

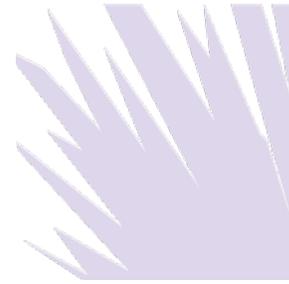
The Sedona Conference Journal

Volume 27

Forthcoming 2026

Artificial Intelligence in Federal Courts: A Random-Sample Survey of Judges

Anika Jaitley, Daniel W. Linna Jr., Hon. Xavier Rodriguez,
V.S. Subrahmanian & Siyu Tao



April 2026



Published in conjunction with the New York City Bar
Association Presidential Task Force on Artificial
Intelligence and Digital Technologies

Recommended Citation:

Anika Jaitley, Daniel W. Linna Jr., Hon. Xavier Rodriguez, V.S. Subrahmanian & Siyu
Tao, *Artificial Intelligence in Federal Courts: A Random-Sample Survey of Judges*, 27 SEDONA
CONF. J. ____ (forthcoming 2026),

https://thesedonaconference.org/publication/Artificial_Intelligence_in_Federal_Courts.pdf

©2026 The Sedona Conference. Licensed under CC BY-NC 4.0. To view a copy of this
license, visit <https://creativecommons.org/licenses/by-nc/4.0/>

For this and additional publications see: <https://thesedonaconference.org/publications>.

ARTIFICIAL INTELLIGENCE IN FEDERAL COURTS: A RANDOM-SAMPLE SURVEY OF JUDGES

TABLE OF CONTENTS

Acknowledgements	2
New York City Bar Association Presidential Task Force on Artificial Intelligence and Digital Technologies	3
Executive Summary	5
Methodology.....	8
Findings.....	13
Most Judges Have Used an AI Tool in Judicial Work	13
Top AI Use Cases by Judges are for Legal Research and Document Review	20
Top AI Use Case by Others in Chambers is Legal Research.....	25
Judges’ and Others’ AI Use is Highly Correlated	29
AI Training Has Not Been Offered to Most Judges	32
One in Three Judges Permit AI Use in Chambers; One in Four Have No Official Policy.....	34
38% of Judges Use AI Daily or Weekly Outside of Work	42
Judges’ Use of AI in their Work is Correlated with Personal AI Use.....	47
Judges are Equally Divided on Optimism vs Concern about AI in the Judiciary	49
Limitations	60
Appendix: Survey Questions	64

ACKNOWLEDGEMENTS

Authors:

Anika Jaitley, Northwestern University

Daniel W. Linna Jr., Professor of Instruction & Director of Law and Technology Initiatives, Pritzker School of Law & Computer Science Department, McCormick School of Engineering, Northwestern University

Hon. Xavier Rodriguez, United States District Judge, Western District of Texas

V.S. Subrahmanian, Walter P. Murphy Professor of Computer Science, Buffett Faculty Fellow at the Buffett Institute for Global Affairs, Northwestern University.

Siyu Tao, Pritzker School of Law, Northwestern University

Collaborators:

New York City Bar Association Presidential Task Force on Artificial Intelligence and Digital Technologies

Principal Investigators:

Professor Daniel W. Linna Jr. and Professor V.S. Subrahmanian, Northwestern University

Supported by:

This research is supported by the Law and Technology Initiative at Northwestern Pritzker School of Law & McCormick School of Engineering and the Northwestern Security & AI Lab at the McCormick School of Engineering and the Buffett Institute for Global Affairs at Northwestern University

NEW YORK CITY BAR ASSOCIATION PRESIDENTIAL TASK FORCE ON ARTIFICIAL INTELLIGENCE AND DIGITAL TECHNOLOGIES

The New York City Bar Association (City Bar) Presidential Task Force on Artificial Intelligence and Digital Technologies (Task Force) is composed of approximately 250 representatives of more than 50 committees, councils, and other task forces of the City Bar and adjunct members, including lawyers, academics, computer scientists, trade association representatives, consultants, technologists, roboticists, neurotechnologists, ethicists, and others.

Virtually every City Bar practice area is represented on the Task Force. The primary mission of the Task Force is to create a Center of Excellence and Thought Leadership in artificial intelligence and digital technologies, including innovative technologies that create new or enhanced products and services, digital assets, and more efficient, convenient, and effective ways of doing business.

For more information about the Task Force, please visit the Task Force dashboard at <https://www.nycbar.org/committees/task-force-on-digital-technologies/> and <https://www.nycbar.org/wp-content/uploads/2026/02/PTFAIDT-Leadership-260219.pdf>.

The main contributions from the Task Force were made by the Task Force Subcommittee on the Use of Artificial Intelligence in the Judiciary (AI Judiciary Subcommittee). The primary mission of the AI Judiciary Subcommittee includes addressing state and local courts; federal courts; litigants; court administration; specialized technology courts such as artificial intelligence courts and robo judges; and alternative dispute resolution.

AI Judiciary Subcommittee members include:

- Gus Coritsidis, Legal Advisor, US Court of International Trade
- Abena Darkeh, Judge, Criminal Court of the City of New York
- Kevin M.K. Fodouop, Associate at Wachtell, Lipton, Rosen & Katz
- Paul Grimm, Professor at Duke University School of Law and former U.S. District Court Judge, District Court for the District of Maryland

- Maura R. Grossman, Research Professor, School of Computer Science at University of Waterloo and Adjunct Professor at Osgoode Hall Law School
- Joseph Lamendola, New York State Supreme Court, Presiding Judge, Commercial Division, Onondaga County
- Co-Chair Harut Minasian, Associate at Paul, Weiss and former Judicial Law Clerk, U.S. District Court, Western District of Texas and New York State Supreme Court, Westchester County, Commercial Division
- Xavier Rodriguez, U.S. District Judge, U.S. District Court for the Western District of Texas
- Susan E. Salazar, Assistant Professor and Director, Mediation Clinic & Supervising Attorney, Main Street Legal Services, Inc. at CUNY School of Law
- John F. Werner, former Chief Clerk and Executive Officer, Supreme Court, Civil Term, New York County
- Co-Chair David Zaslowsky, Partner at Baker McKenzie

EXECUTIVE SUMMARY

The purpose of this study is to understand how, and to what extent, federal judges and other personnel who work in their chambers use artificial intelligence (AI) tools in their judicial work. We selected a stratified random sample of 502 federal bankruptcy, magistrate, district court, and court of appeals judges from a population of 1,738 current federal judges. Of the 502 judges that we surveyed via email, 112 responded (22.3% response rate).

Although a majority of responding judges at least occasionally use AI tools in their judicial work, relatively few report using AI on a daily or weekly basis. Approximately 38% of judges reported that they did not use AI at all in their work. This pattern suggests that AI is present in federal judicial chambers but not yet a routine, embedded part of most judges' decision-making processes.

Respondents report more frequent use of legal-specific AI tools integrated into established research platforms (such as Westlaw's AI-Assisted Research and similar tools) than of stand-alone, general-purpose AI tools such as ChatGPT, Copilot, or Gemini. This pattern indicates that vendor familiarity and perceived reliability may strongly shape which AI tools judges are willing to deploy in chambers.

Judges' attitudes toward AI are almost evenly split between optimism and concern. Many respondents simultaneously recognize AI's potential efficiency gains and express unease about hallucinations, "zombie cases," and skill atrophy. When AI training is offered by court administration, most judges attend, but a sizeable majority have not been offered such training or are unsure whether training has been available, suggesting unmet demand for high-quality, judiciary-specific education on AI.

KEY FINDINGS

The survey reveals the following main insights:

1. **AI adoption is broad but infrequent:** More than 60% of responding judges reported using at least one AI tool in their judicial work. However, only 22.4% reported using these tools on a weekly or daily basis. Conversely, 38.4% of judges have never used any of the AI tools listed on the survey in their work.

2. **Preference for legal-specific AI tools:** Judges are more likely to use "AI for Law" tools (AI tools specifically designed for legal use cases) than general-purpose AI platforms and use them more frequently.

3. **Legal research dominates use cases:** The most common AI use case for judges by far is conducting legal research, reported by 30.0%. This is followed by reviewing documents at 15.5%. Similarly, the top AI use case for others working in chambers is also legal research, at 39.8%, followed by reviewing documents at 16.7%.

4. **AI training has not been offered to most judges:** 45.5% of judges said that AI training had not been provided by court administration, and 15.7% said that they were not sure. Three out of four judges offered AI training attended the training.

5. **One in three judges permit AI use in chambers:** One in three judges permit or permit and encourage the use of AI by those working in their chambers. Only 20.4% of judges formally prohibit AI use, and 17.6% of judges discourage but do not formally prohibit AI use.

6. **One in four judges has no official AI use policy:** About one in four judges (24.1%) reported that they have no official AI policy. If the 17.6% of judges who discourage but do not formally prohibit AI use are also counted as not having an official policy, that number climbs to 41.7% of judges not having an official AI use policy.

7. **Personal and professional AI use are correlated:** A judge's use of AI in their personal life is correlated with their use of AI in their judicial work. Overall, 38% of judges use AI daily or weekly outside of work. A majority of judges reported that they rarely (26.9%) or never (25.9%) use AI outside of work.

8. **Judicial outlook on AI is evenly split:** Judges are nearly evenly divided between those who are optimistic about AI's potential for the judiciary and those who are concerned.

METHODOLOGY

To understand how federal judges are using AI, we conducted a stratified random sample survey of bankruptcy, magistrate, district court, and court of appeals judges. Our goal was to create a representative snapshot of federal judges. Our study population is defined below.

STUDY POPULATION

We began by defining our study population as every active federal judge serving as of August 2025. Because there is no definitive list of all federal judges, we compiled our study population from several sources:

Bankruptcy Judges: We compiled a list using Ballotpedia¹ and verified the names against individual court websites.

Magistrate Judges: We built this list using the *Almanac of the Federal Judiciary*² and removed anyone who was no longer in active service.

District and Court of Appeals Judges: We used the Federal Judicial Center's *Biographical Directory*.³ We removed judges who were on senior status, retired, resigned, or deceased. To ensure we covered the entire U.S., we manually added judges from the territorial courts in Guam, the Northern Mariana Islands, and the U.S. Virgin Islands, using publicly available information from their websites.⁴

Note that we primarily focused on Article III courts but did not include Supreme Court justices or Court of International Trade judges. Given the small number of each, surveying these courts would have resulted in small samples and non-representative results for those types of judges. We also included Article IV territorial court judges in the District Court judges

1. *Federal bankruptcy judge*, BALLOTPEDIA (last visited Feb. 22, 2026), https://ballotpedia.org/Federal_bankruptcy_judge.

2. *Almanac of the Federal Judiciary* (Wolters Kluwer 2025).

3. *Biographical Directory of Article III Federal Judges*, <https://www.fjc.gov/history/judges/biographical-directory-article-iii-federal-judges-export> (last visited Feb. 22, 2026).

4. District Court of Guam, <https://www.gud.uscourts.gov/> (last visited Feb. 22, 2026); District Court for the Northern Mariana Islands, <https://www.nmid.uscourts.gov/> (last visited Feb. 22, 2026); District Court of the Virgin Islands, <https://www.vid.uscourts.gov/> (last visited Feb. 22, 2026).

category.

We included bankruptcy judges, as bankruptcy courts are established as adjuncts to district courts,⁵ but did not include judges in other Article I courts. Our findings are limited to judges in the types of courts we surveyed.

STRATIFIED SAMPLE SELECTION

We gathered the individual U.S. court's email address for each judge through a manual process. Given the time and resource constraints, we aimed for a final sample size of approximately 500 judges. To achieve a representative stratified sample, we calculated that we needed to survey approximately 29% from the pool of judges. We applied this sampling rate evenly across all four categories of judges to ensure each category was proportionally represented. To pick the specific judges in each category, we assigned each judge a random number between 0 and 1 (with 15 decimal places) and selected judges from each category starting with the greatest number until we reached the target numbers below. This stratified random sample process produced a list of 502 judges.

<i>Judge Type</i>	<i>Total Available</i>	<i># Selected</i>
<i>Bankruptcy</i>	319	92
<i>Magistrate</i>	612	177
<i>District Court</i>	630	182
<i>Court of Appeals</i>	177	51
<i>Total</i>	1,738	502

Table 1: Judge Types and Number Selected

5. Bankruptcy Act of 1978, Pub. L. No. 95-598, 92 Stat. 2549 (codified in scattered sections of 11, 28 U.S.C.). Bankruptcy courts are units of the district courts under § 201(a), 28 U.S.C. § 151(a).

SURVEY PROCESS

Once the list of 502 judges was finalized in October 2025,⁶ we designed a Qualtrics survey intended to take approximately five minutes to complete. The survey featured a mix of multiple-choice and free-response questions. It asked judges about their familiarity with AI, current use of AI tools, and perspectives on AI's potential impact on the judiciary. For example, respondents were asked whether they or others in their chambers used specific AI tools daily, weekly, monthly, rarely, or never, and for what purposes. Additional questions inquired whether judges had received formal training on AI or if their chambers had established formal policies regarding AI use. This survey was approved by Northwestern University's Institutional Review Board. Only the Northwestern researchers had access to the Qualtrics survey and the unprocessed data comprising the judges' responses. To protect confidentiality, the Northwestern researchers shared with the other authors and collaborators only (i) visualizations and tables aggregating judges' responses and (ii) judges' individual responses without personally identifiable information. The survey questions are available in the **Appendix**.

Research indicates that sending a pre-notification email prior to survey distribution can increase response rates.⁷ Accordingly, a pre-notice email was sent to judges on Monday, December 1, 2025. The survey itself was distributed about 24 hours later. Follow-up reminder emails were sent one week and two weeks after the initial distribution, with a final reminder sent one day before the survey closed on December 19, 2025.

6. During this step, we found that one judge on our list was no longer serving. We replaced that judge with the next randomly selected judge from the same category. We also successfully identified one missing email address by using the standard email format for that specific court.

7. Roberta Sammut et al., *Strategies to Improve Response Rates to Web Surveys: A Literature Review*, 123 INT'L J. NURSING STUD. art. no. 104058 (2021), <https://doi.org/10.1016/j.ijnurstu.2021.104058>.

SURVEY RESPONSES

We collected a total of 112 responses. These responses consisted of 100 complete responses, eight substantively complete responses,⁸ and four email responses from judges who stated that they do not use AI in their work but did not answer any other questions. Partial responses are only used in the analysis of the relevant questions.

Response Rates by Judge Type. The overall response rate for the survey was 22.3% (112 out of 502). Participation varied significantly, however, depending on the type of judge. Bankruptcy judges had the highest response rate at 33.7%. Court of Appeals judges had the lowest response rate at 11.8%. While our list covers the federal judiciary broadly, we received only six responses from Court of Appeals judges. Given the relatively low number of responses, results related to Court of Appeals judges should be viewed as anecdotal and not representative of the full population of Court of Appeals judges. Additionally, while the response rates and numbers of responses were greater for the other categories of federal judges, results may not be representative of the full population of judges in those categories, as discussed further below, including in the “Limitations” section.

<i>Judge Type</i>	<i>Total</i>	<i># Selected</i>	<i># of Responses</i>	<i>Response Rate</i>
<i>Bankruptcy</i>	319	92	31	33.7%
<i>Magistrate</i>	612	177	32	18.1%
<i>District Court</i>	630	182	43	23.6%
<i>Court of Appeals</i>	177	51	6	11.8%
<i>Total</i>	1,738	502	112	22.3%

Table 2: Stratified Random Sample and Response Rate

8. These survey responses were complete except for a technical issue preventing a final email field from being validated. We corrected the technical issue and manually retrieved these answers and included them.

NOTE ON LIMITATIONS

It is important to consider inherent limitations when interpreting any survey results. Although our choice of a stratified random sample addresses problems related to non-random sampling, other factors such as self-selection bias, non-response bias, and social desirability bias can still result in unrepresentative survey data. For example, judges may have been more, or less, likely to respond based on their personal interest in AI, their comfort in revealing their usage, or their strong opinions regarding AI's likely impact on the judiciary.

