.njcourts.gov/sites/default/files/notices/2024/01/n240125a.pdf last checked 1/2/2025.

Similarly, Rule 5.3 governs the responsibilities regarding non-lawyer assistants, which implicitly includes AI tools. Lawyers must make reasonable efforts to ensure that AI, as an assistant, conducts itself in a manner compatible with the professional obligations of the lawyer. This entails regularly monitoring the AI's outputs, verifying the accuracy and reliability of the information it generates, and taking necessary corrective actions if the AI system produces misleading or incorrect data. The lawyer's supervisory duties require a blend of continuous oversight and human judgment to prevent over-reliance on AI, thus maintaining the integrity and quality of legal representation.

Also keep in mind the related issue of substituted judgment. The California opinion explained that issue as follows:

> A lawyer's professional judgment cannot be delegated to generative AI and remains the lawyer's responsibility at all times. A lawyer should take steps to avoid over-reliance on generative AI to such a degree that it hinders critical attorney analysis fostered by traditional research and writing. For example, a lawyer may supplement any AI-generated research with human-performed research and supplement any AI-generated argument with critical, human-performed analysis and review of authorities.[3]

Also, remember the issue of bias in algorithms. Racial bias in generative AI is a significant concern that can impact the fairness and justice of legal practices. AI systems can inadvertently perpetuate or amplify existing biases present in the data they are trained on. It is essential for lawyers to recognize and address these biases to ensure equitable legal representation.

Supervisory lawyers must implement measures to detect and mitigate racial bias in AI outputs. This includes training AI systems on diverse and representative data sets, regularly auditing AI-generated content for biases, and applying human judgment to review and correct biased information. By taking these steps, lawyers can minimize the risk of racial bias and uphold the principles of fairness and justice in their practice.

[3] https://board.calbar.ca.gov/docs/agendaItem/Public/agendaitem1000031754.pdf, last checked 12/28/2024.

7. Returning the client file and 1.16. DC Ethics Opinion 388

Returning a client's file and adherence to Rule 1.16(d) are critical concerns when lawyers use generative AI in their practice. Rule 1.16(d) mandates that upon termination of representation, lawyers must surrender all papers and property to which the client is entitled. The DC authorities explained:

> When a representation is terminated, Rule 1.16(d) requires a lawyer to do several things, including "surrendering papers and property to which the client is entitled." As we discussed in D.C. Legal Ethics Opinion 333 (2005), this rule requires production of the "entire file," including "copies of internal notes and memoranda reflecting the views, thoughts and strategies of the lawyer." Although lawyers are not required to retain every piece of paper or electronic datum generated or received during a client representation, a lawyer should consider whether specific interactions with GAI in connection with a client matter should be retained as part of the client file.[4]

8. Participating in a fraud

California said all you need to know about this concept. Regarding California Rule 1.2.1, which is its unique version of Rule 1.2(d):

> A lawyer must comply with the law and cannot counsel a client to engage, or assist a client in conduct that the lawyer knows is a violation of any law, rule, or ruling of a tribunal when using generative AI tools. There are many relevant and applicable legal issues surrounding generative AI, including but not limited to compliance with AI-specific laws, privacy laws, cross-border data transfer laws, intellectual property laws, and cybersecurity concerns. A lawyer should analyze the relevant laws and regulations applicable to the attorney or the client.[5]

9. Know the rules in other jurisdictions (Rule 8.5)

In the rapidly evolving landscape of legal technology, the integration of generative AI tools into legal practices presents unique challenges, particularly regarding conflicts of laws and Rule 8.5.

-Different jurisdictions, different laws

[4] District of Columbia Ethics Opinion 388, Attorneys' Use of Generative Artificial Intelligence in Client Matters
[5] https://board.calbar.ca.gov/docs/agendaItem/Public/agendaitem1000031754.pdf, last checked 12/28/2024.

