

Meeting the AI Moment in Asset Management:

An Agenda for Industry Lawyers

As lawyers serving asset managers, many of us have thought to ourselves at some point, “If I wanted to be a spreadsheet expert, I would have gone to business school instead of law school.” Such reflection about our occasional struggles for technical mastery may be even more prevalent now at the dawn of the AI age. Grasping AI can feel like the spreadsheet problem on steroids. But like it or not, a key part of the job has long been to enable the deployment of investment innovations in this highly regulated industry so that investment professionals can take full advantage of new techniques to enhance returns, while clients and regulators can have confidence that the upgraded investment process remains grounded in transparency, risk management and fiduciary principles. The art is in striking a balance between valid goals that can be in tension—just like a car needs both gas and brake pedals to get where it’s going safely and on time—such that the balance struck can be defended as reasonable when inevitable bad investment outcomes occur due to market forces (like car accidents in bad weather).

There are certainly high-profile examples from recent decades of legal liability attaching where investment innovations were arguably deployed without the right balance of gas and brakes. The practice of allowing market timing in registered funds in the early 2000s and widespread investment in mortgage-backed securities in the mid-2000s raised many questions for industry lawyers when the dust settled—including whether they held a deep enough real-time understanding of these practices and products to allow them to assess and manage the relevant legal risks they posed. Of course, there have also been many counterexamples of the industry adopting and integrating innovative products and processes (e.g., ETFs, quant models), where the new capabilities and their attendant risks have been successfully distilled and explained to clients and regulators, allowing the managers to rebuff legal challenges that were predictably raised when market circumstances led to periods of underperformance.

In this context, how should we think about and prepare for the emerging use of AI in asset management? At one level, AI just represents the next generation of investment innovation, and (as in the past) many of us likely need to do some catching up to the investment and operations professionals in our understanding of AI’s capabilities and limitations. A basic understanding of the technology is table stakes for meaningful participation in the coming conversations about how AI can and should be incorporated into the investment process; how to describe to clients and regulators AI’s features and risks; and how to devise usage and marketing guardrails to mitigate those risks.

At another level, however, AI is fundamentally different from earlier technical advances and poses entirely new legal challenges, including for industry lawyers. Significantly, AI tools can execute tasks at dazzling super-human levels, using autonomous machine learning processes that humans currently *cannot fully understand and explain* in many cases. These tools can likewise fail rather spectacularly, also for reasons that are not always *understood or even detectable* by humans. While this black-box condition may be an uncomfortable spot for an engineer, it is a legally precarious position for a fiduciary. It is an overstatement of the technology (at least for now) to think of AI tools as truly autonomous agents, akin to external subadvisers entrusted with independent decision making. But as a legal matter, the difference between an agent and a very complex tool may be a distinction without a difference: a fiduciary asset manager’s use of a non-transparent and unexplainable tool in making investment decisions arguably poses no less legal risk of a fiduciary duty breach than employing a human agent whom the fiduciary is unable to monitor and supervise.

In addressing these and other unique challenges posed by AI, we propose below a three-part agenda for asset management lawyers to frame their efforts to establish a workable and flexible balance in the incorporation of AI tools:

1	Understand and be able to explain AI capabilities and mechanics to effectively advocate for the adoption of AI in the investment process.
2	Understand and be able to explain AI limitations and vulnerabilities to describe investment risks and reduce liability exposure.
3	Utilize and leverage AI tools for legal work to maximize efficiency and consistency, bearing in mind the limitations and vulnerabilities of these tools in a legal setting.

Given the speed with which AI usage is expanding and evolving, each of these agenda items can feel like a moving target. In the discussion below, we introduce basic AI technical principles as a frame of reference, which we hope industry lawyers can use as a platform for organized exploration and self-education in the aspects most relevant to their company or firm. We also summarize the types of legal claims the industry might anticipate facing as the private plaintiffs' bar and regulatory enforcement authorities focus their sights on the increased use of AI in the investment process—and in particular on whether fund and adviser disclosures are keeping pace with evolving practices.

AI Capabilities and Mechanics for Asset Management Use Cases

Asset management lawyers will increasingly be called upon to serve as advocates for the adoption of AI in the investment process and related manager functions (*e.g.*, shareholder communications). This advocacy takes several forms, directed at various potentially AI-wary audiences like the SEC (who will review and comment on registration statements for registered funds utilizing AI processes); institutional investors and broker-dealer intermediaries (who will press for detailed descriptions of an adviser's AI offerings in RFI and DDQ responses); and insurance carriers (who will want to understand liability exposure for AI-based failures).