Additionally, we note the sample size of 112 responding judges compared to the total population of 1,738 federal judges, as well as the varying response rates across different categories of judges. While the overall sample is large enough to suggest that it is broadly representative subject to a margin of error, the findings must be viewed with a recognition of these constraints. Despite these limitations, our methodology and response rate should provide a good foundation for understanding how federal judges are using AI.

For a comprehensive discussion of these limitations, please refer to the "Limitations" section at the end of this article.

FINDINGS

MOST JUDGES HAVE USED AN AI TOOL IN JUDICIAL WORK

Question 1.1: How frequently do you use the following AI tools in your work as a judge?
ChatGPT (OpenAI)
Claude (Anthropic)
Copilot (Microsoft)
Gemini (Google)
Grok (X.ai)
Perplexity
CoCounsel (Thomson Reuters)
Westlaw AI-Assisted or Deep Research (Thomson Reuters)
Protégé or Lexis+ AI (LexisNexis)
Vincent AI (vLex)
Harvey
Legora

As shown in **Figure 1**, more than 60% of responding judges reported using one or more AI tools in their work as a judge, although only 22.4% reported “weekly” or “daily” usage. Another 19.6% of judges said that they use at least one of the AI tools “monthly”, and a further 19.6% of judges said that they “rarely” use any AI tools. The remaining 38.4% of judges responded that they have “never” used any of the AI tools listed in their work.

Click on image to download and enlarge.

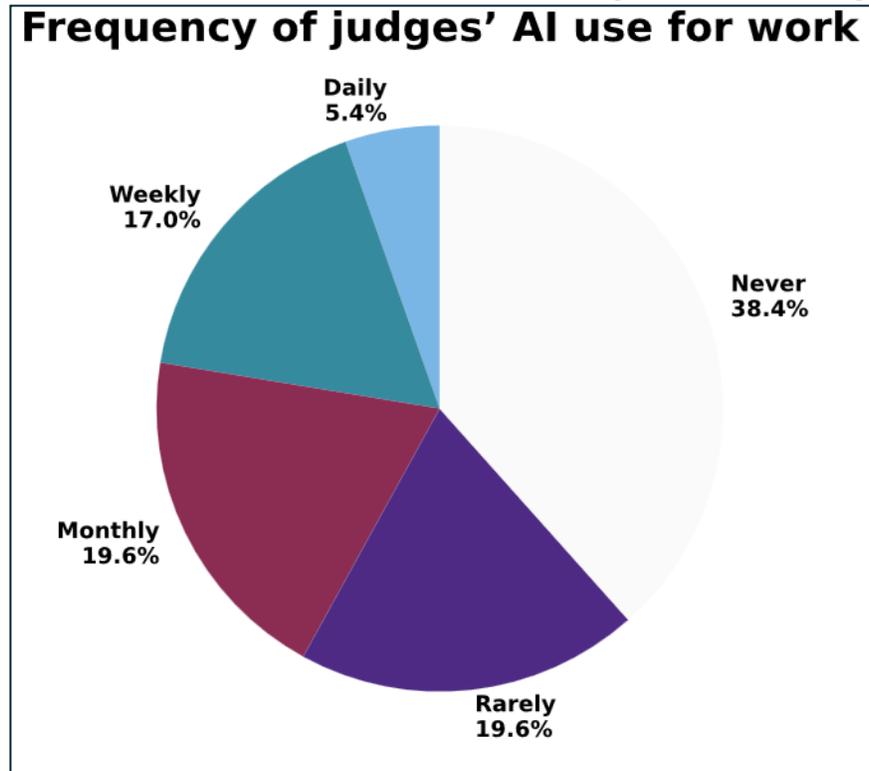


Figure 1: Frequency of judges' AI use for work

For each of the 112 judges who responded to Question 1.1, their frequency of use categorization was determined by the AI tool that they reported using most frequently.

If we categorize the AI tools in Question 1.1 into "AI for Law" tools (i.e., CoCounsel, Westlaw, Protégé/LexisNexis, Vincent AI, Harvey, Legora) and "General AI" tools (i.e., ChatGPT, Claude, Copilot, Gemini, Grok, Perplexity), we observe that the "AI for Law" category of tools is used by more judges and more frequently. As shown in **Figure 2**, judges are more likely to use "AI for Law" tools daily, weekly, and monthly than "General AI" tools. Nearly 20% of judges reported that they rarely use "General AI" tools while approximately 15% reported that they rarely use "AI for Law" tools.

Click on image to download and enlarge.

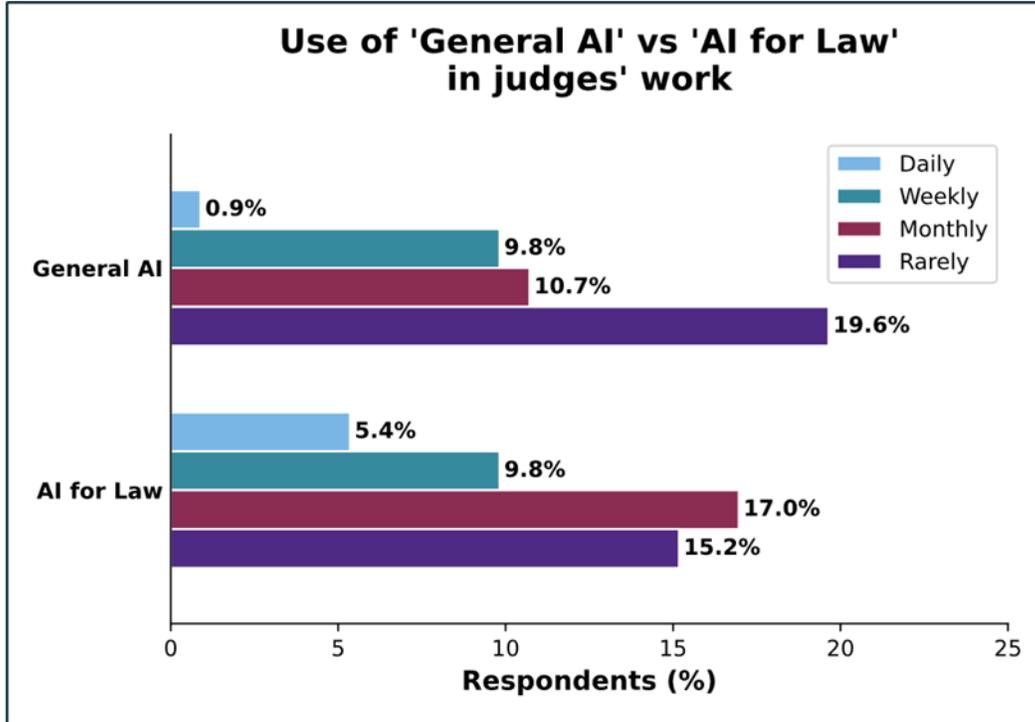


Figure 2: Use of “General AI” vs “AI for Law” in judges' work

“General AI” includes ChatGPT, Claude, Copilot, Gemini, Grok, Perplexity.

“AI for Law” includes CoCounsel, Westlaw, Protégé/LexisNexis, Vincent AI, Harvey, Legora.

For each of the 112 judges who responded to Question 1.1, their frequency of use categorization was determined by the AI tool that they reported using most frequently in each category (i.e., “General AI”, “AI for Law”).

Judges’ reported use of specific AI tools is shown in **Figure 3**. The AI tools judges reported using the most were “Westlaw AI-Assisted or Deep Research (Thomson Reuters)” at 38.4% and “ChatGPT (OpenAI)” at 28.6%. **Table 3** shows, however, that daily and weekly use of AI tools by judges is significantly lower.

Click on image to download and enlarge.

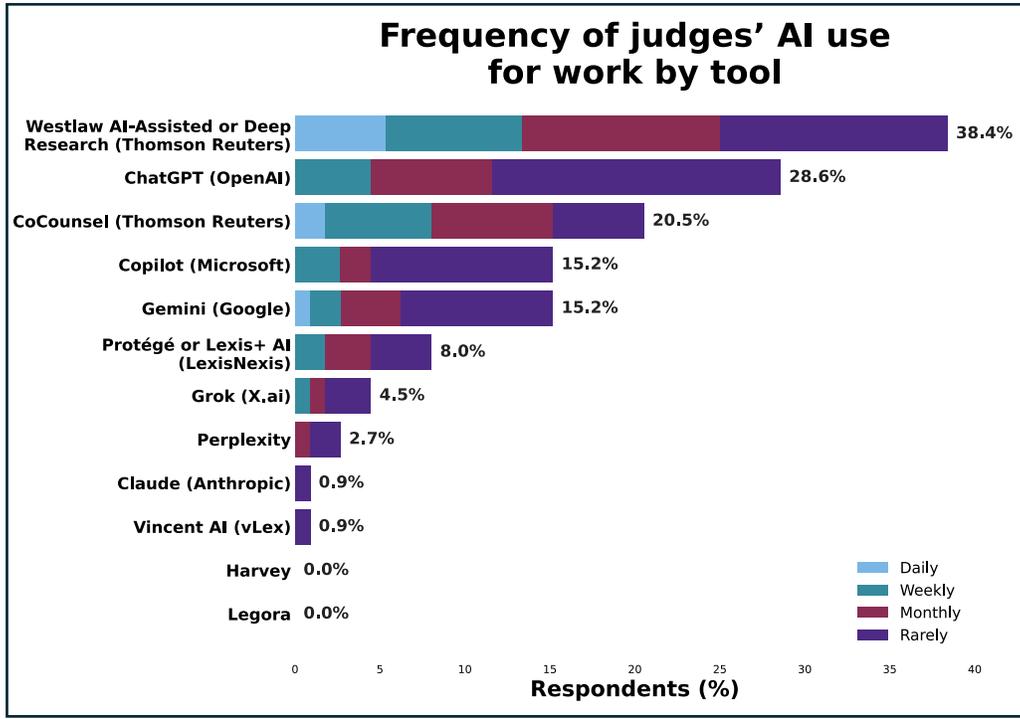


Figure 3: Frequency of judges’ AI use for work by tool

For each AI tool, the figure above shows the percentage of total respondents who reported using that tool either daily, weekly, monthly, or rarely. A total of 112 judges responded to this question. (“Never” responses are not shown on bars.)

	<i>Daily</i>	<i>Weekly</i>	<i>Monthly</i>	<i>Rarely</i>	<i>Never</i>
ChatGPT (OpenAI)	0%	4.5%	7.1%	17.0%	71.4%
Claude (Anthropic)	0%	0%	0%	0.9%	99.1%
Copilot (Microsoft)	0%	2.7%	1.8%	10.7%	84.8%
Gemini (Google)	0.9%	1.8%	3.6%	8.9%	84.8%
Grok (X.ai)	0%	0.9%	0.9%	2.7%	95.5%
Perplexity	0%	0%	0.9%	1.8%	97.3%
CoCounsel (Thomson Reuters)	1.8%	6.3%	7.1%	5.4%	79.4%
Westlaw AI-Assisted or Deep Research (Thomson Reuters)	5.4%	8.0%	11.6%	13.4%	61.6%
Protégé or Lexis+ AI (LexisNexis)	0%	1.8%	2.7%	3.6%	91.9%
Vincent AI (vLex)	0%	0%	0%	0.9%	99.1%
Harvey	0%	0%	0%	0%	100%
Legora	0%	0%	0%	0%	100%

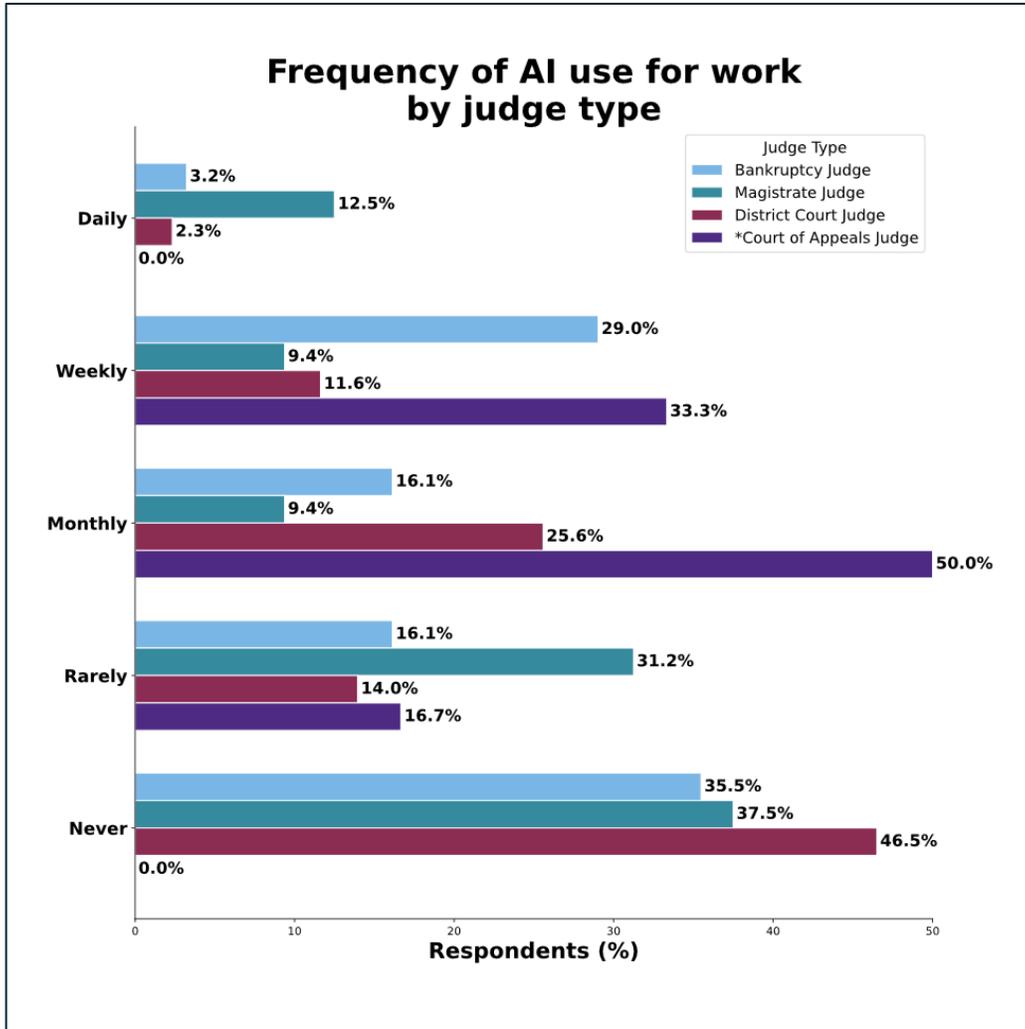
Table 3: Judges' Frequency of Use of AI Tools in their Work (Question 1.1)

Figure 4 compares the reported usage of AI tools in work by different categories of judges, although the results here must be viewed with particular caution due to the low response rate, especially among Court of Appeals judges in our study sample.

With these caveats in mind, 12.5% of Magistrate judges reported “daily” use of AI tools, a much higher percentage than Bankruptcy (3.2%) and District Court (2.3%) judges. An additional 29.0% of Bankruptcy judges reported using AI tools “weekly”, compared to 9.4% of Magistrate judges and 11.6% of District Court judges. Combining daily and weekly usage,

32.2% of Bankruptcy judges reported using AI tools “daily” or “weekly” compared to 21.9% of Magistrate judges and 13.9% of District Court judges. At the other end of the spectrum, 46.5% of District Court judges reported “never” using AI tools in their work, compared to 35.5% of Bankruptcy judges and 37.5% of Magistrate judges.

Click on image to download and enlarge.



*Low response rate. See discussion of “Limitations.”

Figure 4: Frequency of AI use for work by judge type

For each of the 112 judges who responded to Question 1.1, their frequency of use categorization was determined by the AI tool that they reported using most frequently. Each percentage reflects the percentage of that judge type that reported a particular frequency of AI use. The percentages of responses per judge type add up to a total of 100%.

Question 1.2: If you use other AI tools in your work as a judge, please identify the tools and specify for each if you use it Daily, Weekly, Monthly, Rarely, or Never.

