The Regulatory Review

# AI Providers Should Not Be Liable for Users' Securities Violations

Jack Solowey

April 8, 2024

If you are a policymaker looking to assert jurisdiction over a hot new sector, you might try to put deep-pocketed players on the hook for harms within your purview.

This strategy could explain a bill introduced in the U.S. Senate Banking Committee—the Financial Artificial Intelligence Risk Reduction (FAIRR) Act—that would make artificial intelligence (AI) providers liable for uses of their tools that violate securities laws, unless those providers take reasonable preventive steps.

This liability regime might be understandable given policymakers' incentives. Nonetheless, it would produce poor public policy by clumsily assigning blame in a way that clashes with standard economic and legal frameworks for determining when producers should be liable for harms associated with the use of their products.

Ideally, liability rules would discourage violations of individuals' rights without discouraging the productive enjoyment thereof. Economists try to strike this balance by assigning liability to the "least cost avoider"—the party for whom preventing or ameliorating rights-infringing activity is the cheapest.

It appears unlikely that AI tool providers are the least cost avoiders of securities law violations. If, for example, a registered securities firm, such as a broker-dealer, investment adviser, or investment company, uses a generic AI model, it is far likelier that the firm, not the model developer, would have an advantage in mitigating securities law risks—including through a dedicated compliance team.

It may be more tempting to view the AI provider as the least cost avoider when an individual, not a registered securities firm, uses the tool. But even when that person is no securities law expert, it is still probably cheaper for the individual to handle the legal issues than it would be for the AI provider. Paying for legal advice about a known and specific activity is usually going to be less expensive than devising a compliance program to guard against every possible permutation of securities violation.

This is not to say that policymakers cutting the Gordian knot of AI liability by putting model providers on the hook have nothing going for them. For instance, identifying a clearly liable party could present lower administrative costs for the legal system. Richard Epstein, professor at the New York University School of Law, writes of the "*twin* objectives" of legal rules: "reducing administrative costs" *and* "setting desirable incentives." When these aims are in tension, though, Epstein observes that the question becomes "whether the savings in administrative costs is dissipated by the creation of inferior incentive structures."

Making AI providers liable for securities violations generally would produce inferior incentives. Tyler Cowen, economist and professor at George Mason University, argues that "placing full liability on AI providers for all their different kinds of output, and the consequences of those outputs, would probably bankrupt them." That would reduce valuable AI innovations that benefit securities market participants.

In addition, in situations where an AI user is directly responsible for a securities violation and the AI provider merely failed to prevent that violation, it still would be likely more enticing to launch enforcement and private actions against the AI provider wherever possible, as such a provider would tend to be a higher profile target with the seeming ability to pay a large judgment or settlement. For this reason, making AI firms liable for such securities violations would perversely have those companies regularly pick up the tab for the parties clearly and directly responsible for violations.

AI provider liability for securities violations does not look any better through the lenses of the common law, such as products liability or agency law. Neither doctrine suggests universal AI provider liability for resulting harm is appropriate.

For one, in most of the United States, provider liability for faulty products only covers physical injuries to people and property, not "purely economic losses," which would typically characterize securities violations. Moreover, important (but not uniformly applied) product liability considerations ask whether the user modified the product or used it in an abnormal fashion, either of which could forestall the provider's liability. When a user prompts an AI tool to produce a securities violation, there is a reasonable argument that the user engaged in modification or abnormal use. At the very least, this is the type of argument that courts should consider.

Agency law reveals similar issues. The typical rule is that principals (on whose behalf agents act) are liable for acts of their agents that are within the scope of the agents' employment. In the AI context, even if we assume that the AI provider could be the principal to the digital AI agent, there remains the question of whether the AI agent acted within the scope of "employment." Because these questions are often litigated with mixed results, a liability regime that ignores them would clash with the time-tested subtleties of the common law.

Key principles of U.S. securities law also weigh against uniform AI provider liability. Specifically, securities laws have varying state of mind requirements. Notably, private fraud actions under the Securities Exchange Act and its implementing regulations have been interpreted by the Supreme Court to require allegations of the defendant's "intent to deceive, manipulate, or defraud." The FAIRR Act would circumvent such requirements by deeming the AI provider to have the relevant state of mind. Notably, a recent speech by U.S. Securities and Exchange Commission Chair Gary Gensler interpreted the FAIRR Act as proposing to impose a strict liability standard.

Undermining intent requirements not only would clash with existing law, but it would also reveal the inefficiency of placing blame with AI providers. State-of-mind requirements can, in part, codify important intuitions about whether a party was in any good position to avoid harms. For example, in asking whether harm was foreseeable, the tort law's negligence standard essentially asks whether it would have made any sense for the defendant to act in a way that would have averted the harm. In removing, or at least downgrading, AI provider state-of-mind requirements for securities violations, the FAIRR Act would allocate liability for conduct in ways that have not previously made sense to policymakers or courts.

Grappling with liability questions in the age of AI is critical. But taking the lazy way out by blaming those building AI tools that all of us stand to benefit from is not the answer, particularly when doing so clashes with both economics and longstanding legal principles.

*Author Biography: Jack Solowey is a Policy Analyst at the Cato Institute's Center for Monetary and Financial Alternatives, focusing on financial technology.*