Generative AI tools often operate across multiple jurisdictions, each with its own set of laws and regulations. This creates potential conflicts of laws, where the rules and requirements of one jurisdiction may differ from or contradict those of another. Lawyers must be vigilant in understanding and reconciling these discrepancies to ensure compliance with all applicable legal standards. Some key considerations include:

- Privacy Laws: Different jurisdictions have varying privacy laws that govern the handling of personal data. Lawyers must ensure that the use of generative AI complies with these laws to protect client confidentiality and avoid legal repercussions.
- Cross-Border Data Transfers: The transfer of data across borders can trigger compliance issues with data protection regulations such as the GDPR in Europe. Lawyers must be aware of these regulations and implement necessary safeguards when using AI tools that process or transfer client data internationally.
- Intellectual Property: AI-generated content may raise questions of intellectual property ownership and rights. Lawyers must navigate these issues to avoid potential infringements and ensure proper attribution and usage of AI-created materials.
- Cybersecurity Concerns: The use of AI introduces additional cybersecurity risks. Lawyers must implement robust security measures to protect against data breaches and unauthorized access to sensitive information.

-Rule 8.5: Jurisdiction

Rule 8.5 addresses the disciplinary authority and choice of law in cases involving multi-jurisdictional practices. This rule is particularly pertinent when lawyers use generative AI, as it clarifies which jurisdiction's rules apply and under what circumstances. Key issues include:

- Disciplinary Authority: Rule 8.5(a) states that a lawyer is subject to the disciplinary authority of the jurisdiction in which they are admitted, regardless of where the conduct occurs. This means that lawyers using generative AI must comply with the ethical rules of their home jurisdiction, even if they are practicing across borders.
- Choice of Law: Rule 8.5(b) provides guidance on which jurisdiction's rules to follow when a lawyer's conduct involves significant contacts with multiple jurisdictions. This is critical for lawyers using AI tools that may operate or impact clients in different regions. Lawyers must determine the appropriate jurisdiction's rules to apply based on the specific circumstances of their practice and the location of their clients or AI operations.

10. Frivolous Claims and Rule 3.1

Rule 3.1 states that a lawyer shall not bring or defend a proceeding, or assert or controvert an issue therein, unless there is a basis in law and fact for doing so that is not frivolous. This rule is crucial in maintaining the integrity of the legal system by ensuring that legal actions are grounded in merit and not pursued for improper purposes, such as to harass or maliciously injure another party.

Generative AI tools, while powerful, can sometimes produce outputs that lack the necessary legal foundation. When lawyers rely on these tools, there is a risk that they may inadvertently use AI-generated content that does not meet the standards required by Rule 3.1. Here are some key concerns:

- Automated Document Drafting: Generative AI can streamline document drafting by automating repetitive tasks. However, if the AI generates content that includes unsupported claims or arguments, it can lead to the submission of frivolous claims. Lawyers must meticulously review and verify AI-generated documents to ensure they comply with all legal standards.
- Due Diligence and Verification: Lawyers have an obligation to conduct due diligence and verify the information and arguments they present. Using AI-generated content without proper scrutiny can result in the assertion of claims that lack factual or legal basis. This is particularly critical given that AI tools can sometimes produce plausible-sounding but inaccurate or irrelevant information.
- Ethical Obligations: Lawyers must uphold their ethical obligations and ensure that their use of AI tools does not compromise their duty to the court or their clients. Rule 3.1 emphasizes the importance of grounding legal actions in law and fact, and this duty extends to the use of any tools or technologies employed in legal practice.

11. Chatbots—consultation and 1.18, Formation of the A/C relationship, 2.1

Chatbots are used in legal marketing to help lawyers find valuable clients. The technology is basically a computer program that is powered by artificial intelligence and it simulates conversation with people. Potential clients who visit a firm's site can type questions and comments into a chatbox and, when doing so, they think they are speaking with a real person (or at least it's supposed to seem that way). Meanwhile, the bot collects contact info as well as other information about the potential client's case, analyzes it, and gives that info to the lawyer. The chatbot companies say that their AI allows them to sift out the tire kickers, identify the valuable prospects, and improve conversion rates from visitors to actual clients.