We asked judges to identify other AI tools that they use in their work and to specify the frequency of use. We received four responses identifying other uses of AI tools. One judge identified Speechify (an AI-based text-to-speech and voice-to-text tool) as an AI tool used weekly. This highlights how other AI tools that may not directly involve legal research or reasoning may be used in judicial work. Another judge pointed out that their response would depend on the definition of AI (e.g., whether text prediction would also be considered AI use). In the “Limitations” section, we discuss why we did not attempt to define AI and instead started with a list of AI tools and asked open-ended questions about the use of AI tools.

Judges’ responses⁹ included:

- Speechify—weekly.
- AI tools are predominately used by law clerks.
- I don’t use any AI tools, but my [spouse] uses it and shares research findings with me.
- It depends on how you define AI tools. I assume you’re referring to Generative AI. Even assuming it’s Gen AI you’re concerned with, would text prediction be included?

9. Judges’ responses throughout this paper are exact quotes of their responses to the survey, except that the responses have been cleaned up for consistent formatting, grammar, and punctuation. Specific text in responses that created a discernable risk of linking that response to a specific judge has been removed and replaced with generalized text in brackets.

TOP AI USE CASES BY JUDGES ARE FOR LEGAL RESEARCH AND DOCUMENT REVIEW

Question 1.3: Do you use AI for the following use cases in your work as a judge? Select all that apply.

- Review documents (e.g., searching, asking questions about documents)
- Conduct legal research
- Draft documents filed in cases (e.g., orders, opinions, judgments, jury instructions)
- Draft documents not filed in cases (e.g., letters, emails, articles)
- Edit documents filed in cases (e.g., orders, opinions, judgments, jury instructions)
- Edit documents not filed in cases (e.g., letters, emails, articles)
- Inform decisions
- Make decisions
- Prepare case timelines or chronologies
- Transcribe text or audio
- Summarize text or audio
- Analyze text or audio
- Prepare questions for a proceeding (e.g., hearing, trial, status conference)
- Manage case files and administrative tasks
- Other (please describe below)
- I don't use AI in my work

The most common AI use case, reported by 30.0% of judges, was to “conduct legal research.” Next, 15.5% of judges reported using AI to “review documents (e.g., searching, asking questions about documents).” The other individual use cases were all below 10%, although 13.6% of judges reported using AI for “Other” use cases. The majority of judges (51%) reported, “I don’t use AI in my work.” This number is greater than the 38.4% of judges who responded “never” for each of the tools identified in Question 1.1, as fifteen of the judges who reported “I don’t use AI in my work” in response to Question 1.3 responded to Question 1.1 that they use some AI tools “rarely” (7 judges), “monthly” (5 judges), or “weekly” (3 judges). The discrepancy may be an unintentional omission or reflect the respondents’ overall infrequent usage of AI tools, or both.

Judges reported higher usage rates of AI to “draft” (7.3%) and “edit”

(4.5%) “documents **not** filed in cases” (emphasis added) compared to using AI to “draft” (1.8%) and “edit” (2.7%) “documents filed in cases.” Only 1.8% of judges reported using AI to “make decisions” while 4.5% reported using AI to “inform decisions.”

Click on image to download and enlarge.

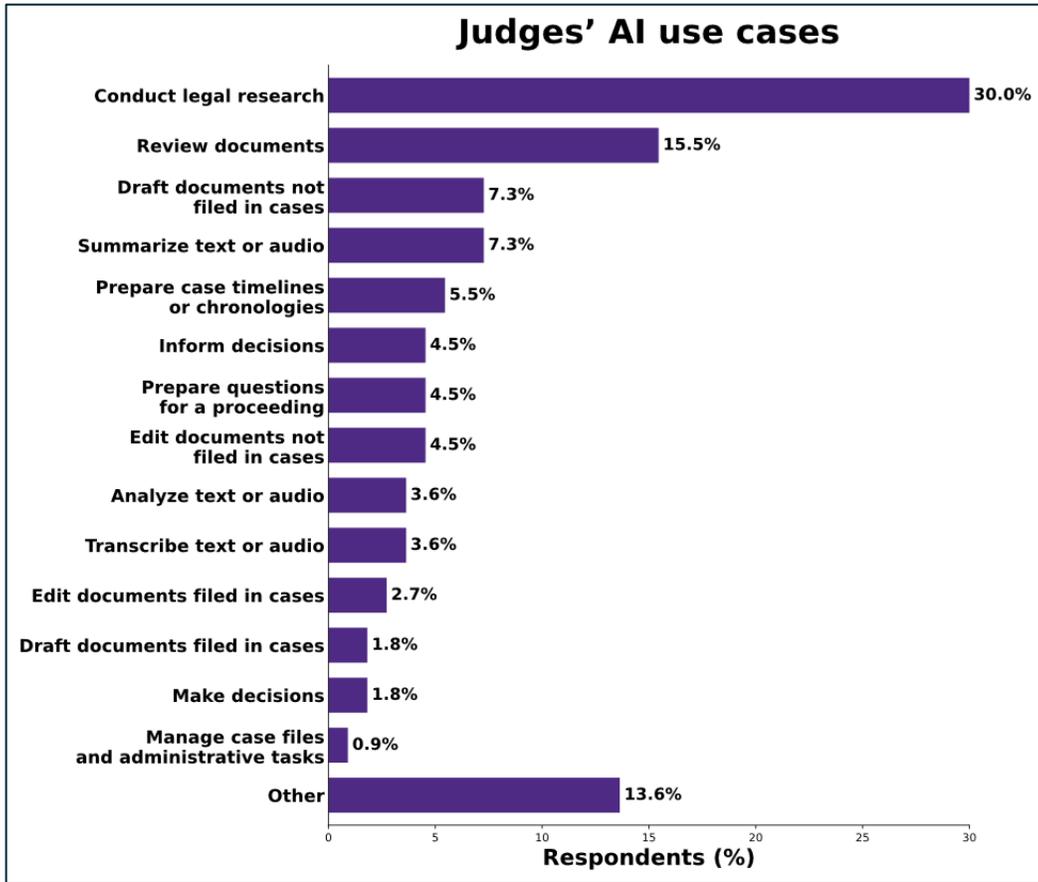


Figure 5: Judges’ AI use cases

110 judges responded to Question 1.3.

Note: 51% of respondents reported “I don’t use AI in my work.” This bar is not included.

We took the judges' responses about these specific use cases and recategorized them into broader use cases. **Table 4** lists the specific use cases and the broad categories we grouped them into. Some use cases could have fit into more than one of the broad categories.

Specific Use Cases	Use Case Categories
Review documents	Review, search, and analyze documents
Prepare case timelines or chronologies	
Summarize text or audio	
Analyze text or audio	
Conduct Legal Research	Conduct Legal Research
Draft documents filed in cases	Draft and edit documents
Draft documents not filed in cases	
Edit documents filed in cases	
Edit documents not filed in cases	
Inform decisions	Decision making
Make decisions	
Prepare questions for a proceeding	
Manage case files and administrative tasks	Manage case files and administrative tasks
Transcribe text or audio	
Other	Other

Table 4: Categorization of Specific Judicial AI Use Cases

Based on our aggregation of judges' responses into broader categories, 31.8% of judges reported uses of AI that fall under "review, search, and analyze documents," and 30.0% reported using AI to "conduct legal research." Next, 16.4% of judges reported uses of AI that fall under "draft and edit documents," 10.9% reported uses of AI that fall under "decision making," 4.5% reported AI uses that fall under "manage case files and administrative tasks," and 13.6% described an "other" use case. It bears repeating that 51% of judges said that they do not use AI in their work.

Click on image to download and enlarge.

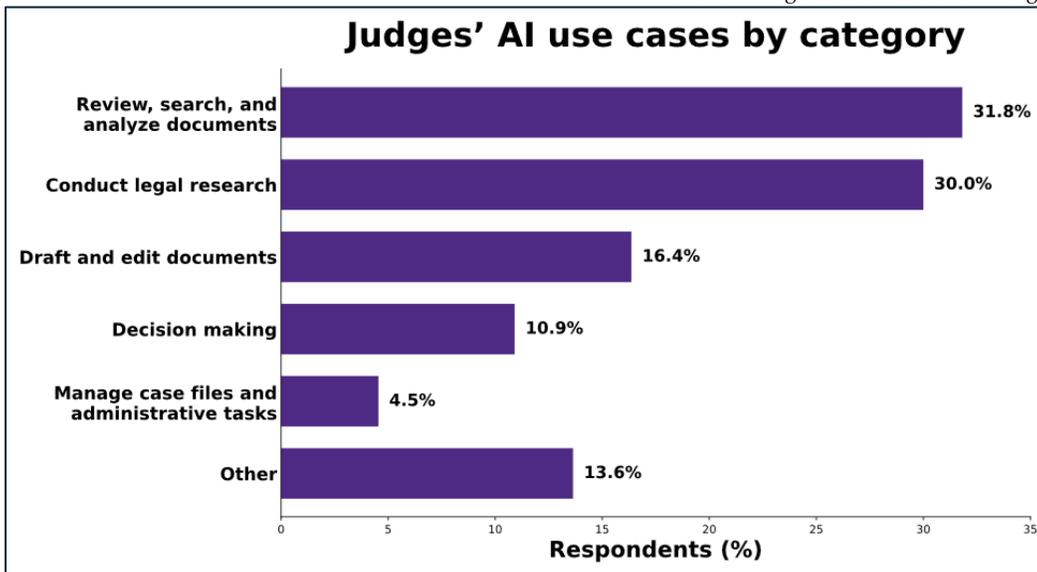


Figure 6: Judges' AI use cases by category

Judges' responses to Question 1.3 have been grouped into the categories described in Table 4.

Note: 51% of respondents (110) reported "I don't use AI in my work." This bar is not included.

The overwhelming majority of judges chose either no AI use case (50.9%) or only one (25.5%). As shown in **Figure 7**, 9.1% of judges chose two use cases, 6.5% chose three, and 8.2% of judges chose four or more AI use cases.

Click on image to download and enlarge.

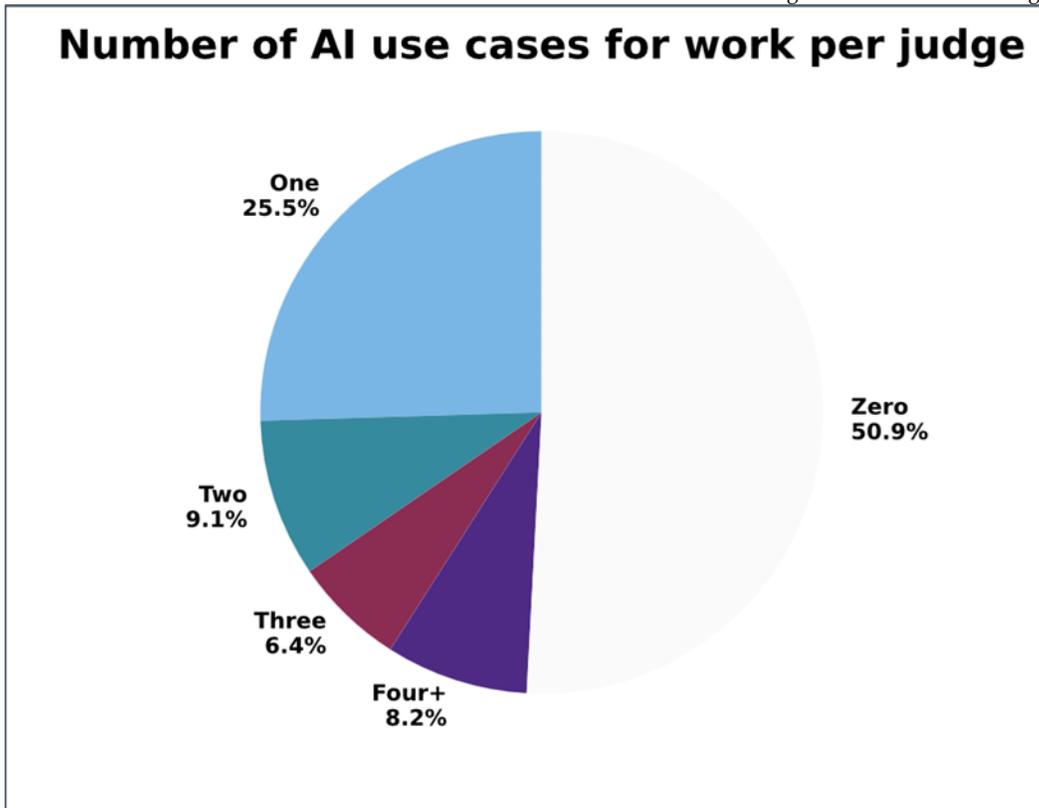


Figure 7: Number of AI use cases for work per judge

The figure above displays the percentage of judges who use AI in their work for each use case listed in Question 1.3. 110 judges responded to this question.

Question 1.4 (only if “Other” chosen for Q1.3): If you use AI for other use cases in your work as a judge, please describe the use cases.

Judges who chose “other” for Question 1.3 were asked to describe the other use cases.

Judges’ responses included:

- WestCheck to verify cases/citations.
- Fact research.
- General fact questions – seeking authoritative sources, not relying on AI answers.
- As a tool, like a treatise, before beginning legal research.

- I guess the use I make of it might be described as "legal research." I use it only when I have a pretty broad, open-ended question to help steer me towards appropriate authority or precedent.
- Reading briefs and court opinions.
- Summarize large documents.
- Preparing speeches, slides for presentations, and emails.
- Prepare materials for talks and presentations.
- I mostly use it in editing to answer word choice/grammar questions I have.
- I have used ChatGPT to prepare a first draft of a CLE program outline ... I have not used it for any case related work.
- I occasionally do a Google search and look at Gemini when it pops up.
- Look up non-legal info quickly.
- I sometimes use Copilot for work tasks that are entirely unrelated to deciding cases. For example, in preparation for attending a recent talk by a jurist from a small foreign country, I asked Copilot for some historical resources on that country.
- I use ChatGPT to understand the technology. I do not use it in my work.

TOP AI USE CASE BY OTHERS IN CHAMBERS IS LEGAL RESEARCH

Question 1.5: Do other individuals in your chambers use AI for the following use cases? Select all that apply.

- Review documents (e.g., searching, asking questions about documents)
- Conduct legal research
- Draft documents filed in cases (e.g., orders, opinions, judgments, jury instructions)
- Draft documents not filed in cases (e.g., letters, emails, articles)
- Edit documents filed in cases (e.g., orders, opinions, judgments, jury instructions)
- Edit documents not filed in cases (e.g., letters, emails, articles)
- Inform decisions
- Make decisions

- Prepare case timelines or chronologies
- Transcribe text or audio
- Summarize text or audio
- Analyze text or audio
- Prepare questions for a proceeding (e.g., hearing, trial, status conference)
- Manage case files and administrative tasks
- Other (please describe below)
- They don't use AI in their work

Figure 8 shows how judges reported AI use cases by others in their chambers. The resulting distribution is broadly similar to how they reported their own use cases of AI. Judges reported higher AI usages by others in their chambers than by themselves: whereas 50.9% of judges reported that they do not use AI in their own work, a slightly lower 45% of judges reported that others in their chambers do not use AI in their work.

Click on image to download and enlarge.

