

Automated Tax Planning: Who's Liable When AI Gets It Wrong?

by Benjamin Alarie, Rory McCreight, and
Cristina Tucciarone

Reprinted from *Tax Notes Federal*, September 25, 2023, p. 2297

Automated Tax Planning: Who's Liable When AI Gets It Wrong?

by Benjamin Alarie, Rory McCreight, and Cristina Tucciarone



Benjamin Alarie



Rory McCreight



Cristina Tucciarone

Benjamin Alarie is the Osler Chair in Business Law at the University of Toronto and the CEO of Blue J Legal Inc., and Rory McCreight and Cristina Tucciarone are lead analysts at Blue J Legal.

In this article, the authors examine who is accountable for erroneous or flawed AI-generated tax advice and caution that responsibility will ultimately rest with the human tax professionals who are best positioned to ensure that the advice has been properly vetted and verified.

Copyright 2023 Benjamin Alarie, Rory McCreight, and Cristina Tucciarone. All rights reserved.

advice based on the output of an AI, to what extent will I be held professionally responsible? How do I navigate situations in which the AI's tax analysis differs from my own, even if I struggle to document or even explain why I expect a different outcome? As AI becomes more powerful and is responsible for informing a greater number of important decisions, the challenge of assigning and apportioning this liability becomes progressively more difficult.

Inevitably, tax professionals will increasingly turn to AI-driven tools for assistance. Forty percent of legal professionals use or plan to use generative AI,¹ and 50 percent believe that generative AI will transform legal practice.² What do tax practitioners need to know before they dive into using generative AI in their daily work? This installment of Blue J Predicts explores the novel liability considerations that will arise as professional firms implement AI tools for automated assistance with tax planning and analysis. It examines the ways in which AI frustrates and upends traditional legal concepts of liability by complicating who is liable when an AI-informed analysis is flawed. It also explores approaches that regulators and tax practitioners can take to shield themselves against adverse AI consequences as they integrate computational tools into their workflow. Ultimately, it notes that the blame game is perhaps the same as it ever was — the responsibility for competent advice lies with the tax professionals who employ these and other tools.

I. Introduction

Considering the continued proliferation of and rapid advancement in artificial intelligence technology, tax professionals are increasingly finding themselves confronted with novel accountability questions. If I render erroneous tax

¹ Bob Ambrogi, "Survey Released Today Finds That 40 Percent of Legal Professionals Use or Plan to Use Generative AI," LawSites, Aug. 21, 2023.

² Ambrogi, "Nearly Half of Legal Professionals and Consumers Believe Generative AI Will Transform Law Practice, LexisNexis Survey Finds," LawSites, Aug. 22, 2023.

II. Novel Liability Concerns Raised by AI

A. Who Is Liable?

In a recent incident involving Uber, a self-driving car hit and killed a pedestrian in Arizona. Prosecutors determined that Uber was not criminally liable for the individual's death but said charges could be pressed against the vehicle operator, who was allegedly watching TV on her phone at the time of the collision.³ Corporations have taken note of this case and others like it involving AI-powered tools and are keen to offload corporate liability through responsibility-shifting. Society's losing struggle to pin down responsibility for algorithmically derived wrongdoing will continue to fan the flames of automation.⁴

This situation draws parallels to two established challenges in corporate liability known as the "many hands problem" and the "no hands problem."⁵ The many hands problem refers to operations like open-source software development that involve many elements working toward a common goal. The elements do not all necessarily interact with one another directly, and the actions of one bad actor cannot easily be isolated from the group. This is one reason tax firms are diligent in ensuring their employees are credited and responsible for all advice given to clients. Conversely, the no hands problem argues that even when no bad actor is present in a large group of individuals, systemic or process failings can lead to unforeseen, negative outcomes. In the Uber example, whether the accident was caused by the operator's negligence or by a bug in the AI's code, fault cannot be easily ascribed to any single agent. However, the driver was the designated backstop for the AI and was obligated to provide oversight of the operation of the vehicle. Thus, Uber was able to avoid liability.