As these audiences become more sophisticated regarding AI, we expect they will require descriptions of not only what the AI tools claim to do, but also *how they do it*—that is, a description of the underlying technology that is used to deliver the proffered results. It will therefore be incumbent on lawyers to understand both the capabilities and basic technical mechanics of the AI systems being adopted.

Crucially, AI is not a monolith. It is a label that embodies several different technologies through which computers are used to analyze data and emulate human thought and learning to solve problems in real-world environments. Different problems involve different data inputs (*e.g.*, numeric, language, visual images), and thus require very different technologies. We are therefore approaching the point where a generic reference to “using AI” to solve a problem will sound as quaint (and uninformative) as an older relative suggesting that you “ask the computer” to answer a question or make a dinner reservation.

In the asset management space, this means that different AI use cases will rely upon different underlying technologies, and lawyers advocating for those use cases will need to understand the capabilities and technical mechanics of the AI models well enough to describe them in relatively plain terms. Following are the current core AI technologies used in the types of tools that might be adopted by asset managers, and some illustrative examples of current and potential use cases:

- **Machine learning (“ML”)**—a type of AI that allows systems to learn from data and improve their performance over time without such revisions being explicitly programmed. The datasets used must typically be quite structured and

consistently coded so that a rules-based set of algorithms can be applied to accomplish relatively narrow and specific tasks.

- Asset management use cases of ML have included portfolio optimization management through real-time risk assessment and automated trading. ML algorithms identify market trends and anomalies and even facilitate personalized investment advice by tailoring recommendations to individual client profiles. Ultimately, ML can facilitate the evaluation and selection of stocks.
- **Deep learning (“DL”)**—a specialized subset of ML using multi-layered artificial neural networks (patterned on the human brain) to process patterns in complex and often unstructured datasets, to enable sophisticated tasks like image recognition, robotics, medical diagnosis, and predictive modeling. This requires much greater computational capacity than standard ML.
 - Asset management use cases of DL include analyzing large and diverse datasets, such as news articles, financial reports, earnings call transcripts, and social media to reveal more subtle relationships and predict future market movements and trends with greater accuracy.
- **Natural language processing (“NLP”)**—technology that allows AI systems to understand, interpret, and generate human language, forming the basis of tools like chatbots, translation services, virtual assistants, search engine results, and spam filters. This technology is expected to grow at the fastest rate in the AI asset management market in the coming years.
 - Asset management use cases of NLP include analyzing textual data (largely from articles) by pulling repetitive keywords and phrases from the data that indicate specific emotions or sentiments that are being reported on a widespread basis.
- **Large language models (“LLM”)**—a specific type of advanced NLP model, built on DL and trained on massive datasets, which excels at understanding and generating more human-like text. It requires much greater computational capacity than NLP, and it can be used for more complex tasks like content creation, writing code, engaging in free-form conversations and summarizing long documents. ChatGPT is an LLM.
 - Asset management use cases of LLM include a manager’s announced implementation of a proprietary ChatGPT model allowing its employees to perform the function of research analysts.

AI Limitations and Vulnerabilities in Asset Management Use Cases

In addition to advocating for AI adoptions, asset management lawyers will also be the first line of defense against legal liability risks arising from AI’s inherent limitations and vulnerabilities. This means understanding the potential shortcomings of AI tools well enough to assess whether and how the risks can be mitigated and in turn to explain the relevant risks in understandable terms keyed to the specific use case.

The Holy Grail for AI developers is human level “artificial general intelligence” (“AGI”), allowing an AI system to readily apply what it learns in one setting to different settings beyond its coding—like a human brain. But contrary to the imaginations of science fiction writers and filmmakers, AI science is many years away from achieving AGI. AI systems can emulate human thought and “learn” in many settings, but these settings (even those using deep learning) are usually very precisely defined and delimited by human-developed parameters.