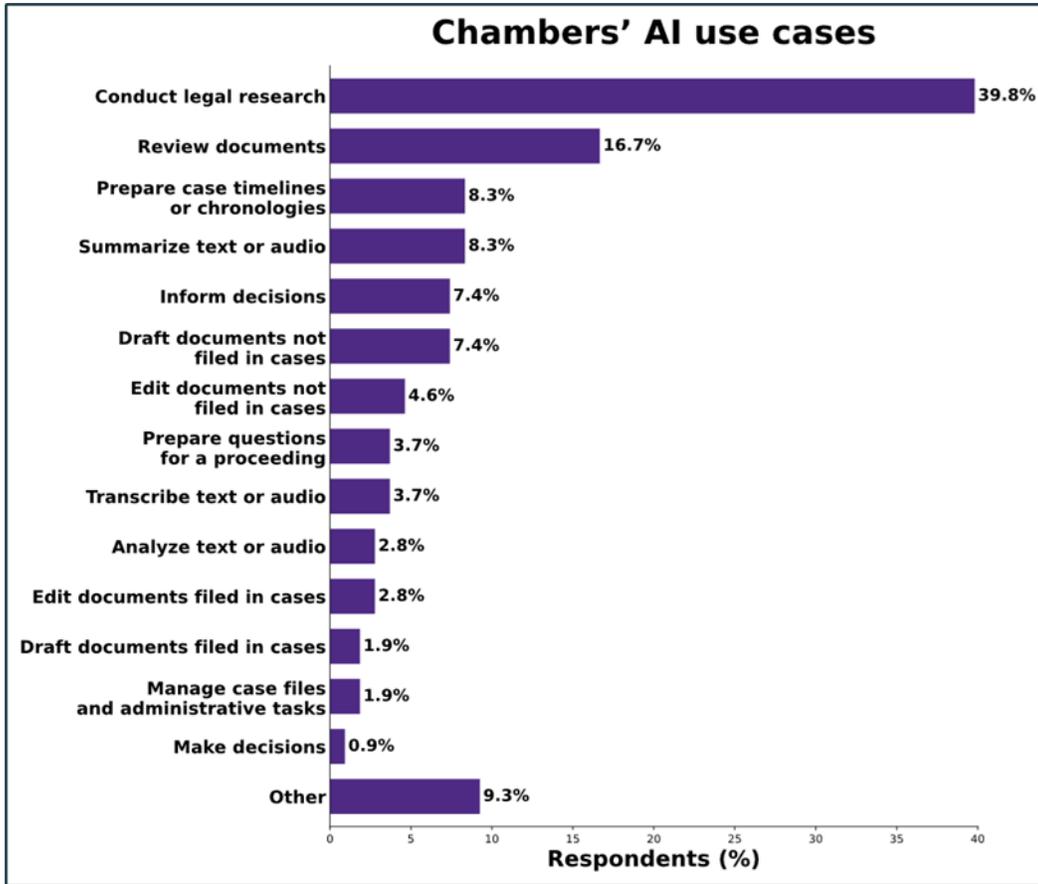


Figure 8: Chambers' AI use cases

108 judges responded to Question 1.5.

Note: 45% of respondents reported "They don't use AI in their work." This bar is not shown.

When the AI uses for others in chambers are placed into the categories introduced above in Table 4, the result is similar. As shown in Figure 9, "conduct legal research" remains the top AI use case compared to the other categories of AI use.

Click on image to download and enlarge.

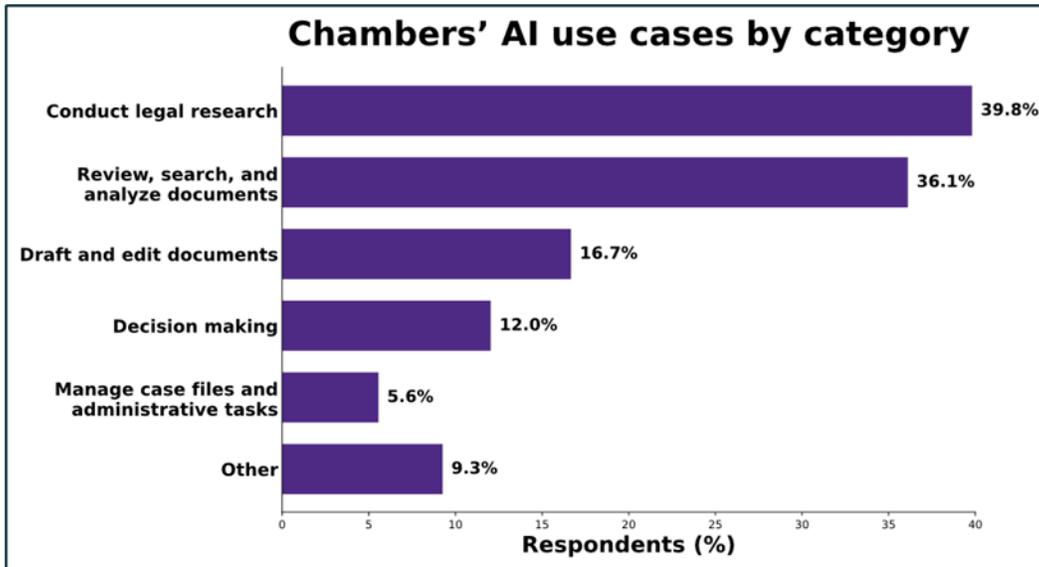


Figure 9: Chambers' AI use cases by category

Judges' responses to Question 1.5 have been grouped into the categories described in Table 4.

Note: 45% of respondents (108) reported "They don't use AI in their work." This bar is not shown.

Question 1.6 (only if "Other" chosen for Q1.5): Please describe the use cases for which other individuals in your chambers use AI.

Judges who chose "other" for Question 1.5 were asked to describe the other AI use cases in their chambers. Notably, multiple responses state or suggest that the judge does not know whether or how other individuals in their chambers use AI. One response discusses an AI tool that cited "fake" cases when it was used by a law clerk to write a legal memo, but the judge does not know which AI platform was used. The responses raise questions about the extent to which judges, lawyers, and other legal professionals understand the benefits and risks of both "General AI" tools versus "AI for Law" tools that are designed for legal tasks. For example, "AI for Law" tools provided by legal information providers have access to cases and other resources that allow them to use techniques to mitigate the risks of citing "fake" cases, although it is more difficult to eliminate other types of error (e.g., summarizing the holding of a case).

Judges' responses included:

- Translation.
- For research to the extent of using Westlaw or Lexis AI formats.
- If they use it all, they occasionally use Westlaw's AI tool to identify AI-composed briefs and documents.
- My law clerks at times have used AI to create presentations or draft remarks for presentations.
- Initial summary of an unfamiliar legal issue before conducting research on standard platforms (Westlaw or LEXIS).
- My law clerk wrote a memo for me, and then after she finished, out of curiosity, she asked AI to write a memo on the same question. Of the 11 cases AI cited in its version, 10 of them were fake. I don't know what AI platform she used, but it clearly wasn't reliable.
- Per my chambers policy, they may use it in the same way that I do.
- I am not certain whether they use any type of AI.
- I don't know.
- I have no law clerks or anyone else in my chambers.

JUDGES' AND OTHERS' AI USE IS HIGHLY CORRELATED

Figures 10 and 11 show the significant overlap in judges' reports of their own use of AI in their work and their reports of the use of AI by others in their chambers, including judicial clerks. For nearly every use case, judges reported moderately greater use by others in their chambers. Judges reported that others in their chambers use AI for legal research 9.8% more frequently than they do. The only categories in which judges reported greater use by themselves are to "prepare questions for a proceeding" and to "make decisions."

A limitation of comparing the use of AI by judges and others in their chambers is that judges are more likely to under-report AI use by others than to over-report it. The status quo had been that AI is not used in chambers. Judges may not know about the specific ways in which others in their chambers have begun to use AI. It is possible that judges think others use AI when they do not, but this seems less likely.

Click on image to download and enlarge.

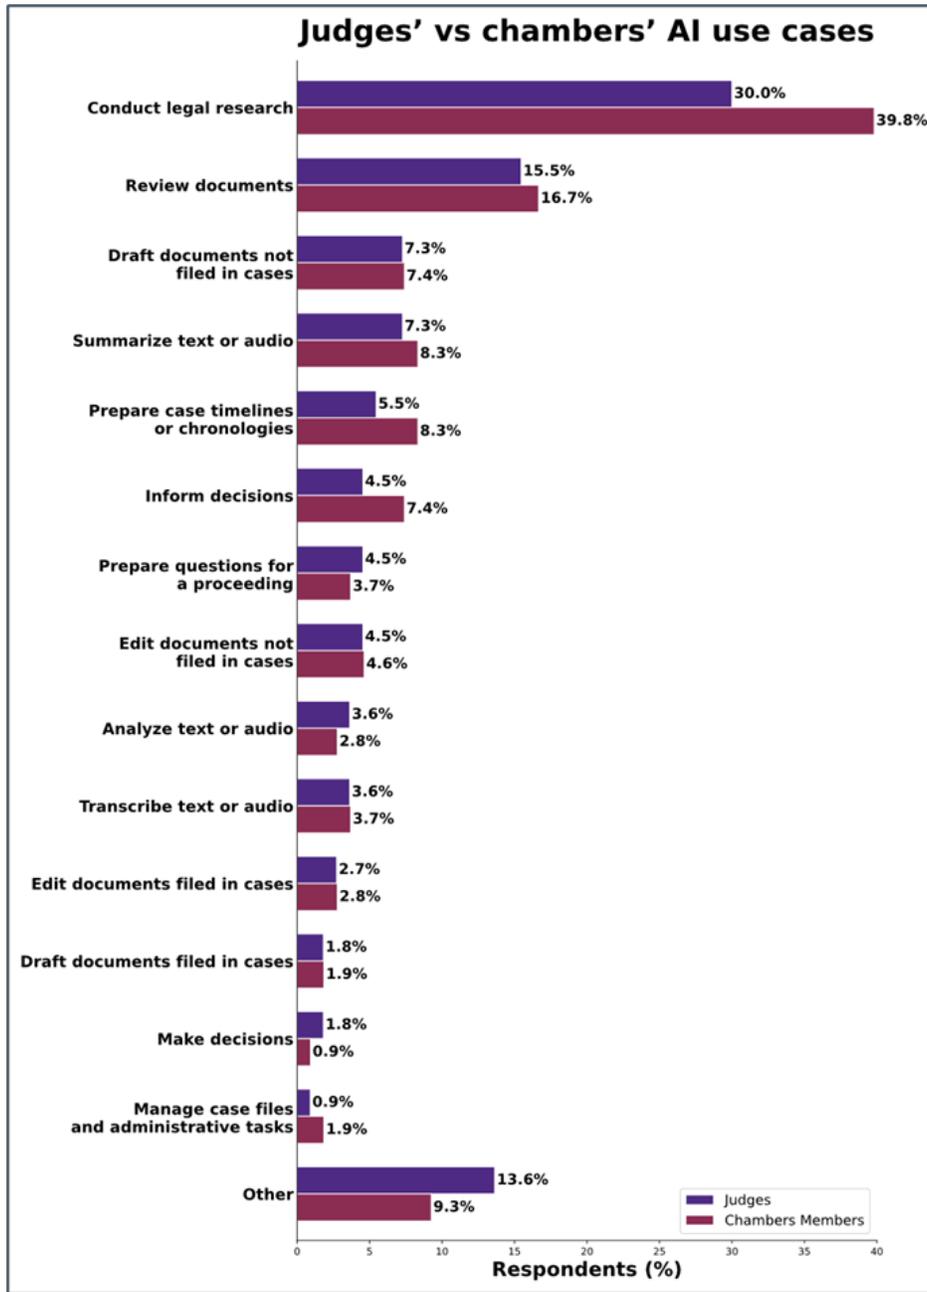


Figure 10: Judges' vs chambers' AI use cases

The figure above displays data from Question 1.3 (110 responses) and Question 1.5 (108 responses) side by side.

Note: 51% of respondents (110) in Fig. 5 reported "I don't use AI in my work." 45% of respondents (108) in Fig. 8 reported "They don't use AI in their work." These bars are not shown.

Click on image to download and enlarge.

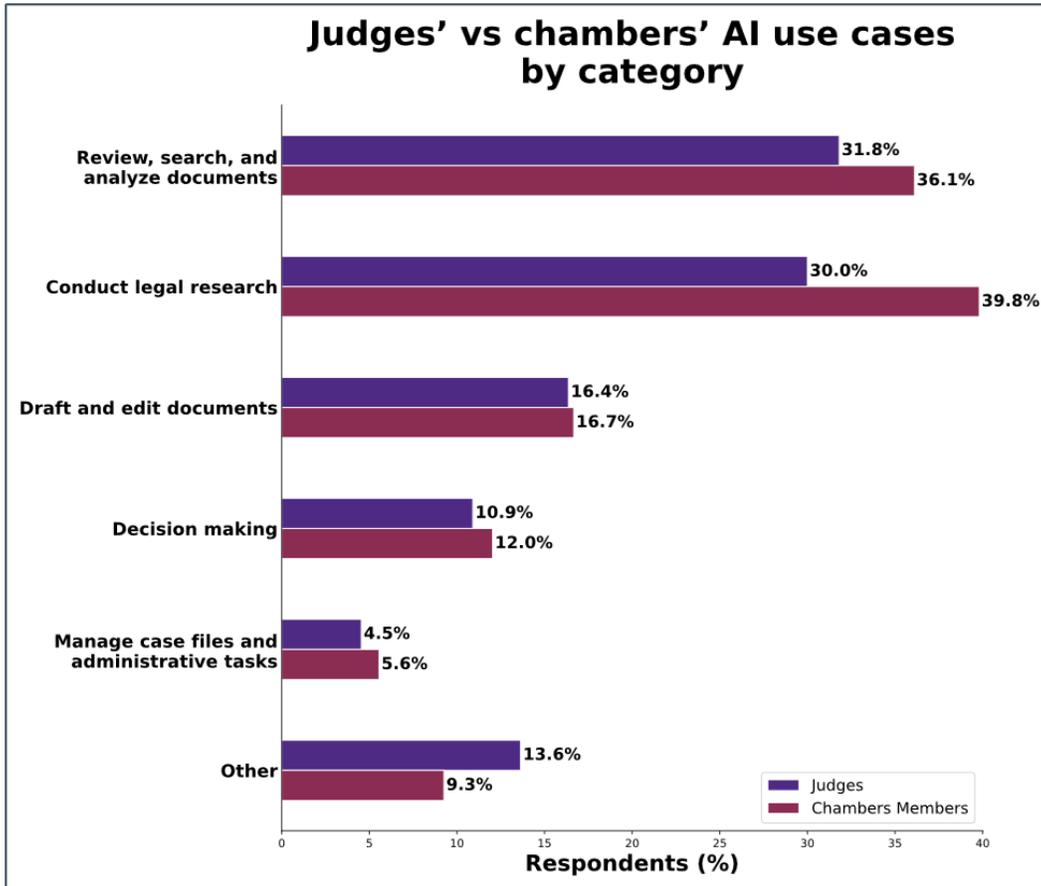


Figure 11: Judges' vs chambers' AI use cases by category

The figure above displays data from Question 1.3 (110 responses) and Question 1.5 (108 responses) side by side. Judges' responses to both questions have been grouped into the categories described in Table 4.

Note: 51% of respondents (110) in Fig. 6 reported "I don't use AI in my work." 45% of respondents (108) in Fig. 9 reported "They don't use AI in their work." These bars are not shown.

AI TRAINING HAS NOT BEEN OFFERED TO MOST JUDGES

Question 1.7: Has the court administration provided you training on the use of AI tools?
Yes, and I attended the training
Yes, but I have not attended the training
No
I am not sure

As shown in **Figure 12**, most judges (61.1%) reported either that no AI training has been provided or that they are uncertain whether such training was offered. Of the 38.9% of judges who recall AI training being provided by the court administration, a significant majority (73.8%) reported attending the training. Only 10.2% of judges responded that AI training was provided but they did not attend. This suggests an unmet demand that more AI training should be offered and specifically in a visible manner.

