FAIRVIEW

# AI Compliance FAQs for Registered Investment Advisers

Artificial intelligence can make advisory firms faster and more efficient, but it doesn't change an adviser's fiduciary duty or regulatory obligations. These frequently asked questions may help serve as a resource as you work on adopting AI at your firm while maintaining a sound compliance program.

## The Regulatory Landscape

### Is there an SEC rule specifically governing how RIAs use AI?

Not currently. The SEC proposed a rule in 2023 addressing conflicts of interest from predictive data analytics but withdrew it in 2025. That doesn't mean AI is unregulated. Existing rules apply fully to AI-assisted activity, and examiners are actively testing how firms apply them.

### Which existing rules apply to AI use?

Although there is no specific rule regarding AI usage, there are some existing rules where AI does come into play. The most relevant are:

- **Fiduciary Duty.** Advisers owe clients duties of care and loyalty whether a decision is made by a person, a model, or both.
- **Compliance Rule (Rule 206(4)-7).** Firms must maintain written policies and procedures reasonably designed to prevent violations, which now includes supervising AI use.
- **Marketing Rule (Rule 206(4)-1).** Statements about AI capabilities must be accurate, fair, and substantiated.
- **Books and Records Rule (Rule 204-2).** AI-generated communications and records may need to be retained.
- **Regulation S-P.** Client information processed by AI tools must be protected, and AI vendors may be service providers subject to oversight requirements.

State laws also matter, particularly state privacy laws and call-recording consent requirements.

### Do my regulatory obligations shift to the AI vendor?

No. An adviser's obligations don't transfer to an AI vendor. If you send a client an AI-drafted email that contains an error, the error is yours. Approach AI tools the way you approach any critical vendor: with a risk assessment, appropriate controls, and training.

# Disclosures and Marketing

### What is “AI-washing”?

AI-washing means overstating or misrepresenting how a firm uses AI, such as claiming AI drives investment decisions when it doesn’t. The SEC has brought enforcement actions against advisers for misleading AI claims, and examiners now routinely ask firms to substantiate AI statements in their marketing and disclosures.

### What should we do before mentioning AI in marketing?

Make sure every statement is accurate and that you can prove it. Keep a file supporting each AI-related claim on your website, in pitch materials, in social media, and in audio or video content. Review those claims whenever your AI use changes, and route AI-related marketing through your normal Marketing Rule review.

### Do we need to disclose AI use in Form ADV?

It depends on how you use it. If AI plays a material role in investment advice, portfolio management, or other client-facing services, consider whether Form ADV Part 2A should describe that use and its associated risks. Examiners have requested all disclosures relating to AI, including risk disclosures, so disclosures should match actual practice. Firms using AI only for internal efficiency, such as drafting or summarizing, may have less to disclose, but should still document that analysis.

### Can AI create conflicts of interest?

Yes. An AI tool could, for example, be configured or trained in a way that favors products or outcomes that benefit the firm over the client. Examiners have asked for documents showing how firms manage conflicts related to AI outputs. Identify any such conflicts in your AI inventory and address them through disclosure or mitigation.

# Approving AI Tools

### What due diligence should we do on an AI vendor?

Treat AI products like any other vendor, with added focus on data. Review contracts, terms of service, and data processing agreements. Confirm where data is stored, who can access it, and whether firm or client data may be used to train the vendor’s models. If the vendor qualifies as a service provider under amended Regulation S-P, apply your service provider oversight procedures. Also ask about AI use in *all* vendor due diligence, since many vendors don’t heavily advertise their AI features.

### Which subscription level should we use for platforms like ChatGPT, Claude, or Gemini?

Use the highest enterprise or business tier available. Higher tiers typically offer stronger security, contractual commitments about not training on your data, and administrative controls such as single sign-on, retention settings, and audit logs. Free and consumer versions generally lack these protections and should be prohibited for firm business.