To illustrate, AI-driven computers can now beat human grandmasters at complex spatial logic games like chess and Go, and even at stylized language-based competitions like TV’s *Jeopardy*. But these games occur in confined, rule-bounded space. By contrast, no computer today could defeat even a child in a game of charades, which requires a mix of human brain capacities—like recognizing the implied meaning of a physical action in space and analogizing to a different context, with an overlay of cultural knowledge woven in—that current computers cannot remotely approach. Humans experience the physical world through our senses and develop basic common sense, which allows us to easily apply what we learn in one setting to many others, almost from birth. It is natural for us to generalize from one experience to many others and to analogize across contexts, but we don’t yet understand exactly how our brains do that—much less how to code computer algorithms to mimic it.

This is a dramatic gap between human and current artificial intelligence, and achieving AGI is still considered to be far in the future. The inability of AI to generalize what it learns in one context and apply it in other settings—and the related lack of a human level “common sense” check on the outputs of AI models—informs many of the vulnerabilities of AI and how they can be expected to manifest in specific uses cases.

The following are key areas of vulnerability for AI models, which we anticipate as points of focus for regulators and the private plaintiffs' bar as they formulate potential legal theories critical of AI usage in asset management, especially in the event of underperformance of an AI-driven investment strategy:

LACK OF TRANSPARENCY AND EXPLAINABILITY

As AI systems have become more complex, and especially with the advent of deep neural networks supporting DL and LLM, human developers of AI are increasingly unable to understand exactly *how* these systems learn and then make “decisions” (*i.e.*, arrive at model outputs). While the growth of this autonomous learning capacity is technically very impressive, it is legally quite concerning—especially for a fiduciary money manager. As an asset manager, if I employ or contract with a team of human portfolio managers and research analysts to carry out an investment strategy, I expect that team to be able to explain (in a way the clients and I can understand) how they made their investment decisions based on the available data. This allows me to monitor and ensure that they are following the established strategy in a professional and prudent way (*i.e.*, fulfilling the duty of care) and acting with the sole purpose of furthering the clients' interests (*i.e.*, fulfilling the duty of loyalty). But if instead I deploy an AI model to carry out parts of the strategy, and I can't understand—and the model can't explain to me—how it made its investment recommendations based on the data inputs, it may be very challenging for me to have confidence I'm carrying out my fiduciary obligations of care and loyalty (or prove it if challenged by clients or regulators). While I may know that my AI developers coded the model to optimize for factors I believe to be in the clients' best interests consistent with the agreed strategy, ultimately I may not be able to confirm either *whether or how* the model did so when deployed. In other words, the AI models have a large “black box” aspect, because they are no longer simply algorithms specifying a sequence of decision-making rules to be rigidly applied to data inputs (like a traditional quant model, with its complicated but ultimately transparent coding). The remarkable power of autonomous AI learning brings with it reduced control over the outputs and how they are achieved—control that fiduciaries don't lightly give away when using human agents to carry out tasks.

What's more, if I don't understand and can't explain how AI models learn and make decisions, it is challenging not only to predict when and how they might make mistakes, but also even to detect whether those mistakes have already occurred. Compounding this problem is AI's inherent lack of human level “common sense” as a check on model outputs—without AGI, AI frequently cannot comprehend why a result is just nonsensical, even though it might be obvious to a human brain. Asset management lawyers will need to assess carefully which of these vulnerabilities apply to the particular technology being deployed by their company and how best to explain the attendant risks in fund and adviser disclosures.

TRAINING DATA BIAS

While some AI tools are still trained on confined data sets compiled and curated by coders, increasingly sophisticated AI models (such as LLM) are trained on massive data sets harvested from “real-world” sources such as news reports, internet search results, mobile phone user data, and social media content. In the asset management space, these sources might also include earnings reports and public company regulatory filings. Such real-world data sets can naturally reflect the societal biases of the millions of people and thousands of companies who generated the data, meaning that AI model outputs trained on that data can enshrine and amplify those societal biases in their outputs.

TAIL EVENTS

This issue is a specific instance of the training data bias problem. Real world training data sets understandably tend to under-represent rare but catastrophic scenarios (since they occur so seldom), which can result in models trained on those data being under-prepared to deal well with those scenarios. A good current example is self-driving cars, which are trained on huge data sets of real-world driving experiences that don't include very many instances of unusual but especially dangerous situations, such as construction zones and extreme weather events—situations where lane lines and traffic signs that the driving algorithms depend upon may be obscured. Human driver brains can use common sense learned from other real-world contexts they have experienced to avoid hazards in these new situations, while AI models struggle with this. Of course, tail event risk is not new to asset managers, as traditional spreadsheet-driven models have long faced the same challenge—how to generate outperformance (alpha) in a strategy while adequately protecting against catastrophic loss in a rare but dramatic market decline (a “black swan” event). Although DL models may increase the sophistication with which historical market data can be analyzed for portfolio optimization, it is not clear that these improvements in analysis can solve for the basic input problem—that the training data set includes only a small number of very idiosyncratic black swan events to analyze. Potentially compounding this risk is the black-box issue discussed above. A spreadsheet-based modeler can usually “run” targeted scenarios to see (and report) how the model would be expected to perform assuming black swan-magnitude market moves, whereas an AI model coder may not be able to “run” a similar targeted test. In short, the former is math; the latter is not.