[Click on image to download and enlarge.](#)

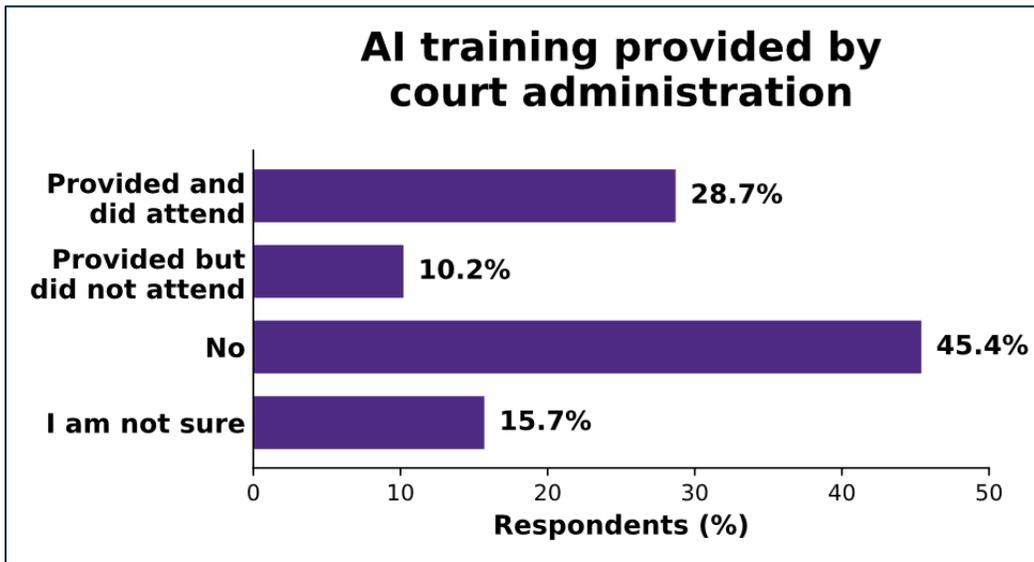
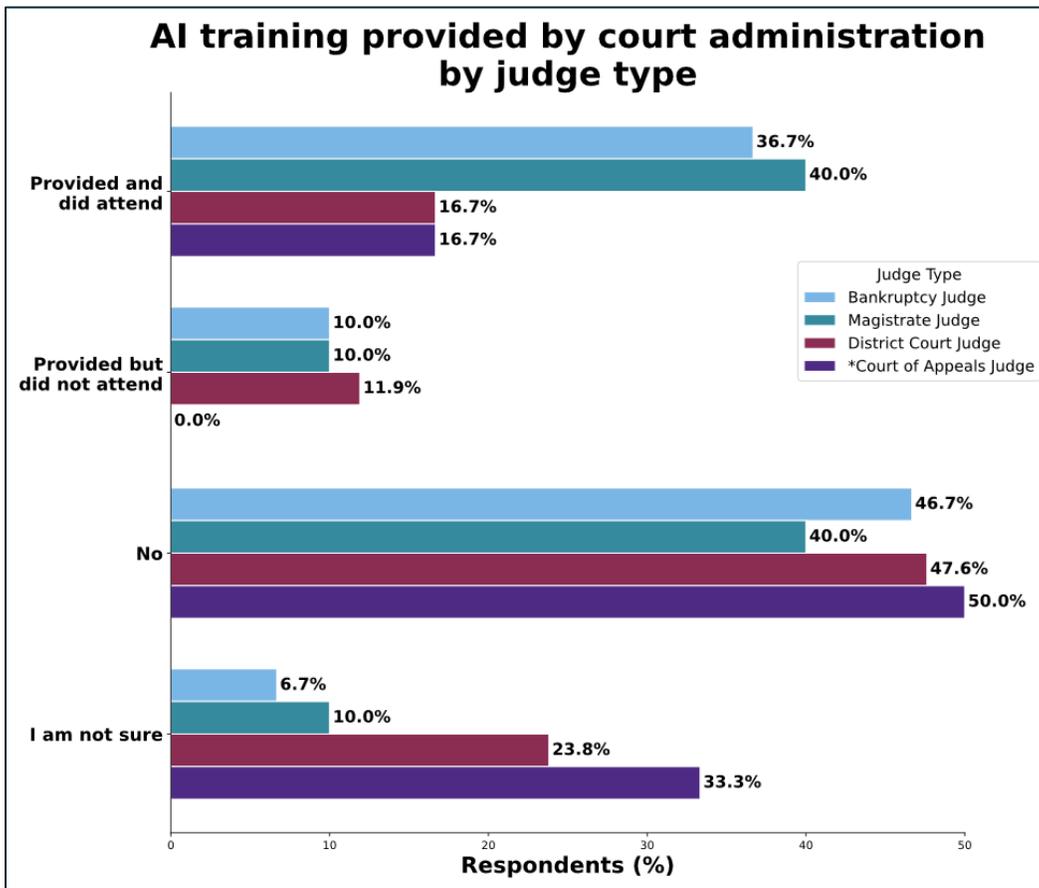


Figure 12: AI training provided by court administration

108 judges responded to Question 1.7.

Figure 13 separates the responses by judge type. Magistrate judges (40.0%) and Bankruptcy judges (36.7%) reported the highest rates of attending AI training. For District Court judges, only 16.7% reported attending AI training. In contrast, 23.8% of District Court judges responded that they were not sure whether training was offered and 47.6% said that training was not provided. Similarly, 46.7% of Bankruptcy judges and 40.0% of Magistrate judges said that AI training was not provided.

Click on image to download and enlarge.



*Low response rate. See discussion of "Limitations."

Figure 13: AI training provided by court administration by judge type

108 judges responded to Question 1.7. Each percentage displayed in the figure above reflects the percentage of that judge type that selected a particular response. The percentages of responses per judge type add up to a total of 100%.

ONE IN THREE JUDGES PERMIT AI USE IN CHAMBERS; ONE IN FOUR HAVE NO OFFICIAL POLICY

Question 1.8: What is your policy for how those who work in your chambers may use AI?
Permitted and encouraged
Permitted
Discouraged but not formally prohibited
Formally prohibited
I have no official policy
No one works in my chambers

As shown in **Figure 14**, one in three judges responded either that they “permitted and encouraged” (7.4%) or “permitted” (25.9%) the use of AI by those who work in their chambers. Slightly more than one in three judges either “formally prohibit” (20.4%) or “discourage but do not formally prohibit” (17.6%) the use of AI by those who work in their chambers.¹⁰ Meanwhile, 24.1% of judges “have no official policy” on AI use for those who work in their chambers. The distribution of chambers policies—ranging from formal prohibitions to explicit encouragement within narrow guardrails—indicates that the judiciary is in an early phase of AI governance, with no dominant model yet emerging across chambers.

10. In hindsight, “discouraged but not formally prohibited” is ambiguous. A less ambiguous response choice would have been, “Prohibited but no formal policy.” In the responses to Question 1.9, some judges who chose “discouraged but not formally prohibited” described use cases for which individuals in chambers were permitted to use AI or for which they could request permission to use AI.

Click on image to download and enlarge.

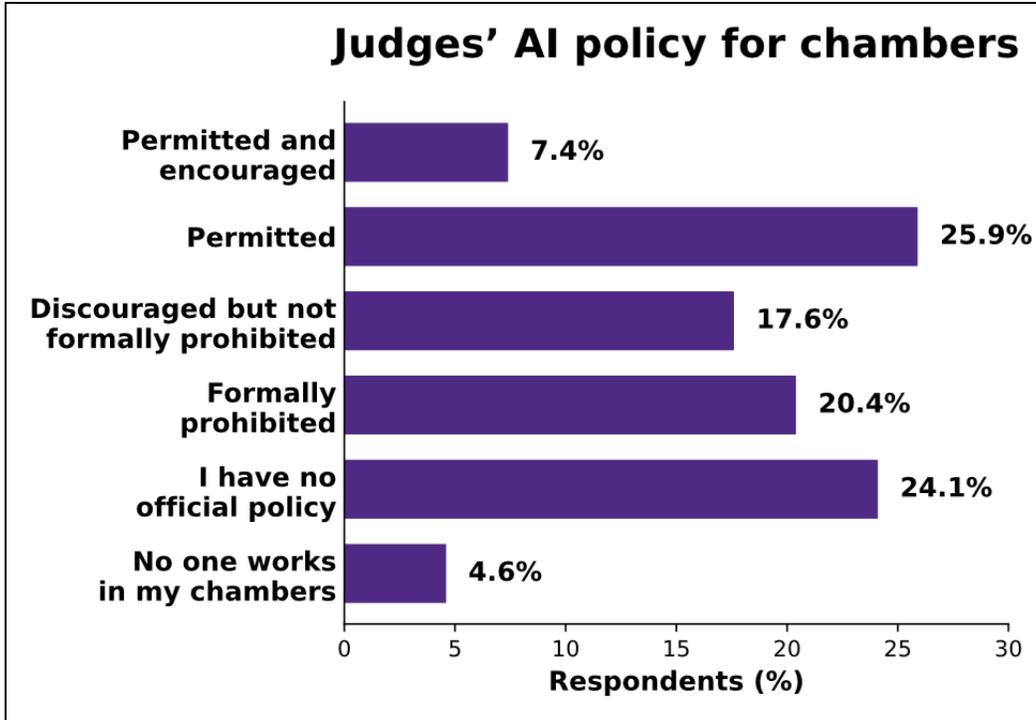


Figure 14: Judges' AI policy for chambers

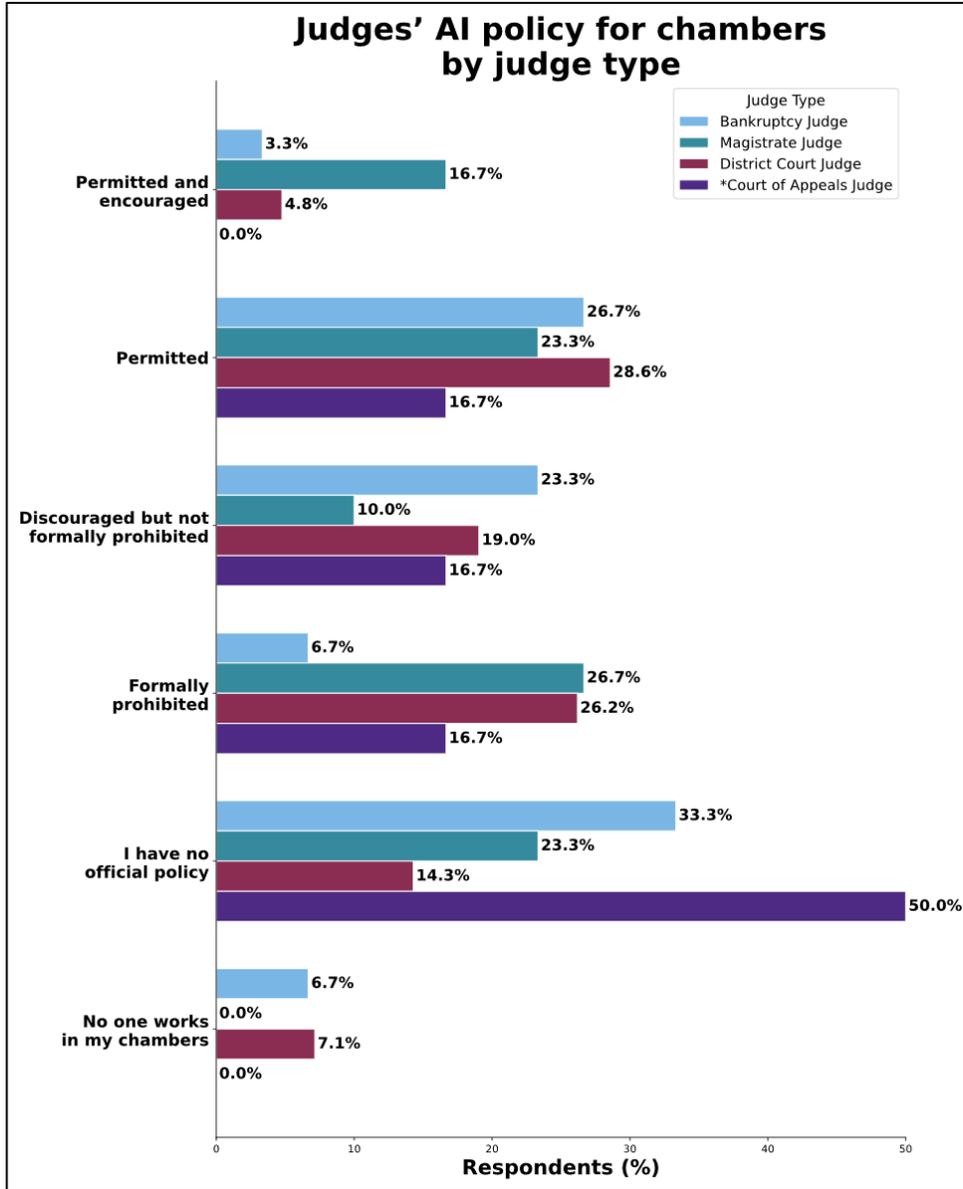
108 judges responded to Question 1.8.

This question about AI policies revealed much greater variance across different judge types than other questions, illustrated by **Figure 15**. By judge type, 40% of Magistrate judges “permit and encourage” (16.7%) or “permit” (23.3%) AI use, 33.4% of District Court judges “permit and encourage” (4.8%) or “permit” (28.6%) AI use, and 30% of Bankruptcy judges “permit and encourage” (3.3%) or “permit” (26.7%) AI use. Only 6.7% of Bankruptcy judges “formally prohibit” AI use by others in their chamber, compared to 26.7% of Magistrate judges and 26.2% of District Court judges. In response to Question 1.9, however, several judges indicated in their responses that even if they chose a variety of “permitted” or a variety of “prohibited,” there were exceptions to the general rule. This shows that some judges are carefully considering each AI use case and are not applying an all-or-nothing rule.

One in three Bankruptcy judges responded that they “have no official policy,” while 23.3% of Magistrate judges and 14.3% of District Court

judges “have no official policy.” Three out of six Court of Appeals judges (50%) responded that they “have no official policy,” although we must again advise that findings for Court of Appeals judges be met with caution because only six responded. Due to the low response rate, this finding cannot be viewed as representative of the full population of Court of Appeals judges.

Click on image to download and enlarge.



*Low response rate. See discussion of "Limitations."

Figure 15: Judges' AI policy for chambers by judge type

108 judges responded to Question 1.8. Each percentage displayed in the figure above reflects the percentage of that judge type that reported a particular policy. The percentages of responses per judge type add up to a total of 100%.

Question 1.9: Please feel free to elaborate on how you or others in your chambers use AI.

Judges were given the option to elaborate on how they or others in their chambers use AI. We have reproduced all responses below, organized by their response to Question 1.8 (i.e., “permitted and encouraged,” “permitted,” “discouraged but not formally prohibited,” “formally prohibited,” “I have no official policy,” “no one works in my chambers”).

Several judges clarified that, while AI use is (i) permitted and encouraged or (ii) permitted, even so, certain use cases are not permitted, and permitted uses can be subject to detailed guidelines (such as a requirement to independently verify AI outputs). Likewise, some of the judges who chose (i) discouraged but not formally prohibited or (ii) formally prohibited, identified use cases for which AI could be used by others in their chambers. These responses indicate that many judges are not thinking about AI in binary “yes” or “no” terms but instead are evaluating the benefits and risks of specific AI use cases.

Judges’ responses organized by response to Question 1.8 included:

- Permitted and encouraged
 - I have a firm policy, though, against AI generating content of orders, opinions, or communications.
 - Permitted and encouraged, but within very narrow guardrails. Only as part of Westlaw or Lexis research tools, and only to summarize voluminous materials - using only tools within Westlaw or Lexis so as not to disseminate court documents into the general public.
 - We have adopted strict guidelines on how our employees and law clerks may use AI.
 - Our district hasn't yet received access to any approved AI tools. Those tools remain for use by certain judges or districts based on a pilot program evaluation. Once we receive access to, and training on, these tools I plan to utilize them and allow those in Chambers to do the same.

- Permitted
 - I and my law clerks use AI as an aid to our legal and factual research. My one rule is that it is essential that all AI-generated information be independently verified.
 - Permitted with disclosure and permission consistent with court policy.
 - Cautiously. As one law clerk puts it, he only uses AI that cites its sources (Gemini) and he treats it like an eager 1st year law grad.
 - The approved AI tools in the judiciary are basically nil. I personally use available tools through Westlaw.
 - Chambers' use of AI is largely for initial or background research efforts on a particular body of law or underlying facts (controlled substances, technology, investigative techniques, etc.). We do not use it to draft or for specific citations to cases - more of a jumping off point for issues or areas we lack background knowledge.
 - Mainly used to summarize lengthy documents and establish timelines.
 - It is permitted with limitations -- namely, it can be used to assist in document searches and certain legal research, but particularly as to legal research, any results need to be verified with legitimate case checks.
 - I allow AI use, but we specifically do not allow AI to reason or write for us. In other words, we make the decision. We write the opinions. If my clerks want to use AI to help conduct research, more power to them.
 - Permitted to use CoCounsel.
 - In discussing the issues we are addressing, I will sometimes, in a brainstorming session, raise a question and then suggest that AI may be a sensible way to get a first cut at the answer to the question.
 - Primarily for legal research.
 - I allow clerks to use AI as an initial point for research within Westlaw or Lexis. I make clear to clerks that AI is a starting point and that any work submitted must reflect their own efforts.
 - In my chambers we discuss how AI might be able to help as a research aid (e.g., by helping find a jumping off point when

conventional research isn't turning up results), but I emphasize that AI is never to be trusted and always must be verified. Citing hallucinated cases or nonexistent law is a terminable offense.