### An application we already use just added AI features. Are they automatically approved?

No. New AI features should go through the same review as a standalone AI tool. Review updated agreements, confirm how data is used, and determine whether the feature can be disabled. If a feature hasn't been reviewed and approved, disable it. Unreviewed add-ins are a common source of "shadow AI."

### What are AI connectors, and what risks do they pose?

AI connectors link a generative AI platform to other systems, such as email, cloud storage, or your CRM. They create real efficiencies, such as drafting a client follow-up from CRM meeting notes, but they also raise questions. Was the output reviewed before it was sent? Was it captured in your books and records? What data can the AI now access? Approve connectors centrally, and require human approval for any connector action that writes, sends, or deletes data.

### How is agentic AI different from generative AI?

With generative AI, a person reviews the output and decides what to do. With agentic AI, the system makes decisions and takes actions, sometimes across multiple systems, before a person sees the result. Agentic tools warrant tighter controls: restricted access, a more rigorous approval process, ongoing monitoring, and periodic audits to confirm they still work as intended. Some platforms enable agentic features by default, so make a deliberate decision about whether to turn them on.

## Recordkeeping and AI Notetakers

### Are AI-generated content and chat logs considered books and records?

They can be. The test is the same as for any other record: if AI output is a written communication relating to advice, recommendations, or other covered topics, or if it is used to satisfy another recordkeeping requirement, it must be retained. Set your AI platform's data retention to match your books and records policy rather than accepting the default, and where available, enable the platform's compliance or audit API so content can feed your archive.

### Do we have to retain AI notetaker transcripts and recordings?

Advisers aren't subject to the call-taping requirements that apply to some broker-dealers. The question is how you use the output. If a transcript documents investment instructions or decisions, or if you send a transcript or summary to clients, it is likely a record that must be retained. Firms generally choose between archiving everything the notetaker captures or retaining only what is shared or relied on. The second approach reduces volume but risks gaps.

### Do we need consent to use an AI notetaker on client calls?

Often, yes. Most states require one-party consent to record, but several require consent from all parties. Because clients may be located in different states, many firms obtain consent from all participants as a matter of practice. Transparency also protects the relationship, since clients may be surprised or uncomfortable to discover a recording. If a client declines, take notes another way.

### What other risks do AI notetakers create?

- **Sensitive data.** Calls may include personal information, material nonpublic information, or other sensitive details that must be protected.
- Anything recorded or transcribed can be requested by regulators or in litigation.
- **Attorney-client privilege.** Sharing privileged communications with a third-party AI vendor may waive privilege. Consider disabling notetakers for calls with counsel.
- AI summaries can misstate amounts, timing, or risk ratings. A person should review summaries before they are shared or relied on.

## Data Security and Privacy

### Can employees enter client information into AI tools?

Only into tools the firm has approved for that purpose, under an enterprise agreement that protects the data. Your AI acceptable use policy should specify which tools may receive client information, nonpublic personal information, or MNPI, and which may not. Consumer and free tools should generally be off limits for any confidential data.

### How does Regulation S-P affect AI use?

Amended Regulation S-P requires advisers to maintain an incident response program, notify affected individuals of certain breaches, and oversee service providers that receive customer information. An AI vendor that processes client data may be a service provider, which means due diligence, contractual protections, and ongoing monitoring. AI-related incidents involving customer information should be covered by your incident response program.

### What is prompt injection?

Prompt injection occurs when malicious instructions are hidden in content an AI tool processes, causing it to act against its intended purpose. With *direct* injection, a user types instructions designed to override the tool's safeguards. With *indirect* injection, an attacker plants instructions in an email, document, web page, or CRM note that the AI later reads. Indirect injection is harder to prevent because it can occur within routine workflows, and it could lead to data leakage or unauthorized actions.

### How can we reduce prompt injection risk?

- Train employees on what prompt injection is and how it works.
- Restrict AI browsing and code execution to approved sites.
- Require human approval for higher-risk actions, such as sending emails, executing code, or finalizing client documents.
- Follow the principle of least privilege for both users and AI tools.
- Keep AI tools and integrations updated and patched.
- Monitor activity with endpoint detection and response, SIEM, and intrusion detection or prevention tools.

