

ARTIFICIAL INTELLIGENCE

Why Companies Should Consider Requiring Internal Disclosure of AI Use

BY: CHARU A. CHANDRASEKHAR, AVI GESSER, KAREN LEVY, WILLIAM SADD AND PATTY (VIRTUAL AI SPECIALIST) - FEBRUARY 22, 2026



SHARE

Most companies that allow generative AI (“GenAI”) in the workplace have a policy requiring users to check AI outputs to ensure that they are accurate, complete, and fit for purpose. But once an employee reviews the AI output and is comfortable with it, the employee can thereafter circulate that document internally or externally as if they drafted it themselves, without disclosing that a





portion (which could be significant) was generated by AI.

Based on our experience representing over 100 companies with their AI adoption over the last year, we think companies should consider adding an AI disclosure requirement like this one to their AI policies:

To the extent you share any work product, either internally or externally, and (1) a substantial portion of the document was generated by GenAI, (2) the work product may be relied on in making decisions, and (3) mistakes or omissions in the work product could impact those decisions, then you must identify which parts were generated using GenAI.

While some might worry that this kind of disclosure could stigmatize AI use and slow work down, we think these concerns are overstated.

Why Require Disclosure of Internal AI Use?

There are at least three reasons to require internal AI-use disclosure.

1. Visibility and Disclosure Strengthen AI Governance

One of the goals of AI governance is to give the company visibility into where AI is being used, so it can (a) identify the good use cases that should be encouraged, given additional resources to improve, and be scaled across the company, and (b) identify the bad use cases that should be discouraged and addressed in future AI guidance and training as examples of how not to use AI.

But in some companies, that visibility is lacking, in part because employees receive mixed messages about AI use. Employees are given access to AI tools, but AI policies and training focus primarily on the risks of AI use. Senior leadership encourages use of AI, but some members of middle management still seem skeptical. As a result, some employees hear that AI use is tolerated, but not something to embrace openly, so both good and bad use cases remain hidden. Requiring internal disclosure of AI provides greater visibility into use cases, both good and bad, thereby improving overall AI governance.

Requiring internal disclosure also normalizes use of AI as legitimate, appropriate, and beneficial. When respected culture carriers—senior leaders, high-performing managers, trusted subject-matter experts—disclose their AI use, it sends an important message in companies that are trying to increase their AI use: “This is how we work now.” It becomes less professionally risky to use AI appropriately, and more professionally risky to hide it.

2. People Review AI Work Differently, And They Should

Even if an employee has thoroughly checked and signed off on an AI output, the supervisor reviewing that document may evaluate it differently depending on whether AI was used to generate it.

In most professional settings, supervision of work involves calibrated deference. If a trusted colleague spent many hours building a model, reviewing a complex regulatory regime, drafting a



negotiation strategy, or synthesizing technical research, our review of that work is usually limited, in part because we assume that our colleague will flag any aspects of the work where they are unsure or there wasn't an answer. We may ask them some questions, review a sample of the source materials, or spot check some of the conclusions that are not familiar to us. But if the task is within our colleague's area of competence, and they have expressed confidence in the analysis and conclusions, we tend to defer to them.

With AI, that assumption may no longer be warranted. First, AI allows employees to create work product quickly that looks like it was the result of subject-matter expertise-driven research and careful drafting. Second, unlike our colleagues, the AI rarely says things like "I'm not sure about this point" or "I could not find an answer to that question." A reviewer who is unaware that AI played a significant role in drafting a document may therefore mistakenly believe that they are largely deferring to a trusted and diligent colleague, when, in reality, they are unknowingly largely deferring to AI.

For example, an AI-generated summary may confidently cite a non-existent exception, omit a limitation, or misstate a standard—all of which are errors that may not get caught by a "reasonableness" review because the AI's writing is fluent and confident, and its mistakes are often not obvious. So while we may have confidence in a colleague's ability to research and draft a particular document from scratch, we may not have that same level of confidence in their ability to review an AI-generated draft of a document—particularly outside of their subject matter expertise—and successfully identify all the errors and omissions.

3. Proper Attribution of Good (and Bad) Ideas

Sometimes a reviewer sees something in a draft document that is remarkable. It could be remarkably good or remarkably bad, but either way, it is important for a company to know whether that remarkable idea, argument, or insight originally came from AI, from an employee, or a combination of both. If it came from the employee, it might be important for that employee's overall evaluation. If it came from AI, it may reveal a capability or weakness of the AI that was previously unknown and should be communicated to others. If it resulted from effective prompting by an employee who iterated with the model to achieve a remarkable result, that prompting technique may be valuable for future AI training programs.

Conclusion

Requiring internal disclosure of AI use in certain circumstances provides companies with increased visibility and normalizes AI use, while ensuring that AI-generated outputs receive the appropriate amount of review and scrutiny before being relied upon for making important decisions. This disclosure framework promotes responsible and beneficial AI adoption across the enterprise.

* * *

This week marks the sixth anniversary of the Debevoise Data Blog. On February 20, 2020, we published our first post: [Fifteen Ways to Reduce Regulatory and Reputational Risks for Your AI-Powered Applications – Lessons from Recent Court Decisions and Regulatory Activity](#), which



actually holds up pretty well. A big thank you to our almost 4,000 subscribers for their continued interest and support.

To subscribe to the Data Blog, please [click here](#).

[The Debevoise STAAR \(Suite of Tools for Assessing AI Risk\)](#) is a monthly subscription service that provides Debevoise clients with an online suite of tools to help them fast-track their AI adoption. Please contact us at STAARinfo@debevoise.com for more information.

The cover art for this blog post was generated by ChatGPT 5.2.



Charu A. Chandrasekhar

Charu A. Chandrasekhar is a litigation partner based in the New York office and a member of the firm's White Collar & Regulatory Defense and Data Strategy & Security Groups. Her practice focuses on securities enforcement and government investigations defense and artificial intelligence and cybersecurity regulatory counseling and defense. Charu can be reached at cchandra@debevoise.com.



Avi Gesser

Avi Gesser is Co-Chair of the Debevoise Data Strategy & Security Group. His practice focuses on advising major companies on a wide range of cybersecurity, privacy and artificial intelligence matters. He can be reached at agesser@debevoise.com.



Karen Levy

Karen Levy is the Chief Information Officer at Debevoise and serves on the firm's AI Governance Committee.



William Sadd

William Sadd is the Head of Practice and AI Systems at Debevoise. He can be reached at wjsadd@debevoise.com.





Patty (Virtual AI Specialist)

Patty is a virtual AI specialist in the Debevoise Data Strategy and Security Group. She was created on May 3, 2025, using OpenAI's o3 model.



PREV POST

Update: Judge Rakoff Issues Written Opinion that AI-Generated Documents Are Not Protected by Privilege

- 7 MINS READ

NEXT POST

Debevoise Featured in Law360 Story on Vibe Coding

- 2 MINS READ



Related Posts



Debevoise Featured in Law360 Story on Vibe Coding

FEBRUARY 24, 2026



Update: Judge Rakoff Issues Written Opinion that AI-Generated Documents Are Not Protected by Privilege

FEBRUARY 17, 2026



AI Is Making Everyone a Potential Competitor – Are Your Data Contracts Ready?

FEBRUARY 16, 2026