- Using it intermittently to understand its use and perhaps be able to better identify legal submissions generated by AI.
- Discouraged but not formally prohibited
 - AI may only be used as a research tool, and only if the sources and authorities are independently verified.
 - Use it to summarize and transcribe court hearings.
 - I use AI to quickly look up how it would answer legal questions on issues that are very ancillary to the core issues in my cases. My clerks do the same, but they sometimes use it for deeper legal research. They also use it to "AI check" briefs that have been filed.
 - Discouraged for legal analysis. Encouraged for data summarization and non-legal work.
 - We don't use AI.
 - We don't have access to AI through our district administrators.
 - Absolutely not used to conduct research or draft orders and opinions.
 - Allowed with special permission.
 - Given concerns over fakes and inaccurate information, I am reluctant to use and as of this time have not found our chambers at a great disadvantage for failing to utilize.
- Formally prohibited
 - It's fine to use for something like a poem celebrating a birthday or anniversary. But I do not permit it for case-related work.
 - It is not prohibited for general information purposes, but it is not permitted for drafting, primary research or anything else.
 - My clerks can use AI for legal research (Westlaw) but not for other functions.
 - In my Chamber's policies, I expressly prohibit the use of AI.
 - I permitted my law clerks to participate in a trial of WestLaw CoCounsel. But after that trial, we agreed as a team that the risks of using AI tools for case research and decisions outweighed the benefits, and I established a formal prohibition on using AI for those purposes.

2026]

ARTIFICIAL INTELLIGENCE IN FEDERAL COURTS

41

- I have a formal policy that nobody can use Generative AI, again depending on how you define that term. ... There are several federal judges who are beta testing Generative AI, including Judge [REDACTED]. After the beta testing is completed and we get clear directions from the AO nobody can use Generative AI in chambers. I highly recommend you speak with Judge [REDACTED]. He's doing some great stuff with Generative AI as well as learning about the dangers. So I'm going to school on him. One of the best ways to learn in life is by other people's mistakes. ...
- I have no official policy
 - We just use the Lexis tool to formulate queries etc.
 - I do not allow my staff to use AI to draft court documents or conduct legal analysis. I am investigating tools to use for summarizing and analyzing case materials.
 - They can use the Westlaw or Lexis products for research, but not drafting.
 - Use is minimal, only for broad review of unfamiliar legal issues or for drafting non-case related correspondence.
 - My informal policy is anyone (inside chambers or lawyers filing documents) can use AI if they want to, but they are responsible for making sure the arguments they make to the court are meritorious and cases they cite to the court are legitimate. They can also look at books, search on Google, assign an associate to do research, or call their cousin I really don't care about the source, I care about the result being correct.
 - I have no official policy but I am the only one with official access to AI tools.
- No one works in my chambers
 - I would like to use AI on Lexis or Westlaw because I can trust them to find real cases. But that is not yet allowed.

38% OF JUDGES USE AI DAILY OR WEEKLY OUTSIDE OF WORK

Question 1.10: How frequently do you use AI outside of work?
<ul style="list-style-type: none">• Daily• Weekly• Monthly• Rarely• Never

Figure 16 shows judges' reported use of AI outside of work settings. A majority of judges responded that they "never" (25.9%) or "rarely" (26.9%) use AI outside of work. Only 11.1% of judges reported that they use AI "daily" outside of work, and 26.9% responded that they use it "weekly."

Click on image to download and enlarge.

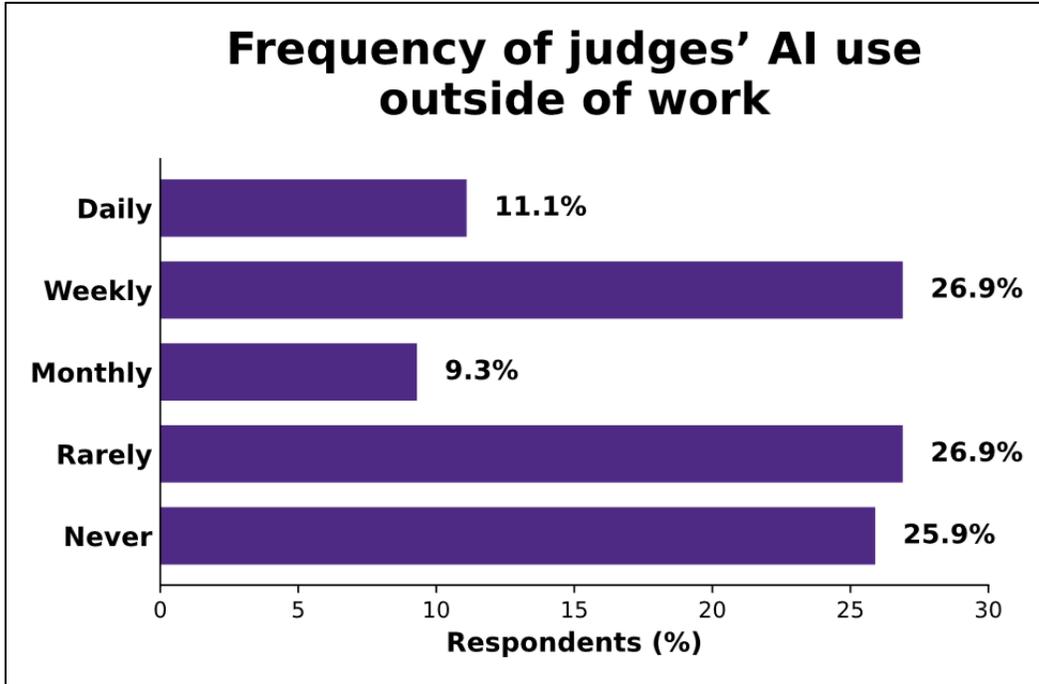
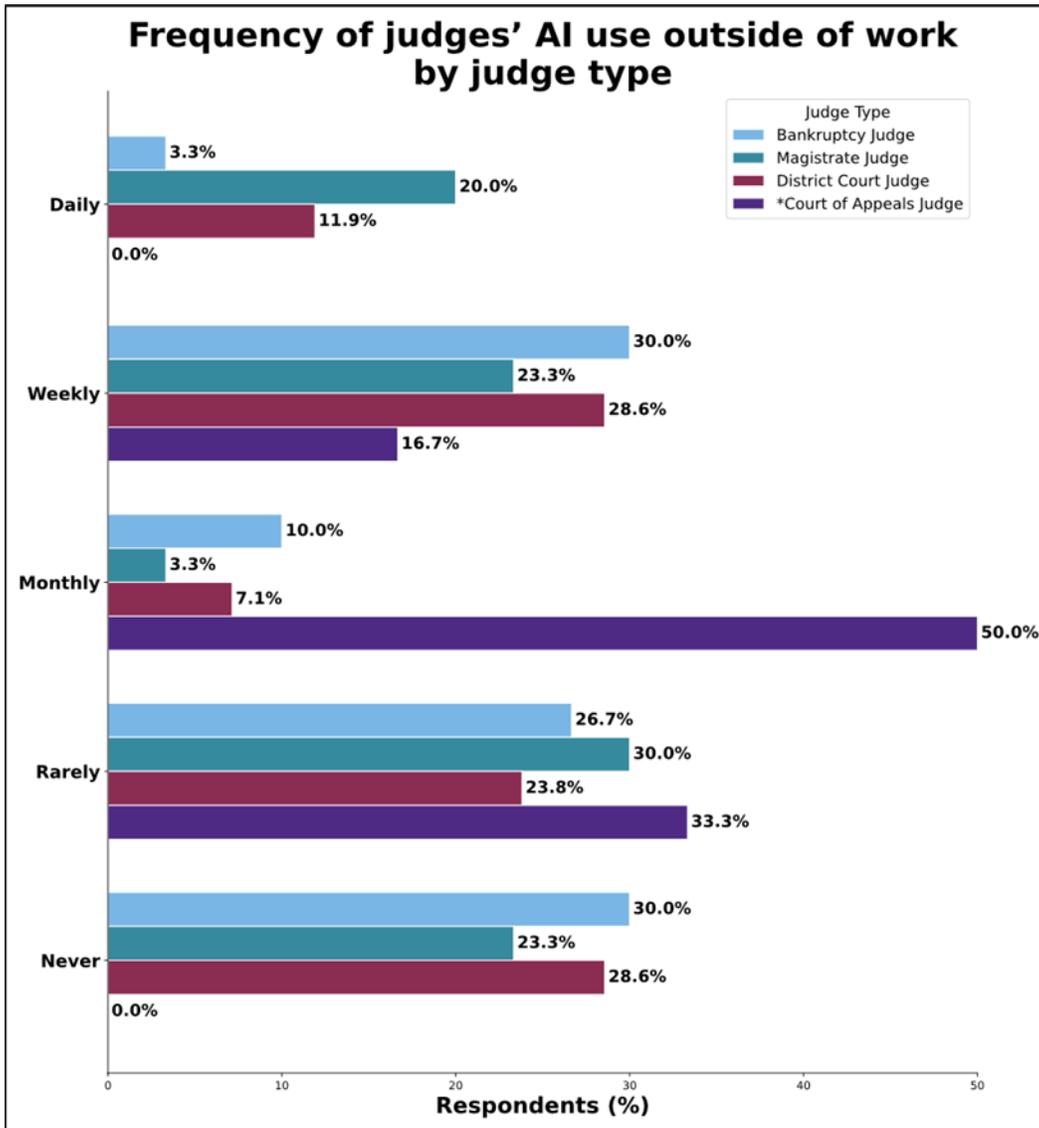


Figure 16: Frequency of judges' AI use outside of work

108 judges responded to Question 1.10.

Click on image to download and enlarge.



*Low response rate. See discussion of "Limitations."

Figure 17: Frequency of judges' AI use outside of work by judge type

108 judges responded to Question 1.10. Each percentage displayed in the figure above reflects the percentage of that judge type that reported a particular frequency of AI use. The percentages of responses per judge type add up to a total of 100%.

Question 1.11 (only if "Never" not chosen for Q1.10): What AI tools do you use outside of work and for what use cases?

Judges were given the option to elaborate on how they use AI outside of work, except if they answered that they “never” use AI in response to Question 1.10. We have provided all responses below, organized by their response to Question 1.10 (i.e., “daily,” “weekly,” “monthly,” “rarely”).

Judges’ responses organized by response to Question 1.10 included:

- Daily
 - Gemini - to answer all manner of questions and generally to ponder complex topics.
 - ChatGPT.
 - Google’s AI searching.
 - Grok and ChatGPT for everyday use - from music and restaurant suggestions to trying to create my own memes for group chats.
 - I use them every day to get answers to questions as they pop up throughout the day. I do not ever use AI to work on my cases.
 - ChatGPT.
 - General questions, prepare presentation outlines or marketing campaigns.
 - I use Gemini to revise personal emails.
 - Gemini when I do Google searches. I read the AI-generated search results.
 - ChatGPT. I use it both for my teaching, writing a book, and personal use, all generally in line with my use cases from work.
- Weekly
 - Just the AI-generated responses that Google generates when I ask Google a question.
 - Google's AI search function.
 - ChatGPT.
 - ChatGPT - social activities; general information.
 - Copilot for personal matters.
 - ChatGPT or Grok for general, personal matters.
 - General knowledge search.
 - Google.
 - ChatGPT and Gemini. Questions related to travel or trip planning. General questions that I might normally ask a Google

search engine. Also, summarizing historical articles or materials for personal education.

- Google searches.
- ChatGPT, sometimes for complicated math, sometimes for information gathering.
- Grok.
- ChatGPT for general and personal queries.
- Google AI. For simple automotive and household how-to questions.
- Gathering information.
- Trivia and social information (best restaurants, hotels, travel info etc.)
- ChatGPT; personal scheduling, planning, etc.
- ChatGPT.
- Google Search.
- ChatGPT, Claude.
- I look at the AI-generated result at the top of a Google search when I perform searches on the internet. And again, I do not trust and always verify if it's important.
- Grok, for [hobby] questions and product comparisons Occasionally, I'll use it for conversion calculations.
- Google searches.
- Miscellaneous research.
- Monthly
 - General factual questions and information.
 - ChatGPT, planning itinerary for travel, drafting documents for not for profit.
 - ChatGPT. Have used to make outlines. Used to plan a trip.
 - ChatGPT. Mainly for travel tips, general knowledge searches, thesaurus.
 - Co-pilot for quick (factual) answers to random questions.
 - I use ChatGPT occasionally for personal tasks like planning a menu or writing a poem for a special occasion.
 - Perplexity, ChatGPT.
 - Follow current events.
 - ChatGPT and Google's AI. Search, restaurant recommendations, plan trip itinerary.

- ChatGPT.
- Rarely
 - Helping plan vacations.
 - ChatGPT.
 - ChatGPT.
 - I use ChatGPT to create prayer or reflection drafts for inspirational messaging during non-profit meetings.
 - Google searches.
 - Copilot. Drafting itineraries, locating helpful phrases and creating clever poems/notes.
 - I might glance at Gemini when it pops up after I do a Google search.
 - ChatGPT. To make funny photos.
 - ChatGPT - problem solving for mechanical issues around the house - appliance/vehicle repair. Drafting sympathy notes and other personal notes.

JUDGES' USE OF AI IN THEIR WORK IS CORRELATED WITH PERSONAL AI USE

Judges' use of AI in their personal lives is correlated with their use of AI in their work, as shown in **Figure 18**. We performed the Chi-Square test to determine numerically whether the association exists, and Cramér's V to measure the strength of that association. The chi-squared test resulted in a very small p-value of 4.30×10^{-6} (much less than 0.05), providing strong statistical evidence that judges' personal and professional use of AI are associated. A Cramér's V of 0.355 suggests a moderate strength association between the two variables.

One in five judges (20.4%) said that they “never” use AI in their personal lives and “never” use AI in their work. More than two in five judges (43.5%) reported that they “never” or “rarely” use AI in their personal lives and that they “never” or “rarely” use AI in their work.

Click on image to download and enlarge.