The chatbots are provided by tech vendors. A lawyer contracts with a vendor that offers the chatbot software, the vendor provides a bit of code that is inserted into the lawyer's

website, and the chatbot becomes a part of the lawyer's site. Someone coming to the website wouldn't know that another vendor is operating it— it simply looks like a chat box that is part of the lawyer's website.

Using a chatbot isn't necessarily a problem. What you need to be concerned about is the nature of the exchange between the bot and the potential client. Of course, it's a problem if a chat bot engages in conversation with a potential client and dispenses legal advice. But that's not likely to happen because that's not what the bots do. They are just supposed to be weeding out the garbage contacts from the good prospects. But in order to do that, the chatbot needs to ask the prospect some questions, and evaluate the data. That is where the problem could arise…

There is a conversation that goes on between the bot and the prospect. During that conversation the prospect will be providing information about their case. What we need to worry about is the potential that people who visit the lawyer's site and engage in a conversation with the chatbot end up being considered "prospective clients" under Rule 1.18. If they do attain that status, the lawyer could have conflict problems. To see what I mean, first understand how the rule works.

a. How Rule 1.18 Works

Rule 1.18 says that if a person "consults with a lawyer about the possibility of forming a client-lawyer relationship" they could be a prospective client. All they need to do is consult about the possibility of forming the lawyer client relationship. But what does that mean? Why should a lawyer care if someone is technically considered a "prospective client?"

First, you can't tell anyone about the information that the prospective client gave you. Rule 1.18(b) explains that "Even when no client-lawyer relationship ensues, a lawyer who has learned information from a prospective client shall not use or reveal that information…" Second, you might be conflicted out of representing people in the future. Even if you don't take the prospective client and never work on their matter, subsection (c) says that if you received information from the prospective client that could be significantly harmful to that person, and some time in the future a person approaches you to represent that new person against the prospective client in the same matter, you might not be permitted to so so. You would be conflicted out of the representation.

That could be devastating. Think about it— if you have a consultation with someone about a lucrative matter and you decide not to take their case…but later you are approached by someone who wants you to represent them in that very case you can't take that other client. You could be forced to forego a lot of money in fees.

b. The problem with chatbots

So back to the bots and Rule 1.18. What's important is the trigger for becoming a prospective client, and as you saw from the rule above, the trigger is a consultation. The

key question, of course, is, when does something rise to the level of a consultation? The answer is that it depends on the circumstances. But, in my opinion, the key circumstances to focus on are (1) what your website says and (2) the level of detail in the chatbot's communications.

If your website just lists your contact information you're going to be okay. If you simply put your information out there and someone sends you information about a case, that's not going create a prospective client relationship. Comment [2] confirms that: "...a consultation does not occur if a person provides information to a lawyer in response to advertising that merely describes the lawyer's education, experience, areas of practice, and contact information, or provides legal information of general interest." Basically, that comment is saying that if you simply tell someone that you exist and that you are qualified, it's not a "consultation." If someone replies in that situation, the person "communicates information unilaterally to a lawyer, without any reasonable expectation that the lawyer is willing to discuss the possibility of forming a client-lawyer relationship." That person, therefore, is not a prospective client.

However, you're going to have a problem if your website encourages people to offer information and your chatbot follows up by asking for information. The comment explains that "...a consultation is likely to have occurred if a lawyer...through the lawyer's advertising in any medium, specifically requests or invites the submission of information about a potential representation without clear and reasonably understandable warnings and cautionary statements that limit the lawyer's obligations, and a person provides information in response."

If your site specifically requests or invites a person to submit information about a potential representation, and your chat bot provides information in response, then you are risking the creation of a prospective client relationship. Obviously, the ethical danger is dependent upon the responsiveness of the chatbot because the rule says that you have to "provide information in response." Well, the more lengthy, intense, and detailed the chatbot's responses, the more likely there will be a problem.