Similarly, the law will need to adapt as corporate boards integrate AI systems into

processes for significant corporate decisions. This novel approach to corporate governance already has precedent.⁶ A Hong Kong-based venture capital firm, Deep Knowledge Ventures, "has appointed VITAL, a machine learning program capable of making investment recommendations in the life science sector, to its board."⁷ Analogous AI-driven decision mechanisms could be created for tax functions within a corporation. Imagine an AI tool entrusted with an organization's financial recordkeeping and further deployed to use that information to recommend tax positions and advice on potential transactions. Given the rapid growth of AI tools in important decision-making, legislators must consider how these new tools could affect liability when things go awry.

B. New Liability Considerations

In the swiftly evolving realm of AI, determining liability poses a new set of challenges that arise from the inherent characteristics of AI systems, which do not fit neatly into conventional frameworks of legal accountability. A pair of scholars who have studied the specific characteristics of AI that challenge our current conception of liability have raised concerns.

In an article that explores how AI will change liability in the medical malpractice context, Scott J. Schweikart⁸ argues that two key elements of AI make the assessment of liability under the law difficult: AI's "black box" problem and the diffuse development and control of AI. The black box problem is one of verification and explainability.⁹ People who create AI systems are unable to see how the responses are generated because of the complexity of the systems.¹⁰ This can create problems for end-users who rely on AI-generated tax advice. If a tax professional advises a client on tax matters, it is imperative that he explain the

⁶ Martin Petrin, "Corporate Management in the Age of AI," 2019 *Colum. Bus. L. Rev.* 965 (2019).

⁷ "Deep Knowledge Ventures Appoints Intelligent Investment Analysis Software VITAL as Board Member," *GlobeNewswire*, May 13, 2014.

⁸ Senior research associate at the American Medical Association and legal editor of the *American Medical Association Journal of Ethics*.

⁹ Schweikart, "Who Will Be Liable for Medical Malpractice in the Future? How the Use of Artificial Intelligence in Medicine Will Shape Medical Tort Law," 22 *Minn. J.L. Sci. & Tech.* 1 (2021).

¹⁰ Lou Blouin, "AI's Mysterious 'Black Box' Problem, Explained," *University of Michigan-Dearborn News*, Mar. 6, 2023.

³ Urial J. Garcia, "No Criminal Charges for Uber in Tempe Death; Police Asked to Further Investigate Operator," *AZCentral*, Mar. 5, 2019.

⁴ Mihailis E. Diamantis, "The Extended Corporate Mind: When Corporations Use AI to Break the Law," 98 *N.C. L. Rev.* 893 (2020).

⁵ Diamantis, "Employed Algorithms: A Labor Model of Corporate Liability for AI," 72 *Duke L.J.* 797 (2023).

client's obligations and positions. When relying on black box advice, accurate or not, proceeding without understanding the rationale behind a position leaves individuals unaware of their risks.

The second element in Schweikart's analysis delves into how the nature of software development can lead to liability concerns.¹¹ Open-source software development, and open-source AI development in particular, rely on many different individuals, using a variety of different tools, in different locations, to create new applications. This development style is both diffuse and discrete. Developers do not tend to work closely together on projects, and all the components can be assembled asynchronously by various individuals who may or may not consider themselves part of the same team. These factors mean that no one person controls the technology. Principles of liability, however, rest on the presumption of assigning responsibility to an individual entity. Open-source software development's diffuse control, combined with the black box nature of AI, makes it difficult to identify a responsible party to hold accountable.

The diffuse development and control of AI open-source software systems can exacerbate issues arising from poor-quality data. AI systems are trained on large datasets to provide answers to users based on learning from these datasets. The quality of the data is hugely important for the quality of the algorithm. Poor-quality data leads to poor-quality outputs. Like the black box problem, this is also a problem of verification because the end-user generally cannot assess the full dataset. Consequently, an AI system trained for tax professionals should ensure that its training data — be it legal precedent, government documents, or reliable secondary sources — is maintained by removing outdated or revoked legal precedent and ensure that the model is frequently updated with new law and legal documents.¹² Also, when working with specialized datasets, such as tax guidance, the model should be able to identify and retrieve the

documents that it relies on for its answer. Although AI developers should strive to ensure that their systems are returning quality advice with an explanation, the tax practitioner remains professionally responsible for any mistakes resulting from reliance on an AI system trained on poor-quality tax data. For this reason, practitioners must carefully verify the tax analysis received from any AI-driven tool.