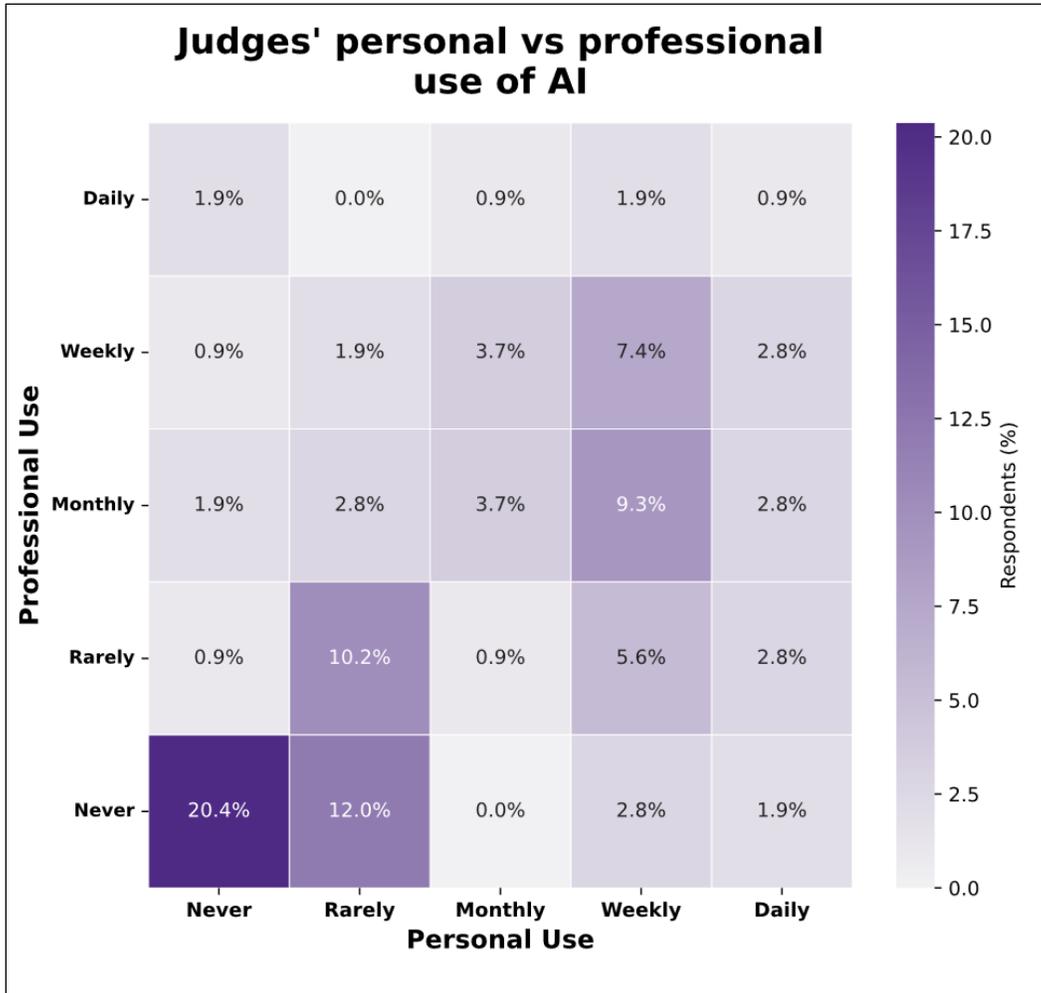


Figure 18: Judges’ personal vs professional use of AI

The figure above displays data from Question 1.1 (108/112 responses) on the y-axis and Question 1.10 (108 responses) on the x-axis.

JUDGES ARE EQUALLY DIVIDED ON OPTIMISM VS CONCERN ABOUT AI IN THE JUDICIARY

Question 2.1: Which of the following best describes your general outlook on AI’s potential for the judiciary?

- Very concerned
- Somewhat concerned
- Neutral
- Somewhat optimistic
- Very optimistic

As shown in **Figure 19**, there is a nearly even split between judges who are optimistic about AI and those who are concerned.

[Click on image to download and enlarge.](#)

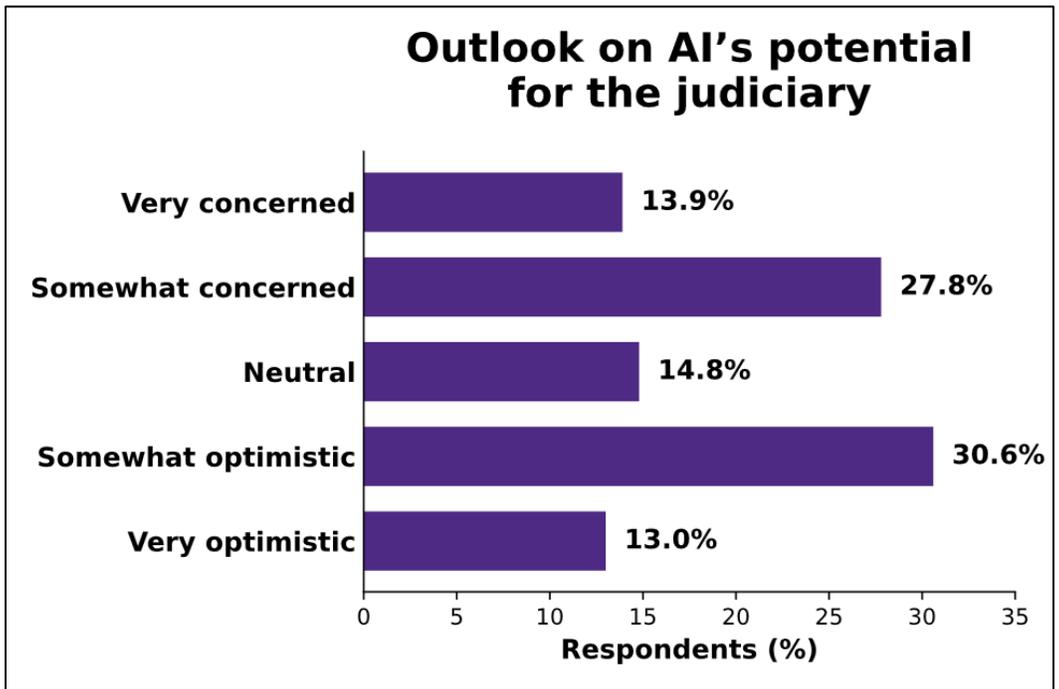
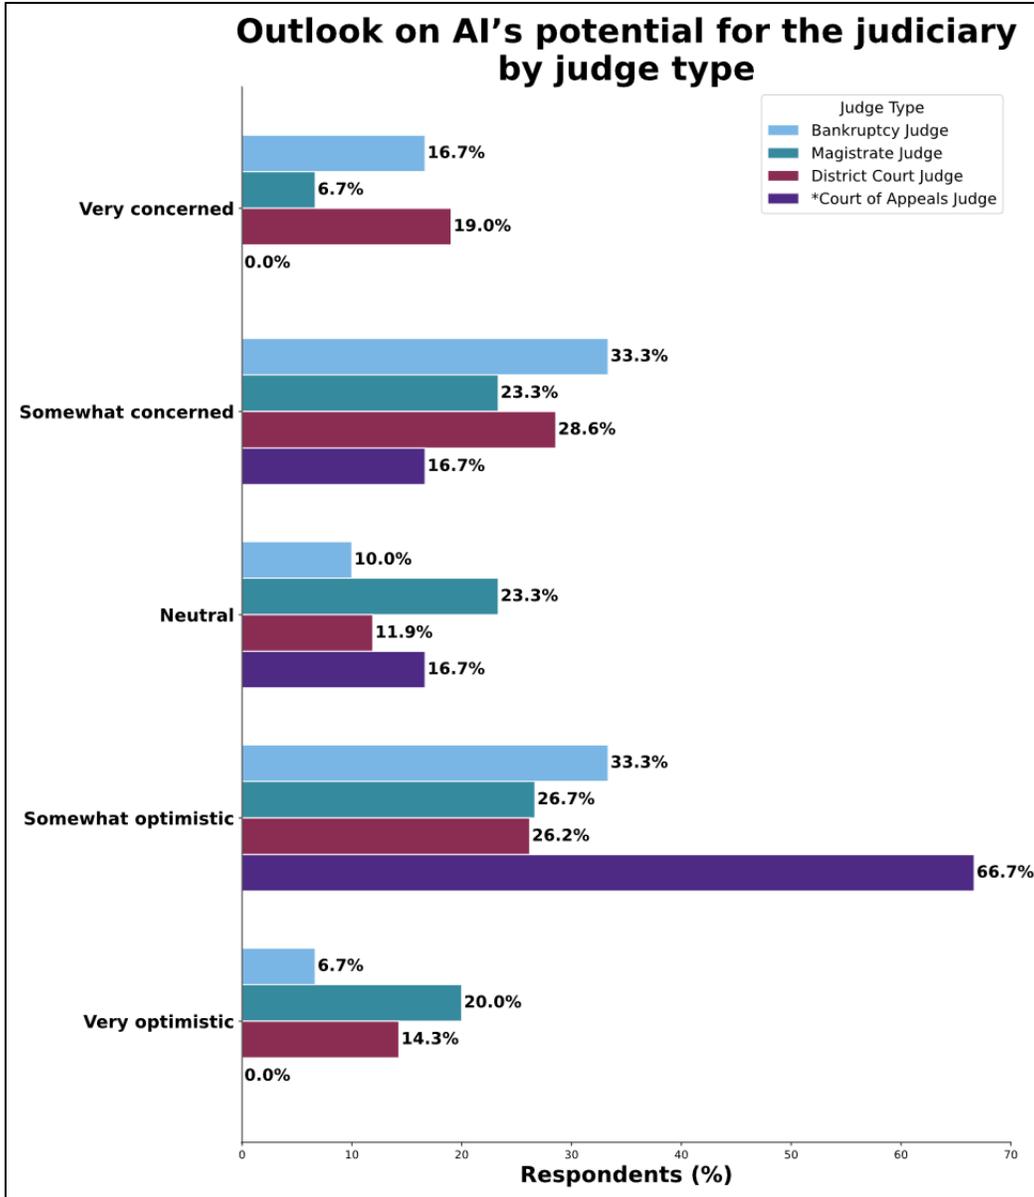


Figure 19: Outlook on AI’s potential for the judiciary

108 judges responded to Question 2.1.

Looking at the outlook on AI's potential for the judiciary by judge type in **Figure 20**, Magistrate judges were more optimistic than concerned, while Bankruptcy and District Court judges were more concerned than optimistic. Among Magistrate judges, 46.7% responded that they were "very optimistic" (20.0%) or "somewhat optimistic" (26.7%) while 30% reported that they were "somewhat concerned" (23.3%) or "very concerned" (6.7%). Among Bankruptcy judges, 50% responded that they were "very concerned" (16.7%) or "somewhat concerned" (33.3%) while 40% responded that they were "somewhat optimistic" (33.3%) or "very optimistic" (6.7%). Among District Court judges, 47.6% reported that they were "somewhat concerned" (28.6%) or "very concerned" (19.0%) while 40.5% reported that they were "somewhat optimistic" (26.2%) or "very optimistic" (14.3%).

Click on image to download and enlarge.



*Low response rate. See discussion of "Limitations."

Figure 20: Outlook on AI's potential for the judiciary by judge type

108 judges responded to Question 2.1. Each percentage displayed in the figure above reflects the percentage of that judge type that reported a particular outlook. The percentages of responses per judge type add up to a total of 100%.

Question 2.2: Please feel free to elaborate on your attitude toward AI's potential for the judiciary.

Judges were given the option to elaborate on their attitude toward AI's potential for the judiciary. We have provided all responses below, organized by their response to Question 2.1 (i.e., "very optimistic," "somewhat optimistic," "neutral," "somewhat concerned," "very concerned").

Judges' responses organized by response to Question 2.1 included:

- Very optimistic
 - Under controlled circumstances to weed out fake stuff.
 - Summarizing trial transcripts and voluminous documents and pinpointing instances of specific testimony in a closed universe environment is a huge time saver.
 - It appears it could be a useful tool for developing case chronologies, findings of fact, or for drafting short routine/common orders.
 - The judiciary is woefully behind in adopting safe AI protocols, tools and training. When used properly AI can be of enormous benefit to the court.
 - I believe it will be a significant benefit to conserving judicial resources. So long as accuracy can be confirmed.
 - We need an AI tool tailored for the judiciary that does not share information. We also need one that NEVER makes up cases.
- Somewhat optimistic
 - It could help with efficiency and ability to accomplish more in less time. But it will not replace people and critical thinking.
 - If fact checked it can be very useful.
 - I think legal-specific AI programs (e.g. Westlaw's) have the same potential now that original Westlaw and other databases had when our profession was using hard copy books for research. But I am wary of programs that are not actually designed for lawyers and judges to use.
 - I think AI can be a useful tool once it becomes more reliable.
 - I think there are tools that can be useful to the judiciary. I

- approach with a trust but verify approach.
- I do think that that the tools are quite powerful. There are obviously risks and challenges. Navigating them will be complex. But I do believe that if used appropriately and cautiously, the tools can help us improve our work.
 - I think there is a role for AI in administrative functions that can make us more efficient.
 - We need training.
 - I think it can be very helpful for summarizing documents. I also like Westlaw's AI summary of negative treatment a case has received.
 - I think for research it can be helpful. I worry about it being used for legal analysis.
 - I would appreciate training on the use of AI, provided the use will increase efficient responses by chambers to litigation, and will offer tools that are as trustworthy as Westlaw or Lexis.
 - I believe AI is inevitable and that the courts should use it.
 - I think it is inevitable that AI will change some work processes and that the judiciary should actively try to figure out how best to use this new tool while maintaining the integrity of our work.
 - Optimistic that the judiciary can responsibly utilize this new tool in its work product.
 - I believe it will enhance research capabilities and will provide time savings on tasks such as summarizing volumes of exhibits and creating first drafts of timelines.
 - As new law clerks with a good grasp of AI come onboard, I expect there will be a "trickle up" effect of usage in the judiciary.
 - AI can be enormously helpful, provided the user knows that it regularly generates fictional responses and therefore needs to be rigorously checked. My principal concern is that pro se litigants now regularly use AI to generate lengthy pleadings that look sophisticated but are often frivolous; this multiplies the amount of work I and my law clerks have to do in those cases.
- Neutral
 - I'm optimistic that AI can help us become more efficient (e.g., by generating skeleton drafts based on template documents that we

use frequently, or in the future automatically adding new motions into our chambers workflow or otherwise helping with the burden of more administrative tasks), but I am highly concerned that AI is causing younger generations of lawyers and laypeople not to think critically and to lose essential research and writing skills. I also am concerned that AI may help litigants with frivolous cases prolong the process.

- In the long term, neutral. It is a tool. So is email and electronic case filing. There will be short term disruption while we all figure it out, then we will wonder how we lived without it. One of my concerns is that it will overburden judges, as aggressive and litigious pro se litigants use AI to churn out large volumes of motions and arguments that judges will need to address. The internet did that, too, but AI will increase how often this happens. And even ordinary pro se litigants who are trying to do the right thing and who rely on AI may not realize that the brief they had AI generate that looks pretty good to them is egregiously wrong about the law.
- It's a tool, but not a rudder. Best use is to analyze attorney's filings. Also, insofar as it could help decipher pro se handwritten pleadings, it would be amazing.
- I am open to the judiciary looking at enterprise-only AI solutions, so that the judiciary's use of the solution would not entail sending information outside of the judiciary.
- If it would work effectively on summarizing or even transcribing hand-written pro se filings, that would be incredibly helpful.
- Somewhat concerned
 - Ease of abuse by litigants.
 - If all judges, clerks, and other assistants were uniformly careful, ethical, and professional, I wouldn't be concerned about the possibility of AI-generated hallucinations corrupting research tasks and the like. Unfortunately, you can't rely on all people to uniformly practice good hygiene in their use of AI. Someone who is a poor writer, poor researcher, or simply wants to take the easy way out can be easily tempted to rely on AI-generated writing or research without taking the necessary precautions against hallucinations or AI's inclination to want to "please" the person

asking the questions. My average caseloads over the years has varied from 3500 cases to 5500 cases at any one time. I am unlikely to uncover hallucinated cases and quotes from lawyers' briefs and motions unless opposing counsel spots those and brings those to my attention. Sanctioning counsel for ethical lapses is my least favorite activity as a judge, but I believe that stiff sanctions are generally appropriate when counsel fails to police their use of AI adequately.

- I have no problem with the use of AI as a tool. It is not a substitute, however, for one's own legal research and writing.
- Concerned that it will degrade lawyers' ability to conduct legal analysis.
- I don't like it, but I don't know or see how we avoid it. We better come up with some guardrails.
- AI hallucinations concern me and the recent instances of judges incorporating incorrect AI findings in their filed opinions prevents me from exploring AI further at this time.
- Given recently publicized events, I am concerned that judges do not have proper guardrails in place to prevent the inappropriate use of AI.
- If used properly it could be of huge benefit, but there do not seem to be sufficient guardrails at this point, and I have seen numerous false citations in briefs filed by attorneys and particularly pro se litigants that appear to have arisen from the use of AI.
- Concerned to the extent individuals do not thoroughly check case cites, facts and holdings. Also from a standpoint of litigants' use-should not rely on AI to formulate initial theories or ideas in a case...
- The consistent reports of zombie cases and other instances where AI conjures law or facts is terrifying and forms the basis for how we use AI in chambers.
- Even with human employees, I worry that when giving assignments or editing their work I'll gloss over some nuance that I would have caught if I did the work myself. I believe that risk is multiplied when using AI, both because of any hallucinations, and also because current AI isn't actually intelligent.

