THE SHARED POLICY QUESTION FOR AI AND CRYPTO:

When are Developers Liable for Writing Code?

DeFi Education Fund

DEFI EDUCATION FUND
1155 F Street NW, Suite 1050
Washington, DC 20004, United States
defieducationfund.org

DEFI EDUCATION FUND
1155 F Street NW, Suite 1050
Washington, DC 20004, United States
defieducationfund.org

The Shared Policy Question For AI and Crypto:
When are Developers Liable for Writing Code?

The policy question most likely to shape the next century of American technology development is a legal one: when, if ever, does building software make you legally responsible for how other people use it? This liability question more than any other (not even those related to compute, chips, or talent) will determine where the next century of innovators choose to build and operate transformative businesses.

Imagine a startup develops an open source AI agent that automates routine cybersecurity work. Companies use it to identify vulnerabilities in their own systems, verify software patches, and test their networks for weaknesses. That same agent can also be downloaded and modified by criminals to scan hospitals for exposed servers before launching ransomware attacks. Though the startup developers never interact with the criminals and have no control over how the agent is used by third parties, should they face legal liability because they should have anticipated that a bad actor would misuse their software in a way they never intended?

Let's put aside the professional software engineer and imagine an average person who "vibe codes" their own agent; for example, a teacher with no technical background that uses an AI coding assistant (e.g., Claude) to create a simple scheduling agent for her neighborhood youth soccer league. She posts the code online so other teachers can also benefit from it. Later, a third party downloads the same code, modifies it to automatically send phishing emails that impersonate the school, and scams people out of thousands of dollars. Should the teacher face liability because she created the software?

This question of third party liability becomes more relevant everyday as blockchain technology and AI rapidly democratize access to software development and make it easy for the average person to build tools. Millions of Americans who have never before written a line of code are now able to create and share sophisticated code simply by "vibe coding" – i.e., describing in plain English to an AI assistant what kind of technology solution they want to create.

Blockchain developers and those advising them have been trying to address this liability question for years. The debate came to a head this summer, when the U.S. Senate deadlocked over the Clarity Act and one provision in particular: the Blockchain Regulatory Certainty Act (BRCA). The BRCA clarifies that a developer who creates software that allows third party users to transfer their own funds is not a "money transmitter" or operating a "money transmitting business." Money transmitting businesses are regulated financial institutions that accept and transmit money and, as a result, must register with FinCEN and comply with applicable anti-money laundering (AML) obligations. For example, PayPal and Venmo are money transmitting businesses and, because they have control over your money and transfer it on your behalf, they must comply with obligations under the Bank Secrecy Act (BSA).

The BRCA codifies the principle that software developers who do not have control over other people's money ("non-controlling developers") are fundamentally different from these kinds of businesses

because non-controlling developers merely create tools that others rely on to engage self-directed financial transactions and the developers have no control over the end users' actions. The BRCA would have benefitted non-controlling developers in every segment of the software development world, not just those building on blockchain. Without it, non-controlling developers building tools that allow others to transfer their own funds do not have clarity as to when or if they will be considered a money transmitter subject to AML obligations.

While the blockchain industry was united in advocating in favor of the BRCA, not all lawmakers were on board. There was a surprising amount of pushback from lawmakers, incumbent financial intermediaries, and law enforcement groups related to the scope of this narrow provision shielding non-controlling developers from one specific type of liability related to money transmission – imagine if a provision actually addressed the full spectrum of third party conduct for which a developer could face liability.

The AI industry is about to feel the full weight of trying to address this third party liability question for its own technology development community: does creating and deploying software make a developer responsible for how third parties use it?

The History of U.S. Digital Technology Protections

American law has traditionally erred in favor of protecting those building innovative technologies. For example, the Computer Fraud and Abuse Act (CFAA), which is the primary federal law criminalizing unauthorized access to computers and computer systems, including by hacking, data theft, and fraud, applies only to those who "knowingly" or "intentionally" gain access to a computer system without authorization. In other words, the law applies to the wrongdoer who engages in the criminal conduct at issue and not the creator of the computer system. Similarly, Section 230 of the Communications Decency Act (CDA) established that the people who build and operate online platforms are not civilly liable for information that unrelated third parties posted to them. While the CDA has been subject to debate in its application to certain illicit conduct, Congress's intent in passing the law was clear: liability generally follows from bad conduct, and providers of software platforms should not be held criminally responsible for the conduct of third parties they cannot control.

That single principle is quite literally the foundation of the freedom to build the modern internet, and the reason it was built onshore in the United States rather than somewhere else.

Some argue that this issue has also been addressed in the context of cybersecurity. For example, in the Napster case, *A&M Records v. Napster*, the Ninth Circuit confronted whether the operator of a peer-to-peer file-sharing service could be held liable for copyright infringement committed by its users. The court did not hold that developing peer-to-peer file-sharing software technology itself created liability. Rather, Napster's liability stemmed from operating centralized servers that maintained a searchable index of files available from users, having actual knowledge that its system was being used for infringement, and retaining the ability to supervise the system and block users.