Andrew Selbst,¹³ in an article analyzing the tort of negligence in AI, lists four additional unique features of AI that can complicate accountability.¹⁴ The first is the inability to predict and account for errors. This is directly related to the black box problem of AI. Without truly understanding how the AI rationalizes, we can never fully know how it will act in a given circumstance. Thus, risk is inherent when relying on advice from AI systems. Second, Selbst claims that there are physical and cognitive limitations when humans and AI interact. The complexity of AI systems can make direct use by individual tax practitioners challenging. These interactions may require mediation because of the intricate nature of AI models and the need for a certain level of computer programming understanding to effectively use it. Next, Selbst discusses how software vulnerabilities in AI systems can lead to new challenges in areas that were not previously prone to software weaknesses. AI systems trained to understand and interact with real-world concepts move beyond traditional computer engineering, which involves working with defined software concepts. Interacting with real-world objects can lead AI to be deceived by a human (adversarial or not) who is changing the world around it.¹⁵ And finally, Selbst expands on the concern about relying on data and statistical methods for decision-making. Any bias in the data will be reflected in the AI's decisions and will be amplified over time as new data is generated from previous AI decisions.

¹¹ Schweikart, *supra* note 9.

¹² Blue J's generative AI tool, Ask Blue J, is regularly updated to incorporate the latest IRC and regulation updates and the most recent IRS tax guidance. Outdated or revoked legal guidance is promptly removed from its dataset.

¹³ Assistant professor, UCLA School of Law.

¹⁴ Selbst, "Negligence and AI's Human Users," 100 *B.U. L. Rev.* 1315 (2020).

¹⁵ Ariel Bogle, "Hackers Tricked a Tesla, and It's a Sign of Things to Come in the Race to Fool Artificial Intelligence," ABC News, Apr. 13, 2019.

While his ideas sometimes overlap with Schweikart's, Selbst introduces important contextual considerations. He explores how the effect of AI liability must be analyzed based on human interactions and interactions with the world, highlighting emerging challenges as AI systems proliferate. Malicious actors may attempt to disrupt AI, posing significant risks and potential liability for designers or users. Selbst underscores the need for AI systems that have accessible interfaces and defenses against deception.

C. The Inadequacy of Legal Frameworks

With the explosion of AI technology in the last year, governments find themselves struggling to enact timely AI-related legislation. Moreover, new AI technologies are emerging with such regularity that existing legal institutions must work hard to keep pace. Courts are ill-equipped to create effective precedent under these rapidly changing conditions. By the time a judge delivers her final ruling regarding damages caused by certain AI technology, interceding evolutions in the technology may well have rendered the decision obsolete for future cases. The demand for legal guidance from the judiciary will continue to predictably outstrip its supply.

Liability law considers not only who was responsible for a wrong but also the intention behind the individual's decisions.¹⁶ Yet AI does not manifest legal intent in the same way a human does. An AI program operates on a defined goal that it is programmed to pursue. For example, an AI tax product could be programmed to minimize tax liability to the extent permitted by the law, subject to several other constraints (for example, transaction costs). If the AI designer's stated goal in this instance is accurately reflected in the model, then the intent of the AI's action cannot be in doubt. If this hypothetical AI system were to suggest a potentially illegal course of action, it would likely reflect a bug in the code or poor training data regarding the legality of the transaction. Because AI cannot be responsible for defining its own goals (at least not yet), one must

look to the AI creator, the individual responsible for the bug in the code, or the user for liability purposes. Put differently, the evaluation of intent undertaken in liability law is illogical when determining the extent to which an AI can be held accountable for its mistakes.

The complexity and opacity of the underlying data present a distinct set of challenges. AI systems developed to perfectly minimize tax liability may be inscrutable to even the most seasoned tax practitioners. "To some extent, the point of artificial intelligence is to develop new approaches in a way that is more effective than human intelligence can manage," Selbst wrote. "Yet that can lead to circumstances beyond what human intelligence can anticipate."¹⁷ These models are designed to find the best way to achieve a stated goal and may derive wholly new reasons for a particular solution that can be unintuitive to a human reviewer. Working within tax rules, the AI system could derive new tax avoidance schemes that follow the letter of the law but may violate the spirit of the law or run afoul of antiavoidance rules. This inscrutability compounds the black box problem.