Oh, and don't get hung up on the fact that your chatbot is not a "person" under the rules. Personally, if the bot provides information I think a tribunal will see the software as an extension of the lawyer. Plus, if the AI software is doing its job correctly, the potential client should believe that they are actually communicating with a real person. For those reasons, I wouldn't be surprised if a tribunal concluded that the AI in the chatbot is the functional equivalent of a "person" for the purposes of the rule.

Of course, there is a huge get-out-of-trouble card. All you have to do is include the disclaimers set forth in the rule. If your site has "clear and reasonably understandable warnings and cautionary statements that limit the lawyer's obligations" as stated in Comment [2], you're probably ok. This, however, is a situation where you can win the

ethical battle, but lose the overall war. What I mean by that is…what if this issue isn't raised in an ethics grievance? What if it is, instead, raised in a disqualification motion?

c. Win the ethical battle, but lost the disqualification war

Let's say you're in a medium sized firm that handles a variety of different types of mattes. Your firm represents Business X and you've been their counsel on various issues for years. Your firm has a website that utilizes a chatbot to evaluate the strength of clients. You have language on the website that properly disclaims Rule 1.18. Someone visits your site and explains that they have a workplace discrimination claim. They provide details of the case to the chatbot. The bot inquiries further and the prospect provides more information, in fact, the client wants to make sure that the lawyer with whom they are chatting has a complete understanding of the case (maybe they don't know it isn't a computer) so they provide a lot of details.

The chatbot sends the info to the attorney at the firm responsible for reviewing the contacts made by the chatbot and that lawyer thinks that the prospect has a great case. After reviewing the information, the attorney contacts the prospect and learns that the adverse party is Business X. However, the lawyer figures that the firm will probably be representing Business X in that matter because the firm does all of their work. As a result the firm doesn't take the potential client.

The prospect finds another lawyer, and they file suit against Business X. As the lawyer anticipated, the firm is representing Business X. The prospect's lawyer files a motion to disqualify you as counsel and you oppose it. You claim that there is no violation of the rule— the prospect never became a "prospective client" under Rule 1.18 because you had the proper disclaimer. And you're probably right. But there is a good chance that a judge will disqualify you anyway.

Remember, the judge isn't deciding discipline — the judge is deciding whether you should be disqualified. They don't necessarily care about the technicalities of the rules, they care about two things — the two things that are at the core of every conflict— loyalty and confidential information.

The critical question that the judge will ask was, during the interaction the firm had with the prospect, did you learn confidential information from the other party? And when the judge realizes that your chatbot gathered information that would ordinarily be considered confidential information and it was passed on to the lawyer in your firm for review, they're going to say you have a conflict and kick you out of the case. You're not going to be saved by the disclaimers because those disclaimers only helped you avoid discipline under Rule 1.18. In the disqualification context the court cares about loyalty and confidential information. And when it finds out that you were privy to a slew of details from the potential client's case, they will disqualify you.

d. How to make chatbots safer

All of this doesn't mean that chatbots can't be used, they just need to be used carefully. What can you do to make the chatbot safer? Here are 5 ideas:

- Use disclaimers
- Make sure the bot is just gathering information and not giving any information. And if it does give information, make sure it's super limited. Keep Comment [4] to Rule 1.18 in mind which states, "In order to avoid acquiring disqualifying information from a prospective client, a lawyer considering whether or not to undertake a new matter should limit the initial consultation to only such information as reasonably appears necessary for that purpose."
- Go over Rule 1.18 with the vendor supplying your chatbot. Make sure they understand it. also explain the disqualification issue. Remember, most tech vendors have no idea about the details rules like 1.18.
- Train the staff/lawyers in your office who are responsible for following up on the leads developed by the bot. Let them know about Rule 1.18 and the issue of disqualification.
- Create a process that limits the exposure the lawyers who review the information provided by the chatbots. It is possible to screen those attorneys per 1.18(d)(2). Here's what that section states, in part:

> (d) When the lawyer has received disqualifying information…representation is permissible if…(2) the lawyer who received the information took reasonable measures to avoid exposure to more disqualifying information than was reasonably necessary to determine whether to represent the prospective client; and (i) the disqualified lawyer is timely screened from any participation in the matter and is apportioned no part of the fee therefrom; and (ii) written notice is promptly given to the prospective client.