EXPOSURE TO ATTACK

AI researchers have shown that AI models can be fooled—and their results distorted—by intentional manipulation of the data inputs by external bad actors. Here again, the bad actor risk is not new to AI tools, as traditional models are also subject to intentional manipulation by those with sufficient access (*e.g.*, the ability to manually alter inputs in a spreadsheet). However, because of the black-box problem, AI model distortions can be

built on extremely subtle manipulations that are very difficult to detect and prevent by human coders, and because of the common sense gap, the AI models themselves cannot readily recognize when their own outputs are nonsensical.

Based on these limitations and vulnerabilities of AI as applied to asset management use cases, what are the types of legal claims we might expect to see plaintiffs' lawyers and regulators assert in response to the expanding use of AI?

- **Traditional securities law claims:** Shareholder suits and enforcement actions are most likely to assert violations of the securities law provisions governing misleading statements and omissions in registered fund disclosures. Such suits are typically brought by shareholders under Sections 11 and 12(a)(2) of the Securities Act of 1933 (the "Securities Act") and Section 10(b) of the Securities Exchange Act of 1934 (the "Exchange Act") and Rule 10b-5 thereunder. The analogous sections enabling SEC enforcement actions are Section 17(a) of the Securities Act and Section 34(b) of the Investment Company Act of 1940 (the "ICA"). Such SEC enforcement actions also typically include claims under the broad anti-fraud provision in Section 206 of the Investment Advisers Act of 1940 (the "Advisers Act"), and related claims under Section 206 for failure to implement policies reasonably designed to prevent violations of the securities laws.
 - All claims referenced above under the Securities Act, Exchange Act or ICA require a showing that the alleged misstatements or omissions were *material* (i.e., that the statements would have been seen by a reasonable investor as important in deciding to buy or sell the security given the total mix of information available about the security). Claims brought by private plaintiffs must also demonstrate loss causation, which means that the plaintiff must prove that there was a drop in the fund's share price and that it was caused by the alleged misstatement or omission in the fund disclosures. Securities Act and ICA claims, and most Advisers Act claims, do not require a showing of *scienter* (i.e., an intent to mislead through intentional, knowing or reckless wrongdoing), while most claims under the Exchange Act do require the plaintiff to prove *scienter*.
 - Potential legal theories for securities law claims include that registered fund prospectuses and marketing materials are materially misleading in that they (i) over-promise what AI models can deliver in terms of enhanced returns, increased insights, portfolio optimization, etc. (so-called "AI-washing" claims of the kind we've already seen the SEC pursue¹); and/or (ii) understate the risks of the models failing to perform as designed and protect against losses.
- **Derivative common law claims:** Available to shareholder-plaintiffs only, these are common law claims asserted under state law by shareholders, ostensibly on behalf of the fund itself. Derivative claims tend to focus on breach of fiduciary duty and breach of contract claims against the fund's adviser, officers and/or directors. Derivative claims must pass several procedural hurdles (which vary somewhat by state) before they may be pursued in court on the merits, including that the potential plaintiff make a demand of the fund's board to pursue the claims on behalf of the fund in the first instance (which must then be unreasonably rejected by the board) or demonstrate that making such a demand would have been futile. As previewed above, relevant derivative claims may include allegations that fund managers and trustees violated their fiduciary duties by delegating investment decisions to AI models without basis to fully understand how the models were making decisions or whether they were doing so consistently with the intended strategy. Unlike traditional securities law claims, plaintiffs here would not necessarily be required to establish a causal link between the alleged AI-related failings and a decline in share price (but would need to establish a causal link between the supposed breach and actual damages of some type).
- **Section 36(b) excessive fee claims:** Under Section 36(b) of the ICA, investment advisers owe a fiduciary duty with respect to their receipt of compensation from the fund, and the provision creates both a private right of action for fund shareholders and an SEC enforcement action for alleged breaches of that duty (i.e., charging excessive fees). This provision might be used to assert claims against advisers that charging higher fees to a fund deploying AI models than those charged to a similar non-AI fund cannot be justified by any additional services, technology and costs required in using AI.