Even if a user has access to an AI's rationale, he may be unable to understand it. This nuance will further influence a court's facts and circumstances analysis for the scope of a taxpayer's liability. As an example, how would the courts handle penalties for taxpayers that relied on an AI system for their tax planning? Should taxpayers be liable for section 6662 accuracy-related penalties when led astray by reliance on a tax-trained AI? Courts will likely be too slow in adjudicating liability in these emerging tax schemes, by which time AI systems will have become commonplace, resulting in many more AI-informed tax filings influenced by this legal rationale. In this scenario, it might be considered reasonable for taxpayers to rely on tax guidance from a hypothetical tool that is widely relied on, potentially shielding them from additional penalties. Legislators will need to take proactive measures to deter irresponsible reliance on AI systems by tax advisers and taxpayers alike.

¹⁶ Robin Feldman and Kara Stein, "AI Governance in the Financial Industry," 27 *Stan. J.L. Bus. & Fin.* 94 (2022).

¹⁷ Selbst, *supra* note 14.

III. Proposed Regulatory Response

Recently, a coalition of AI experts and industry executives advocated for a six-month moratorium on AI development, citing potential societal risks arising from the technology's rapid advancement. The group recommended that this hiatus be used to resolve regulatory issues, ensuring that AI technologies have a positive effect and that their associated risks are manageable.¹⁸ Given the concerns raised by industry leaders and the prevalence of publicly accessible AI tools like ChatGPT, efforts to establish regulations are underway. Today, nearly every state, along with the District of Columbia, has legislation pending that addresses some facet of AI regulation.¹⁹ That regulation is increasingly viewed as essential for promoting responsible AI development and usage while deterring malicious applications and reducing the likelihood of harmful errors.

A. Liability Allocation and Accountability

The integration of AI-generated tax advice necessitates a clear framework for allocating liability in cases of errors, misinformation, or financial losses. Regulatory bodies can offer vital guidance on how responsibility should be divided between tax professionals and AI tool developers.

For effective liability allocation, clearly defined roles and responsibilities for each party are essential. Tax professionals are responsible for supervising AI-generated advice, ensuring its accuracy in the context of intricate tax laws and client-specific circumstances, and validating recommendations before presenting them to clients. They bear the responsibility for any errors or omissions in the advice given. AI developers, on the other hand, should be accountable for the functionality and reliability of their tools and ensuring that they meet certain standards and ethical guidelines. Clients also have a role: They must provide accurate and complete information to the AI system, exercise due diligence in

following recommendations, and communicate any discrepancies to the tax professionals. Existing laws, such as those imposing accuracy-related penalties, already mandate that clients complete documentation for advice they rely on.²⁰

Contracts between tax professionals and AI developers should clearly outline the scope of AI-generated advice, the extent of its accuracy guarantees, and the limitations of liability for each party. These agreements set the legal framework for addressing any errors or adverse outcomes. Contingency plans should also be in place to outline corrective steps, including client communication and potential remedies. Both parties may also consider obtaining insurance coverage specifically for liabilities arising from AI-generated advice.

Current legislative efforts are exploring various models of liability. For example, Maryland House Bill 996 proposes holding AI developers strictly liable for damages "if the software is used to cause personal injury or death."²¹ Another approach suggests giving AI legal personhood, enabling it to bear legal responsibility.²² This would mean that the AI itself could be held liable in instances of malpractice. However, this approach may be impractical. Unlike human entities, AIs do not possess the emotional or conscious capacity to comprehend liability, and unlike corporations, they lack the legal capability to own assets. Implementing that change would require a raft of significant cascading legal adjustments. The practical implications of holding an algorithm accountable remain unclear.

It seems most probable that future regulations will continue to hold tax professionals responsible, given that human expertise and oversight are essential for ensuring accurate and relevant tax advice. Traditional negligence-based approaches have sometimes allowed professionals to disclaim duty of care by citing reliance on advanced technology. However, the

¹⁸"Pause Giant AI Experiments: An Open Letter," Future of Life Institute (Mar. 22, 2023).

¹⁹National Conference of State Legislatures, "Artificial Intelligence 2023 Legislation" (Apr. 18, 2023).

²⁰Benjamin Alarie, Cristina Tucciarone, and Christopher Yan, "Overcoming Accuracy-Related Penalties With Reasonable Cause," *Tax Notes Federal*, Mar. 27, 2023, p. 2145.