12. Generative AI and the Billable Hour

Every few years, the legal ethics world revisits the same question: is the billable hour dying? Generative AI has revived that conversation with renewed urgency, and lawyers across the internet are again proclaiming that hourly billing is finished. That proclamation is wrong. More than wrong, the move away from the billable hour is likely to create the very problems that attorneys are trying to avoid.

Understanding why requires starting with the rule that governs attorney fees.

Rule 1.5(a) prohibits a lawyer from making an agreement for, charging, or collecting an unreasonable fee. The rule then identifies eight factors relevant to the reasonableness

determination: the time and labor required; the novelty and difficulty of the questions involved; the skill requisite to perform the legal service properly; the likelihood that the representation will preclude other employment by the lawyer; the customary fee in the locality; whether the fee is fixed or contingent; the time limitations imposed by the client or the circumstances; the amount involved and the results obtained; the experience, reputation, and ability of the lawyers performing the services; and the nature and length of the professional relationship with the client. Time and labor required leads the list, and that placement is not accidental. It reflects the foundational premise of attorney compensation that has governed the profession for generations.

Rule 1.5(b) addresses the communication of fee arrangements. It requires that the basis or rate of the fee and expenses be communicated to the client, preferably in writing, before or within a reasonable time after commencing the representation. The requirement applies whenever the lawyer has not regularly represented the client. Rule 1.5(b) is a transparency provision, but its practical significance is substantial. An attorney who makes a fee arrangement without specifying its basis leaves the client without information necessary to evaluate what they are being asked to pay.

Rule 1.5(c) governs contingent fee agreements. A contingent fee must be confirmed in a writing signed by the client. That writing must state the method by which the fee is to be determined, including the percentage or percentages to accrue to the lawyer in the event of settlement, trial, or appeal, the litigation and other expenses to be deducted from the recovery, and whether those deductions will be made before or after the contingent fee is calculated. At the conclusion of a contingent fee matter, the lawyer must provide the client with a written statement showing the outcome of the matter and, if there is a recovery, showing the remittance to the client and the method of its determination. Contingent fees are prohibited entirely in domestic relations matters and in criminal defense representations.

Those three provisions frame every conversation about how attorneys charge for their services. They also expose a significant flaw in the emerging discourse about AI-driven billing innovation.

The billing arrangements being promoted as transformative responses to generative AI are, in nearly every case, arrangements that have existed for decades under different names. Outcome-based fees are contingencies. Success fees are contingencies. Value billing, in which the client pays based on the benefit received rather than time expended, is a form of contingency thinking applied outside its traditional context. Project-based billing is flat-fee billing. Subscription legal services are retainers. The vocabulary is new. The structures are not. Rule 1.5 already addressed all of them, and the ethical requirements that apply to

those arrangements did not change because the technology enabling them became more sophisticated.

The real problem with moving away from hourly billing in the AI era is that alternative structures expose attorneys to a reasonableness challenge they are not going to win.

Consider what generative AI actually does to the time a matter requires. Work that previously demanded five hours now demands two. An opinion letter that once took a senior associate most of a day is now drafted, researched, and refined in a fraction of that time. An attorney billing hourly faces a genuine reduction in revenue per matter, and that reduction is proportional to how effectively they use the technology. The attorney who uses generative AI well and bills hourly makes less money per file than the attorney who works slowly by hand.