- I think it can be a good tool, but we need to make sure AI results are accurate.
- AI has value so long as the cited sources are verified.
- Very concerned
 - We have no clue how Generative AI will affect the judiciary. Currently, we're just guessing. And even with our best guesses we are likely to come up short. The proposed FRE regarding deep fakes is a prime example. Years from now how we've tried to anticipate and deal with deep fakes will be used as the example of how the judiciary is stumbling around in the dark. [Some] are doing their best to address the issue, but I'm concerned not enough or the correct people are listening.
 - My [spouse] teaches ... and ... has sensitized me to the harmful effects that AI is having on students' ability to think and write for themselves. The undergraduate students of 2025 are the law clerks of 2030, so yes, I'm concerned.
 - Although I think AI could be useful in summarizing motion papers and assisting with initial research, I am very trepidatious to rely on it for in depth research and writing.
 - I guess it may be helpful as a preliminary tool to organize thoughts prior to a complete from-the-ground-up review by someone, or perhaps as a way to propose alternate legal conclusions after Chambers has arrived at an answer. But use of AI as a substitute for legal research and analysis seems to me the beginning of a sea change of atrophying judicial skills. So my attitude is not positive.
 - I am concerned about the ability of AI to hallucinate cases and quotations from fictional cases. I am very concerned about the prospect that hallucinated evidence will be presented to the court.
 - Making decisions requires deep thinking and understanding.
 - I want my clerks and lawyers who appear before me to hone critical thinking skills as well as research and writing. I am considering some uses of AI to check filed pleadings to make sure the cases are real.

Question 2.3: We encourage you to add other comments and information about how AI is being used in the judiciary, including the possible benefits and risks.

Judges were given the option to provide additional comments and information about how AI is being used in the judiciary, including the possible benefits and risks. We've provided all relevant responses below, organized by their response to Question 2.1 (i.e., "very optimistic," "somewhat optimistic," "neutral," "somewhat concerned," "very concerned").

Judges' responses organized by response to Question 2.1 included:

- Very optimistic
 - The adoption of AI tools is lagging, perhaps because of the lack of understanding and fear about how the tools can be misused. We have seen very public instances of the judiciary getting burned because of the misuse of AI. There is a real need for clear guidelines, training and oversight. But also a real need for the judiciary to provide these tools for judges and their staff.
- Somewhat optimistic
 - Concerned about the litigants, and especially pro se litigants, using AI irresponsibly without any guardrails. There have been many instances thus far where AI documents were filled with "bad" AI.
 - The risk is infection of court email and operating systems.
 - I lack information and experience sufficient to provide additional comment, although I have reservations, based on publicity concerning hallucinated research results.
 - I think it can be useful, but it must be used with caution and great care -- and always verified.
 - Courts need to be aware of and to address the potential and also the risks of AI for the judiciary. It is appropriate to be mindful of the different opportunities and risks for unrepresented individuals, and to consider ways that AI can enhance - or at least not inhibit - access to justice.
 - There is an AI Task Force that is looking into AI policies for the judiciary.

- AI has been used for years - spell check, etc. I think we will continue to use what we are comfortable with and hopefully evolve some over time. A large risk might include affecting credibility (as is in the legal news these days). But, it can be a good jumping off point for many uses.
- Neutral
 - The risks are obvious--AI makes things up, and some judges are not going to double check their work and will be embarrassed when some fake case or laughably stupid argument slips through. I'm personally unlikely to intentionally use AI until I can count on it to not make stuff up. Some people want to use AI to summarize facts or large amounts of data. That sounds useful and intriguing. But, I've heard from one colleague that used AI to summarize facts from a transcript that AI actually made up a witness that never testified and included that in the summary. If the technology isn't reliably accurate, I'm not particularly interested in wasting my time with it.
 - AI should always be seen as at most a research aid, and its results must always be verified. And there is no, absolutely no replacement for reading the materials submitted by the litigants; AI should never be used or viewed as a shortcut for fulsome review of the record.
- Somewhat concerned
 - As AI performance improves, it certainly could be a valuable tool in making certain tasks more efficient; however, its inaccuracies and potential for problems stemming therefrom will take time to instill comfort before using more extensively.
 - I am concerned about AI use by pro se litigants especially.
- Very concerned
 - Risk of not deep thinking through important decisions.
 - Legal research and analysis are skills. Theoretically AI can come up with potential legal answers no one has thought of. But in practice that's not how it works. AI does not make attorneys or judges better in the long term - instead it teaches them to be more reliant on AI in the long-term, to the detriment of their own intellectual development. If the goal long-term is to replace judges with AI, then fine, we should embrace it. But if not, judges

and clerks need a diet of research and analysis to maintain critical thinking skills, not easy AI answers.

- I currently have implemented strict restrictions, and anyone caught using it will be likely lead to termination. However, if our team is trained to use AI judiciously, I would be amenable to trying to implement its use in the future.
- Again, I will use it for "give me the high points of Malawian history" or "tell me who's starting for the [football team] this week," but I won't go anywhere near it for deciding cases. If I had published an opinion with hallucinated citations, I'd have to give serious consideration to resigning.

LIMITATIONS

It is important to consider the limitations when interpreting any survey results. While these limitations must be considered, our methodology, survey questions, and response rate provide a good foundation for suggesting that our findings provide useful information about how federal judges are using AI and their attitudes about AI.

SAMPLE SIZE: 112 OF 1,738 FEDERAL JUDGES

Upon completion of the survey period, 112 judges responded to the survey that we emailed to 502 federal Bankruptcy, Magistrate, District Court, and Court of Appeals judges. This is a 22.31% response rate. Given that we identified a study population of 1,738 federal judges, we surveyed 6.44% of our study population. If 50% of judges had answered a specific question the same way, with 112 judges in our sample out of a study population of 1,738 judges, at a confidence level of 95% the margin of error would be +/- 9.0% (i.e., the actual percentage of judges in the population who would answer the question that way is between 41.0% to 59.0% 95 times out of 100).

Focusing on the sample size of all judges, it is large enough to allow us to suggest in statistical terms some confidence that it is representative of the full population of all judges subject to a margin of error, which would vary depending on the percentage of judges that answered a question in the same way.

SAMPLE SIZE FOR EACH CATEGORY OF FEDERAL JUDGE

Response rates varied by judge type, with particularly low participation among Court of Appeals judges. Estimates for that group should therefore be interpreted as descriptive of the respondents only, and not as representative of all Court of Appeals judges.

Only six Court of Appeals judges responded of the 51 to whom we emailed a survey, a response rate of 11.8%. This is six of 177 Court of Appeals judges, which is too few to suggest that this sample of six is representative of the population of 177 Court of Appeals judges. If 50% of Court of Appeals judges answered a question in the same way, at a confidence level of 95% the margin of error would be +/- 39.4% (i.e., 10.6% to 89.4%).

The second-lowest response rate was from Magistrate judges. Of the 177 Magistrate judges to whom we emailed a survey, 32 responded, for a response rate of 23.6%. This is 32 of 612 Magistrate judges. If 50% of Magistrate judges answered a question in the same way, at a confidence level of 95% the margin of error would be +/- 16.9% (i.e., 33.1% to 66.9%).

The third-lowest response rate was from District Court judges. Of the 182 District Court judges to whom we emailed a survey, 43 responded, for a response rate of 18.1%. This is 43 of 630 District Court judges. If 50% of District Court judges answered a question in the same way, at a confidence level of 95% the margin of error would be +/- 14.4% (i.e., 35.6% to 64.4%).

The highest response rate was from Bankruptcy judges. Of the 92 Bankruptcy judges to whom we emailed a survey, 31 responded, for a response rate of 33.7%. This is 31 of 319 Bankruptcy judges. If 50% of Bankruptcy judges answered a question in the same way, at a confidence level of 95% the margin of error would be +/- 16.7% (i.e., 33.3% to 66.7%).

SELF-SELECTION BIAS, NON-RESPONSE BIAS, AND SOCIAL DESIRABILITY BIAS

Selection bias occurs when the respondents to a survey are not representative of the target population. Our choice of taking a stratified random sample of judges addresses a common cause of selection bias: non-random sampling. But even with random sampling, self-selection bias, non-response bias, social desirability bias, and other biases can result in unrepresentative survey results. For example, perhaps judges who use AI were more likely to respond to a survey on AI use due to their interest. Or perhaps they were less likely to respond, either due to a lack of interest or because they did not wish to reveal their use of AI, even though the survey promised to maintain the confidentiality of personally identifiable information. Judges may also be more likely to respond when they have strong opinions either for or against using AI. Judges might under- or over-report AI use depending on how they see this correlated with prestige, perhaps as protecting the traditions of the judiciary or bringing the judiciary into the modern era. These and other risks to the generalizability of survey data must be recognized and without additional studies and data, it is difficult to say what effect they had on the results.

NOT DEFINING ARTIFICIAL INTELLIGENCE

We decided not to provide a definition of artificial intelligence (AI). The primary risk is that judges may have interpreted AI too broadly or too narrowly when they answered the survey questions. Because AI is not uniformly defined in practice, some judges may under- or over-report AI use depending on whether they regard tools such as text prediction, spellcheck, integrated research features, technology assisted review, etc. as AI.

WHETHER AI TOOLS OR TASKS WERE IDENTIFIED

We created a list of AI tools and AI use cases in the judiciary. Some tools or practices (including emerging or specialized applications) may not have been captured, contributing to measurement error in specific tool or use-case estimates. The risk of underreporting was somewhat mitigated by including a follow-up question allowing respondents to identify other AI tools and the choice “other” in the list of AI use cases. But we would expect under-reporting when it required respondents to identify AI tools or AI use cases when they were not in our list. Over-reporting could also occur due to use of the identified AI tools in unexpected ways, vague categories of AI use cases that include unexpected tasks, etc.

JUDGES AS PROXY REPORTING ON OTHERS IN CHAMBERS

Asking judges to provide information about how others in their chambers use AI is known as proxy reporting, which can produce measurement error. Judges may not always know how law clerks or staff are using AI, which may lead to under- or over-reporting of AI use by others in chambers. Even when judges know generally how others are using AI, they might not have accurate information about the AI tools used or the AI use cases.

RESULTS LIMITED TO FEDERAL JUDGES

Our results are based on a stratified random sample of Bankruptcy, Magistrate, District Court, and Court of Appeals judges. We did not include other types of federal judges.

Judges outside the federal judiciary may have a different approach to AI use. Most judges within the United States are state judges, not federal. States vary in how they are approaching AI in the judiciary, including some

states explicitly stating how AI may or may not be used. Additionally, AI tools and funding available in state courts also varies significantly, which would be one factor that could lead to significant variation in approaches to AI use across different states.

APPENDIX: SURVEY QUESTIONS

Welcome

You have been selected to participate in a research study on the use of artificial intelligence (AI) tools in the judiciary. The goal of this project is to develop an understanding of judicial use and perception of AI tools that will inform judges, court administrators, legislators, lawyers, practitioners, and researchers.

This research is being conducted by Northwestern University professors [Daniel W. Linna Jr.](#) and [V.S. Subrahmanian](#), in collaboration with the New York City Bar Association.

Your responses will be kept confidential and used for research purposes only. Results will be published in a format such that individual judges cannot be identified. Your email address will be collected for quality control, and you will be given the option to indicate if you are willing to be contacted by us for follow-up interviews.

We would like your help to contribute to this important research. The survey is designed to take approximately **5 minutes** to complete.

If you have any questions about this research, please contact Daniel W. Linna Jr. at daniel.linna@law.northwestern.edu and V.S. Subrahmanian at vss@northwestern.edu.

By selecting “I agree” you consent to the conditions in the Consent Form and as summarized above.

- I agree
- I disagree

Skip to the "End" section of the survey if answer = "I disagree"

Part 1: Familiarity and Current Use

Q1.1: How frequently do you use the following AI tools in your work as a judge? *¹¹

Scale: Daily, Weekly, Monthly, Rarely, Never

- ChatGPT (OpenAI)
- Claude (Anthropic)
- Copilot (Microsoft)
- Gemini (Google)
- Grok (X.ai)
- Perplexity (Perplexity)
- CoCounsel (Thomson Reuters)
- Westlaw AI-Assisted or Deep Research (Thomson Reuters)
- Protégé or Lexis+ AI (LexisNexis)
- Vincent AI (vLex)
- Harvey
- Legora

Q1.2 If you use other AI tools in your work as a judge, please identify the tools and specify for each if you use it Daily, Weekly, Monthly, Rarely, or Never.

Q1.3 Do you use AI for the following use cases in your work as a judge? Select all that apply. *

- Review documents (e.g., searching, asking questions about documents)
- Conduct legal research
- Draft documents filed in cases (e.g., orders, opinions, judgments, jury instructions)
- Draft documents not filed in cases (e.g., letters, emails, articles)
- Edit documents filed in cases (e.g., orders, opinions, judgments, jury instructions)

¹¹ Questions that require a response are marked with a red asterisk.

- Edit documents not filed in cases (e.g., letters, emails, articles)
- Inform decisions
- Make decisions
- Prepare case timelines or chronologies
- Transcribe text or audio
- Summarize text or audio
- Analyze text or audio
- Prepare questions for a proceeding (e.g., hearing, trial, status conference)
- Managing case files and administrative tasks
- Other (please describe)
- I don't use AI in my work

Display Q1.4 only if Q1.3 answer includes "Other (please describe below)"

Q1.4 If you use AI for other use cases in your work as a judge, please describe the use cases. *

Q1.5: Do other individuals in your chambers use AI for the following use cases? Select all that apply. *

- Review documents (e.g., searching, asking questions about documents)
- Conduct legal research
- Draft documents filed in cases (e.g., orders, opinions, judgments, jury instructions)
- Draft documents not filed in cases (e.g., letters, emails, articles)
- Edit documents filed in cases (e.g., orders, opinions, judgments, jury instructions)
- Edit documents not filed in cases (e.g., letters, emails, articles)
- Inform decisions
- Make decisions
- Prepare case timelines or chronologies
- Transcribe text or audio
- Summarize text or audio
- Analyze text or audio
- Prepare questions for a proceeding (e.g., hearing, trial, status conference)
- Managing case files and administrative tasks

2026]

ARTIFICIAL INTELLIGENCE IN FEDERAL COURTS

67

- Other (please describe)
- They don't use AI in their work

Display Q1.6 only if Q1.5 answer includes "Other (please describe below)"

Q1.6: Please describe the use cases for which other individuals in your chambers use AI. *

Q1.7: Has the court administration provided you training on the use of AI tools? *

- Yes, and I attended the training
- Yes, but I have not attended the training
- No
- I am not sure

Q1.8: What is your policy for how those who work in your chambers may use AI? *

- Permitted and encouraged
- Permitted
- Discouraged but not formally prohibited
- Formally prohibited
- I have no official policy
- No one works in my chambers

Q1.9: Please feel free to elaborate on how you or others in your chambers use AI.

Q1.10: How frequently do you use AI outside of work? *

- Daily
- Weekly
- Monthly
- Rarely

- Never

Display Q1.11 only if Q1.10 answer is not "Never"

Q1.11: What AI tools do you use outside of work and for what use cases? *

Part 2: Outlook on AI's Potential for the Judiciary

Q2.1: Which of the following best describes your general outlook on AI's potential for the judiciary? *

- Very optimistic
- Somewhat optimistic
- Neutral
- Somewhat concerned
- Very concerned

Q2.2: Please feel free to elaborate on your attitude toward AI's potential for the judiciary.

Q2.3: We encourage you to add other comments and information about how AI is being used in the judiciary, including the possible benefits and risks.

Part 3: Contact Information

Q3.1: Please provide your email address. (Email addresses will be used for quality control, including to identify duplicate responses. Your responses will be kept confidential and used for research purposes only. Results will be published in a format such that individual judges cannot be identified.) *

Q3.2: May we contact you via email if we have follow-up questions or wish to interview you? *

- Yes

2026] ARTIFICIAL INTELLIGENCE IN FEDERAL COURTS 69

- No

Q3.3: Would you like to receive a copy of the research findings when they are published? *

- Yes
- No

End

Thank you for your time and for sharing your perspective with us.