²¹Md. H.B. 996 (Feb. 10, 2023); see also Minn. H.F. 2890 (Mar. 15, 2023).

²²Jason Chung and Amanda Zink, "Hey Watson — Can I Sue You for Malpractice? Examining the Liability of Artificial Intelligence in Medicine," 11 *Asia Pac. J. Health L. & Ethics* 51, 57 (2018).

core principle of negligence, which centers on accountability, argues against that blame-shifting. It suggests that a negligence approach should be employed to attribute liability only when a person or entity could have prevented harm through increased diligence.²³ Some advocate for a strict liability regime, which would hold tax professionals fully accountable for damages stemming from their activities, thereby encouraging responsible behavior.

Without clear accountability, tax firms may mitigate their liabilities by opting to risk algorithmic wrongdoing rather than face potential liability caused by their tax professionals' errors. The existing legal landscape raises concerns, granting corporations immunity as they reduce human involvement, as seen in cases like Uber's self-driving car incident.²⁴ The legal concept of *respondeat superior*, which holds an employer responsible for employees' actions, may inadvertently encourage human involvement to reduce liability.²⁵ By applying the concept of the extended mind theory to tax law, firms can establish accountability for algorithmic errors, attributing them to the firm itself.²⁶

The labor model of liability holds firms accountable for the actions of their algorithms.²⁷ Under this model, tax professionals would be liable for the outcomes they substantially control and benefit from.²⁸ Thus, the bulk of the responsibility for AI-generated work is likely to rest with the tax professional.

B. Disclosure and Informed Consent

Tax professionals have a duty to provide clear information about the tax guidance they offer. Regulatory bodies might consider mandating disclosures about the use of AI tools in generating that advice. By doing so, clients can better understand the role of AI in the advice they receive and decide whether to follow it based on

²³ *Id.*

²⁴ Garcia, *supra* note 3. Lacking a liability theory, prosecutors declined to press charges against Uber when one of its self-driving cars was involved in a fatal pedestrian accident in Arizona.

²⁵ Diamantis, *supra* note 4.

²⁶ *Id.*

²⁷ Diamantis, *supra* note 5.

²⁸ *Id.*

their comfort level with AI-generated insights. That disclosure can prevent potential misunderstandings about the origins of the advice.

Highlighting the use of AI brings attention to potential risks, such as errors from algorithmic limitations or data-quality concerns. With this knowledge, clients can exercise due diligence, possibly seeking further validation or requesting additional consultations. Ethically, disclosure aligns with principles of honesty and accountability, emphasizing tax professionals' duty to communicate any tools, methods, or technological interventions used.

For effective disclosure, tax professionals should do the following:

- clearly communicate AI involvement at the outset;
- update clients on any changes to the AI tools or their capabilities;
- maintain comprehensive records of AI use disclosures;
- explain AI's role in understandable terms;
- provide resources detailing AI's role in tax advice; and
- obtain explicit client consent for AI tool use when necessary.

Legislation in some states is already moving toward that disclosure.²⁹ These bills would require clients to be informed if their personal data is used by AI models in tax advice. Proper disclosure fosters trust and ensures ethical AI use in tax planning.

C. Data Privacy and Security

The importance of data privacy and security in AI-generated tax advice cannot be overstated. The European Union's proposed AI Act, with its risk-based approach, is a notable attempt to address these concerns.³⁰ For instance, it bans systems with unacceptable risks, such as government-run social scoring.³¹ For other applications, consumers must be informed about

²⁹ Mass. H. 1873 (Feb. 16, 2023); N.J. S. 3714 (Mar. 13, 2023); N.Y. A. 3308, A. 3593, S. 2277 (Feb. 2, 2023).

³⁰ Future of Life Institute, The Artificial Intelligence Act (last accessed Sept. 1, 2023).

³¹ *Id.*

AI-generated content to protect against copyright breaches and illicit content.