The rational response, from a revenue standpoint, is to move to flat fees. Set a price for the deliverable rather than charging for the time it takes to produce it. A letter that used to generate $2,500 at $500 an hour over five hours could be priced at $2,000 as a flat fee. The client pays less than before; the attorney earns more per hour of actual work; the arrangement appears, on its face, to benefit both parties.

The arrangement collapses when a client challenges the fee.

A trier of fact evaluating the reasonableness of that $2,000 flat fee under Rule 1.5(a) is going to ask how long the work took. The attorney, assuming truthful testimony, answers: two hours, using generative AI. The trier of fact applies the first factor of the 1.5(a) analysis, which is the time and labor required, and concludes that two hours of work does not support a $2,000 fee. The fee is unreasonable. The attorney loses.

The argument that the flat fee should be evaluated against the value of the work product, or against what the same letter would have cost before AI existed, or against the client's benefit from receiving it, runs directly into the structure of Rule 1.5(a). Time and labor required is a listed factor. It does not disappear from the analysis because the attorney chose a different billing method. A trier of fact who knows that the work took two hours is not going to ignore that fact because the parties labeled their agreement a flat fee.

This is the trap that attorneys are walking toward with (mistaken) confidence. The billing innovation they believe will solve their revenue problem is the same mechanism that will expose them to fee challenges they cannot defend.

The hourly rate, for all its limitations, is the structure most naturally aligned with Rule 1.5(a). If generative AI reduces the time a matter requires, the attorney bills for less time. The bill reflects the actual labor. It is defensible in a fee dispute because its

reasonableness is evaluated on its own terms. An attorney who spent two hours on a letter and billed for two hours at an established rate has a straightforward answer for any tribunal that examines the charge.

One structure warrants genuine reconsideration in the AI environment, and it is not a new one. The classic general retainer, in which a client pays a set amount to secure the attorney's availability and ongoing counsel across a range of matters, becomes more attractive and defensible as AI changes the per-matter time investment. The retainer is not keyed to any single deliverable. It compensates the attorney for being available, for maintaining current knowledge, and for the relationship itself. An attorney who uses AI to serve a retained client more efficiently is not exposed to a fee challenge based on time per matter, because the retainer was never structured around time per matter. The retainer may be the one billing structure whose underlying logic actually makes sense in a world where technology makes individual tasks faster.

The death of the billable hour has been predicted before, and it has not arrived. Generative AI is a reason to examine billing practices carefully, and it is a reason to understand what Rule 1.5 actually requires. It is not a reason to abandon the structure that remains the most defensible under the rule that governs every fee we charge.

13. Generative AI, Privilege, and the Heppner Decision

For the first time in a long time, a decision is living up to the hype. The recent decision about privilege and AI is really as bad as they're saying. It's not clickbait designed to get you to read about some speculative ethics concern, this is the real deal. And what most of these pundits are not telling you—mostly because they don't know—is that it's far worse than it appears on its face. And it aggravates me, because I have been talking about it for some time and there is NO ANSWER.

So first, here's what happened:

A guy named Heppner was being investigated for securities fraud. He learned that was the case because he had discissions with government officials and realize that he was a target off their investigation. That investigation got kicked up several notches when he got a subpoena from the grand jury asking for all kinds of information.

Heppner figured that he was going to be charged with some crimes, so he did what most people do these days when they need professional advice. You do it too—what do you do when you get a blood test result in your medical portal. Do you call the doctor right away? No, you go to Doctor Google. THEN you call the doctor. Heppner did the 2026 legal

equivalent of that: He went to AI for advice. Heppner used the publicly available version of Claude, a generative AI LLM. For the novices out there, it's basically another company's version of Chat GPT.

Heppner's lawyer didn't tell him to do this, he did it on his own.

So, using Claude, he prepared reports that outlined defense strategy, and arguments he could make in response to possible charges he'd be facing from the government. He did a ton of research on his own and then consulted his lawyer.