## Supervision, Policies, and Training

### What should an AI policy include?

At a minimum, an AI acceptable use policy should cover:

- Permitted and prohibited AI tools, platforms, and use cases.
- What data may and may not be entered into AI tools.
- Required human review before AI output is sent, published, or relied on.
- The approval process for new tools and new AI features in existing tools.
- Recordkeeping, monitoring, and incident reporting requirements.
- Training requirements and consequences for violations.

Firms whose employees build AI tools, share prompts, or code should also adopt a software development lifecycle policy, so new tools are reviewed, tested with a small group, and approved before firm-wide release.

### What is "shadow AI," and how do we prevent it?

Shadow AI is the use of AI tools the firm hasn't approved, such as an employee using a free chatbot on a work device or an unreviewed AI feature inside an approved app. Reduce it by clearly listing approved tools, disabling unapproved add-ins, providing a secure enterprise alternative, training employees, and monitoring usage. Identity controls like domain capture on enterprise platforms can also surface employees who signed up on their own with a work email.

### How should we monitor AI use?

Monitor both approved and unapproved use. Options include usage logs from enterprise AI platforms, custom reports built from those logs, device-level software that tracks activity across popular AI tools, and third-party monitoring services. Pair monitoring with periodic access reviews, and audit any AI agents to confirm they continue to work as intended.

### How often should we train employees on AI?

Training should happen at rollout and on a regular schedule, at least annually and whenever tools or policies change meaningfully. Cover approved tools, prohibited uses, data handling, prompt injection, and human review expectations. Keep records of completion. Examiners have specifically requested evidence of AI-related compliance training.

### What should we do if AI fails or produces inaccurate results?

Have a documented contingency plan. Identify which business functions depend on AI, how you would continue without it, and how errors are detected, escalated, and corrected. For AI supporting critical functions, benchmark outputs against expected results, especially after material changes. Log incidents where AI use raised regulatory, ethical, or legal issues, along with how they were resolved.

## SEC Exams and Getting Started

### What are SEC examiners asking about AI?

Recent exam requests cover four broad areas:

- **Disclosures and marketing:** all materials that mention AI, including audio and video, plus a categorized description of every AI use.
- **Models and performance:** descriptions of AI models used in investment decisions, trading signals, performance reports, and validation reports.
- **Data:** data sources, providers, acquisition errors, and data security documentation.
- **Governance:** AI policies and amendments, conflicts, contingency plans, incident reports, responsible personnel, committees, and training records.

Examiners have also asked how firms use AI to write code that implements analytical methods, focusing on the process and oversight rather than the code itself.

### We only use AI for basic tasks like drafting emails. Do we still need to prepare?

Yes, though the scope will be smaller. Even limited use raises questions about approved tools, data security, recordkeeping, supervision, and training. Firms that market AI or use it for core functions such as portfolio management should expect the most detailed requests.

### Where should our firm start?

- **Inventory and risk-assess** every AI tool in use, including AI features in existing software.
- **Adopt written policies**, including an AI acceptable use policy and, if you build tools, a development lifecycle policy.
- **Configure platforms securely**, with enterprise licensing, retention aligned to your records policy, single sign-on, and approval requirements for higher-risk actions.
- **Train employees** on AI-specific risks and expectations.
- **Monitor usage** and validate outputs on an ongoing basis.
- **Review disclosures and marketing** to confirm every AI statement is accurate and supported.

## How can Fairview help?

Fairview's **AI Compliance Services** help RIAs adopt AI with confidence while meeting their regulatory obligations. Our team can conduct AI audits and risk assessments, draft AI acceptable use and development lifecycle policies, deliver AI-specific employee training, and set up AI usage monitoring.

**Contact us** with questions about AI compliance or to discuss support for your firm.

**About Fairview**

Founded in 2005, Fairview®, LLC, provides a full range of back-office support services for investment advisers and other financial institutions. For more information, visit FairviewInvest.com.