While it's promising, the EU legislation will take time to enforce and may need updates as AI evolves. In the interim, existing privacy laws like the EU's General Data Protection Regulation can address some AI-related concerns. For instance, after a temporary restriction, ChatGPT was reinstated in Italy only after complying with the General Data Protection Regulation requirements, which included installing age verification systems and permitting users to opt out of having their personal data processed.³²

Existing data protection laws such as the General Data Protection Regulation and the California Consumer Privacy Act provide a framework for responsible AI data handling. Their features include:

- Data minimization: collecting only necessary data for AI analysis.³³
- Anonymization and encryption: removing personally identifiable information from datasets and securing data during storage and transmission.
- Access controls: limiting data access to essential personnel.
- Vendor due diligence: ensuring third-party AI tools adhere to robust privacy practices.
- Data retention policies: defining data storage durations and ensuring mechanisms for data access and deletion requests.³⁴

By adhering to these privacy and security features, tax professionals can ensure the security of client data, building trust and avoiding potential legal issues.

³² Shiona McCallum, "ChatGPT Accessible Again in Italy," BBC News, Apr. 28, 2023.

³³ Bert-Jaap Koops, "The Trouble With European Data Protection Law," 4(4) *Int'l Data Privacy L.* 250-261, 256 (2014). Separately, Blue J's generative AI solution, Ask Blue J, answers challenging tax law questions using a natural language interface. Additional AI-powered tools for case research and analysis use the selection of values for factors considered relevant to case decisions. This does not require personal identifiable information input. Personal identifiable information collection for our users is limited to the data necessary to manage authentication and authorization for the purposes of using the Blue J platform (email address and name). Further information regarding Blue J's information security program can be found at Blue J Legal, "Information Security Program at Blue J" (2022).

³⁴ As an example, customer data is retained and protected by Blue J indefinitely unless a formal request for removal is received.

D. Algorithmic Accountability

As AI becomes more prevalent in tax services, regulatory bodies should consider mandating that AI systems be designed to elucidate their decision-making processes. This would counteract the opaque nature of many AI algorithms, ensuring that outcomes are understandable, reviewable, and correctable.³⁵ Transparent AI provides insights into its creation, training data, considered features, and decision criteria. For instance, the most influential variables in an AI-generated recommendation can help tax professionals comprehend the underlying factors. Consequently, explainable AI enables tax professionals to articulate the reasoning behind recommendations to clients, aiding clients in making tax decisions.

Understanding an AI system's reasoning enables tax professionals to detect errors before advising clients. They can trace the data to identify areas for improvement, adapt models to changing needs, and address practical challenges.

Maintaining thorough documentation is vital for algorithmic accountability and explainability. Documentation should include data sources, model guidelines, training process specifics, validation techniques, change logs, and guidance for interpreting AI-generated advice. That documentation ensures quality compliance and risk mitigation. Especially in regulated sectors such as taxation, comprehensive documentation is pivotal for demonstrating compliance through transparent decision-making processes.

Canada's proposed Artificial Intelligence and Data Act exemplifies a move toward algorithmic accountability and explainability.³⁶ The act says that entities responsible for high-impact systems must publicly describe the AI system, its intended use, generated content, and decisions, along with implemented mitigation measures.³⁷ At Blue J, we prioritize explainability by revealing the source material for Ask Blue J responses and offering

³⁵ Anat Lior, "AI Strict Liability vis-à-vis AI Monopolization," 22 *Colum. Sci. & Tech. L. Rev.* 90 (2020).

³⁶ Government of Canada, Artificial Intelligence and Data Act (last updated Aug. 2, 2023).

³⁷ Roland Hung, "Regulating Generative Artificial Intelligence: Balancing Innovation and Risks," *Torkin Manes Barristers & Solicitors* (June 20, 2023).

detailed explanations for our tax predictions.³⁸ In essence, algorithmic accountability and explainability are vital for understanding AI decision-making, especially in tax planning, for which clear reasoning and error tracing are essential.

E. Standardized Testing and Validation

As AI tools become integral to tax advice, a pivotal regulatory response could be the introduction of standardized testing and validation protocols. These regulations would mandate thorough testing of AI systems before their deployment in tax advisory roles. The goal is to ensure the accuracy, reliability, and performance of these AI systems and minimize the risk of inaccuracies that have significant financial or legal consequences.

AI tools would undergo mandatory assessments for accuracy and reliability, tested against varied tax scenarios and datasets to evaluate their consistency and precision. Standardized performance evaluations would allow for comparisons against benchmarks and industry standards, ensuring that AI responses align with both human expertise and established guidelines. Key validation metrics such as accuracy, precision, recall, and other measurable criteria, would be standardized, ensuring a uniform approach to performance assessment, especially as AI systems evolve. Emphasis would be placed on proactive error, bias, and inconsistency detection, mitigating potential issues before they reach clients.