Heppner was eventually arrested for securities fraud, and as part of the arrest, authorities seized a whole bunch of devices and papers, and among the things seized were those reports he created with Claude.

One could imagine that there might be some damning information in those documents. I mean, when Heppner asked Claude about the case he surely included some information in his queries that he'd rather not tell the government. You know, he was brainstorming, he was evaluating all of the possible problems he could face. That's why his lawyers tried to get the court to protect those AI generated reports.

Heppner's lawyers argued that the government should not be allowed to have those AI generated documents, and they based their argument on attorney-client privilege.

Privilege says that a client can't be compelled to reveal the confidential legal advice that's given between a lawyer and a client. Heppner's lawyers claimed that the reports he created with the AI were covered by attorney-client privilege. But to be successful, he had to show that the elements of privilege were satisfied, specifically, that it was (1) legal advice given by (1) a lawyer to (3) a client, and that (3) the advice was given confidentially.

Heppner's lawyers had an uphill battle. The first element, the legal advice, seems to be satisfied. The entire report was advice about how to deal with the charges Heppner might face. The second element, however, was a little tricky: privilege only covers information given from a *lawyer* to a client. Is the AI considered a lawyer?

Heppner's actual lawyers made some interesting arguments to try to satisfy that element. They conceded that the lawyers didn't actually run the searches, but they said that Heppner inputted things that he learned from counsel; he created the documents in anticipation of speaking with counsel, and he eventually shared those documents with counsel.

The government, obviously, opposed the motion and said that they should have access to the documents. The court rules for the government. And said that the documents are not protected by privilege. Here's why→

First, the advice was not given by lawyers. Discussions about legal matters with nonlawyers is not covered by privilege. In addition, the court noted that prior cases required that privilege requires a trusted human relationship with a licensed professional who owes a fiduciary duty and it subject to discipline. None of that was present here.

Seems sort of obvious. But that wasn't the part that got people all upset. It was the part of about confidentiality. The court held that the communications memorialized in those documents were not kept confidential because the generative AI platform that Hepner used said that it didn't keep information confidential.

The key is that the written privacy policy for the publicly available version of Claude specifically says that it does not keep your communications confidential. And it's not just Claude. This is the same with all of the publicly available LLMs. The user consents to Claude collecting data on user's inputs, and Claude's outputs and that it uses your data to train it's system. It also tells you that it will disclose that data to third parties, including "governmental regulatory authorities." The system puts users on notice that it will "disclose personal data to third parties in connection with claims. Disputes or litigation. The court noted that AI users do not have substantial privacy interests in their conversations when you voluntarily disclose information to a publicly available AI platform that retains that information in the ordinary course of business. And then they gave us the key phrase that is scaring everyone.

The court said that, for those reasons, Heppner did not have a reasonable expectation of confidentiality in his communications with Claude.

And the AI documents are not like confidential notes that someone prepares with the intent of sharing them with an attorney because they weren't shared with an attorney, they were shared with a third party (the AI platform). For what its worth, that language comes right out of the commentary to Rule 1.6.

VI. Transcription Software and the Limits of Rule 1.6

Legal transcription software has become a fixture of modern practice. Attorneys use it to convert recorded client conversations, deposition summaries, dictated notes, and draft arguments into workable text. The technology is fast, increasingly accurate, and deeply convenient. It is also, in ways that most attorneys have not considered, a meaningful threat to our obligations under Rule 1.6.

Rule 1.6(a) prohibits a lawyer from revealing information relating to the representation of a client. The scope of that prohibition is broad. It covers all information relating to the representation, regardless of whether the client has asked for confidentiality, regardless of whether the information would embarrass the client, and regardless of whether the information is already known to others. Rule 1.6(c) adds an affirmative dimension to that duty. It requires lawyers to make reasonable efforts to prevent the inadvertent or unauthorized disclosure of, or unauthorized access to, information relating to the representation. The word "efforts" is important. The Rules do require that we take the risk seriously and act accordingly.