¹ On March 18, 2024, the SEC announced settled charges against two investment advisers, Delphia (USA) Inc. and Global Predictions Inc., for making false and misleading statements about their purported use of AI. See [In the Matter of Delphia \(USA\) Inc.](#) and [In the Matter of Global Predictions Inc.](#)

Using AI Tools for Asset Management Legal Work

The dawn of the AI age might not be all bad news (*i.e.*, more work) for asset management lawyers. In addition to the need to understand and explain both the pros and cons of using AI in the investment process, asset management lawyers may also have many opportunities to utilize AI tools in performing their jobs more efficiently—maybe even lightening the workload. This will require an openness to incorporating AI technology in tasks we may have traditionally thought of as uniquely suited for human lawyer brains. But then, we used to feel this way about making restaurant reservations and giving directions to cab drivers, so we know we can grow and adjust. Just because our work product is generally in the form of paragraphs rather than spreadsheets and charts doesn't mean that AI can't be put to work for us, too.

Following are some examples of potential AI-driven solutions that creative lawyers might develop and utilize:

- Utilizing a training dataset of extant public disclosures by funds and advisers (*e.g.*, registration statements, Form ADVs), deploy LLM tools to generate and/or sense-check your firm's disclosures on AI use cases and risks, or many other topics beyond AI.
- Prepare draft minutes of fund board meetings, using LLM tools and a dataset of prior exemplar board minutes, agendas, rough meeting notes, and potentially (with proper disclosure and risk mitigation rules) voice recognition software.
- Generate response materials for the annual fee approval process under Section 15(c) of the ICA, based on datasets of prior years' versions and updated input information on performance, personnel, fund lineups, etc.

Several existing commercial and proprietary platforms could be readily adapted to develop generative AI tools like these—in any number of industry-related legal contexts—built on public and internal data sets that could be compiled with relative ease given some organized effort.

Just as the use of AI in the investment process may pose new risk considerations under the securities laws, so too the use of AI by counsel may give rise to challenges and risks under the ethical obligations, procedural rules and practical standards guiding the profession. Examples to bear in mind include:

- **Record over-creation:** The use of voice recognition software to generate meeting summaries and minutes means that many more oral communications are being recorded in some fashion than would have been traditionally. A default practice of recording internal video meetings among colleagues might lead to the creation of unhelpful written records of candid discussions that later could be subject to civil discovery, litigation “hold” requirements, and/or standing record-keeping rules (such as those applicable to investment advisers and broker-dealers), even if a “raw” transcript is later edited and paraphrased. It will be increasingly imperative to have a clear understanding of how these features operate and whether and where any records created are stored. Relatedly, questions will likely arise as to whether a user's prompts and other written exchanges with LLM models like ChatGPT or AI-based legal research tools must be preserved and potentially produced like other now “traditional” electronic communications like email.
- **Privilege ambiguity:** This century's explosion of informal corporate communications via email, text messaging and instant messaging greatly complicated determinations of what communications are protected from production in litigation and regulatory investigations by the attorney-client privilege and related doctrines. AI tools may present the next generation of complexity on this topic. If a non-lawyer employee submits a prompt on a legal question to a company-provided LLM tool (and neither the prompt nor the output will be seen by a lawyer), is there a basis to claim privilege protection for that exchange? If there is privilege, is it waived if the LLM accesses the internet to generate responses and/or uses the prompt and output to train models that may be used by other companies?

EXPLORING OPPORTUNITIES TOGETHER

We would welcome the opportunity to speak with you further on the above topics and any others involving the adoption of AI in asset management. Ropes & Gray, as a thought leader in AI's expanding impact on our clients' business practices and our legal services to them, is excited to partner with forward-looking industry leaders in exploring these opportunities together.

PRINCIPAL CONTRIBUTORS

**AMY D. ROY****Partner***Boston*Amy.Roy@ropesgray.com

+1 617 951 7445

**ROBERT A. SKINNER****Partner***New York and Boston*Robert.Skinner@ropesgray.com

+1 212 596 9927