To further ensure AI system reliability, compliance, and ethical integrity, independent audits and certification standards should be introduced. Third-party audits, conducted by external experts, would evaluate AI systems against best practices and regulatory standards, ensuring ethical, compliant, and error-free advice. The audits promote transparency through comprehensive documentation and bolster accountability by subjecting AI to external review.

³⁸ See, e.g., Alarie, Kim Condon, and Nasreen Rahman, "Unbridled Losses: Harnessing Machine Learning for Tax Analysis," *Tax Notes Federal*, Apr. 24, 2023, p. 637. Alarie et al., "The Rise of Generative AI in Tax Research," *Tax Notes Federal*, May 29, 2023, p. 1509. Alarie and Rory McCreight, "The Ethics of Generative AI in Tax Practice," *Tax Notes Federal*, July 31, 2023, p. 785.

Tax professionals can leverage the audits to assure clients and regulators of their AI system's rigorous evaluation and adherence to predefined standards. In parallel, certification standards would act as quality benchmarks, reflecting industry best practices and regulations. Certification assures tax clients of the reliability of AI-generated advice while holding vendors accountable for their products' performance and security.³⁹ Continuous reevaluation would ensure that certified AI systems stay updated with changing tax regulations.

Legislators are already contemplating mandatory audits and certifications. For example, Connecticut has directed its Department of Administrative Services to inventory and assess AI systems used by state agencies for potential discriminatory effects starting February 1, 2024.⁴⁰ However, a pressing question of who will conduct these audits and certifications remains. While government agencies may seem like ideal candidates, they lag in establishing AI safety standards, given the rapid technological advancements and the slower pace of legislative development.⁴¹ In the interim, business leaders and academics are encouraged to establish nongovernmental regulatory bodies and certification processes, delineating reliable AI applications.⁴²

In summary, regulations enforcing standardized testing, independent audits, and certification standards for AI tax services would significantly enhance the reliability, compliance, and ethical accountability of these systems. Such measures would foster trust in AI-generated tax advice, promote accountability, and ensure that AI systems remain robust, reliable, and compliant with industry standards. Moreover, the iterative nature of these initiatives would enable the ongoing improvement of AI system performance and alignment with evolving tax regulations.

³⁹ At Blue J, we work with an independent auditor to maintain a Service Organization Controls 2 report that objectively certifies our controls to ensure the continuous security of our customers' data. For more information, see Blue J Legal, *supra* note 33.

⁴⁰ National Conference of State Legislatures, *supra* note 19.

⁴¹ Cat Zakrzewski and Nitasha Tiku, "AI Companies Form New Safety Body, While Congress Plays Catch-Up," *The Washington Post*, July 26, 2023.

⁴² Blair Levin and Larry Downes, "Who Is Going to Regulate AI?" *Harvard Bus. Rev.* (May 19, 2023).

IV. Conclusion

AI represents a paradigm shift in many ways. It is positioned to change the way we work and live, empowering individuals to push beyond what was previously thought possible. This shift will carry important implications for the work processes of tax practitioners.

The shift does not come without concerns, and this article has explored the specific concerns AI creates for assignment of liability. Ultimately, addressing these concerns as they relate to AI-generated tax advice requires a multifaceted approach that involves collaboration among regulatory bodies, AI developers, and tax professionals. An ideal approach will strike a balance between leveraging the benefits of AI while safeguarding clients and ensuring adherence to ethical and legal standards.

As a community of tax professionals, it is important for us to understand the effect that AI may have on the profession at large and contribute to a conversation about how to address the challenges it presents while embracing the advantages it can bring. While there are risks inherent in harnessing AI, we believe the advantages outweigh those risks — so long as we are diligent and thoughtful about how to properly oversee the use of AI in the practice of tax law. ■

taxnotes®

Federal | State | International



Read what the leaders read.

Our subscribers include decision-makers, policy advisers and practitioners from the Am Law Top 100 law firms; U.S. and international governing agencies like Treasury, Congress, the IRS Office of Chief Counsel, state finance departments, and the OECD; influential NGOs; the Big Four accounting firms; and the leading law schools.

taxnotes.com

Written by experts,
read by decision-makers.