Transcription software creates a confidentiality exposure at every stage of its operation, and understanding the exposure requires understanding how the technology actually works. When an attorney records a client conversation or dictates a document, the audio file does not stay on the attorney's device. It travels. The audio or text is uploaded to the vendor's cloud servers, where the transcription process occurs. The resulting transcript is then sent back to the attorney. That round trip is the heart of the problem.

During transit, the data may or may not be encrypted. Encryption in transit, typically through Transport Layer Security protocols, protects information as it moves between the attorney's device and the vendor's servers. Many transcription platforms use it. Some do not, or they use it inconsistently, or they use outdated protocols that offer weaker protection. An attorney who uses a transcription product without confirming that all data in transit is encrypted is transmitting client information across the internet without a meaningful guarantee that a third party cannot intercept it. ABA Formal Opinion 477R, issued in 2017 to address evolving threats to electronic communication, identified unencrypted channels as a significant confidentiality concern and called on attorneys to consider the sensitivity of the information before choosing a transmission method. A client's communication about pending litigation, a criminal matter, or a sensitive business transaction represents exactly the kind of information for which unencrypted transmission is not a reasonable choice.

The risks on the vendor's servers present a separate category of concern. Even if data is encrypted during transit, it generally must be decrypted for the transcription process to work. At that point, the information exists in a readable form on systems we do not control. There are several questions that arise, but there are two in particular we will address here—and attorneys almost never ask either of them before signing up for a service.

The first question is whether humans employed by the vendor can access the transcribed content. Many transcription services rely on a combination of automated processing and human review. Human reviewers may be used to improve accuracy, to audit the system, or to handle audio that the algorithm cannot process reliably. Some vendors are transparent about this; many are not. A vendor employee who reads a transcript of a client's confidential communication is a third party who has received information relating to the representation. That violates Rule 1.6. Comment 18 to Rule 1.6 addresses the use of technology in client communications and calls on attorneys to understand the tools they use. An attorney who does not know whether human beings employed by the transcription vendor can access client transcripts has not met that standard.

The second question is what the vendor does with the data after the transcript is returned. Data retention policies vary enormously across transcription platforms. Some vendors delete audio files and transcripts promptly after delivery. Others store them indefinitely. Some use the accumulated data to train and improve their machine learning models, which can mean that a client's words about a contract dispute or a criminal charge become part of a dataset used to build a commercial product. There are Rule 1.15 issues here. That rule requires that we safeguard client property. We are not doing that if we are allowing the client file (which is their property) to remain on that tech company’s servers.

Why? Well, think about data breaches. It seems like every tech company is a target for hackers. If that’s the case, then why should we let sensitive client data sit in a place that it can be exposed to potential hackers? And, yes, all sites can be hacked, potentially. That’s even a greater reason for why we should not leave our client data on servers where it doesn’t need to be. The less it’s distributed, the less of a chance it will be targeted.

Comment 18 to Rule 1.6 directs attorneys to consider, among other factors, the sensitivity of the information, the likelihood of disclosure if safeguards are not employed, the cost of additional safeguards, and the difficulty of implementing them. For legal transcription software, that analysis is not complicated. The information is often highly sensitive. The likelihood of disclosure increases with every unencrypted transmission and every human reviewer with server access. And the safeguard is straightforward: read the terms of service, ask the vendor direct questions about encryption, human access, and retention

practices, and choose a product that provides verifiable answers. The safeguard is not expensive. It is simply a task most attorneys skip.

Rule 1.6(c) does not require us to abandon transcription software. It requires us to understand what we are using before we use it. Convenience is not a defense to a confidentiality violation, and the fact that a software product is widely used in the profession does not establish that it satisfies our obligations under the Rules. The attorney who records a client call, runs it through a transcription service, and never considers where that audio went or who might have heard it has made a choice. Under Rule 1.6, it is not a defensible one.