DEFI EDUCATION FUND
1155 F Street NW, Suite 1050
Washington, DC 20004, United States
defieducationfund.org

Several years later in the Grokster case, *MGM Studios v. Grokster*, the Supreme Court considered decentralized file-sharing software whose developers had little technical control over users' actions once the software was distributed. Unlike Napster, Grokster did not operate central indexing servers and the Court declined to impose liability on Grokster based on its developers having merely created a technology capable of copyright infringement. Instead, the Court relied on a theory of inducement, holding that a technology provider could be liable where it distributed a product *with the intent to promote the unlawful use of copyright infringement*, as demonstrated through affirmative steps encouraging unlawful activity. And the *Grokster* Court expressly stated it sought to avoid "discouraging the development of technologies with lawful and unlawful potential," which is why it held that the law requires evidence of "purposeful, culpable expression and conduct" rather than mere knowledge that a technology could—or actually did—facilitate unlawful activity. Critically, *Grokster* also left intact the *Sony v. Universal City Studios* holding that distributing a technology capable of substantial lawful uses does not, standing alone, create developer liability merely because the developer knows that some users may use it unlawfully.

The distinction drawn in these cases has direct relevance to the emerging debates over blockchain and AI developers. The cases collectively reflect reluctance to impose liability stemming from a developer creating a general-purpose technology capable of both lawful and unlawful uses, and instead, preserve potential liability where additional culpable conduct exists—such as intentionally inducing unlawful copyright infringement.

While courts have recognized these distinctions in the copyright context, the principle should apply more broadly: liability should turn on a developer's own culpable conduct, not on mere authorship or the foreseeable misuse of technology by third parties. This is common sense. We would never expect to hold a hammer manufacturer liable if someone buys their hammer and damages property with it. Yet this basic logical principle is now under assault in the context of blockchain and AI developers.

The Supreme Court on Third Party Liability

Outside the world of software development, the Supreme Court has drawn a clear line on third party liability for the development of other tools. In June 2025, in *Smith & Wesson Brands v. Estados Unidos Mexicanos*, a 9-0 unanimous Supreme Court rejected Mexico's attempt to hold American firearms manufacturers liable for cartel violence.

How does this relate to software development? Mexico's legal theory was that the manufacturers knew some of their products ended up in criminals' hands and they kept selling them anyway. Justice Kagan's opinion held that this was not enough to hold American gun manufacturers liable for the actions of unrelated third parties, explaining that aiding and abetting requires "conscious and culpable participation" in a specific wrong, generally requires an affirmative act rather than a failure to act, and cannot be built out of ordinary lawful commerce plus downstream knowledge.

Swap "firearms" for "software" and you have a direct analogy to the software developer liability debate. Software systems are, of course, fundamentally different than firearms, and *Smith & Wesson* does not itself resolve the legal questions surrounding developer liability, but the Court's broader principle has direct relevance: a tool creator does not become liable for the independent misconduct of third parties merely because they know that end users may misuse its tool. The question now is whether Congress and the courts will apply the same logic when the product is code underpinning blockchain or AI systems.

The Battle Over Developer Protections for Blockchain Developers

Over the last two years, one of the larger battlegrounds in the push to pass market structure legislation for crypto was related to protections for non-controlling developers and infrastructure providers.

The Blockchain Regulatory Certainty Act (BRCA), started as a bipartisan standalone bill in both the House and the Senate, and was later included in the CLARITY Act as a critical part of a package of developer protections. The BRCA would put a seemingly straightforward principle into statute: a non-controlling developer who publishes software — with no control over user transactions or funds — is not a money transmitter under the Bank Secrecy Act. In other words, software developers who create blockchain technology for others to use to transfer their own funds, with no retained control over the tool itself or how it is used, are not liable for how unrelated third parties use that tool. The crypto industry was universally aligned in its support for the BRCA's inclusion in a market structure bill.

The BRCA aligns the law with the Department of Treasury's Financial Crimes Enforcement Network (FinCEN) 2019 Guidance on Convertible Virtual Currencies, which makes clear that only those who act as intermediaries and have "total independent control over" user assets are subject to BSA obligations. The BRCA does not alter existing AML laws (e.g., 18 U.S.C. § 1956) or limit prosecutions for fraud, sanctions evasion, trafficking, or terrorism financing. Critics of the BRCA often painted it as sweeping far more broadly in its protections for developers than it did, but in reality, the BRCA only affected the laws specifically related to licensing and failing to obtain a license for operating a money transmitting business.

At a late stage in the market structure legislative process, law enforcement groups raised vocal opposition to the BRCA. The National District Attorneys Association, National Association of Assistant United States Attorneys, International Association of Chiefs of Police, and National Sheriffs' Association argued in several letters that in their view, the BRCA's protections might create gaps in law enforcement's ability to investigate and prosecute illicit activity involving digital assets, while acknowledging that their concern is not with individuals who only write or publish software code.

Education on the realities of blockchain technology and on the numerous tools currently available to law enforcement and added by market structure legislation helped address many concerns expressed on the BRCA. For example, Senator Chuck Grassley, who initially raised concerns that the BRCA could create "blind spots" for law enforcement, ultimately reached a compromise with Senator Cynthia Lummis that

preserved the BRCA's core protections for non-controlling developers while addressing concerns regarding a specific criminal enforcement provision. Many law enforcement organizations also shifted: the National Fraternal Order of Police ultimately endorsed the Clarity Act and the National Organization of Black Law Enforcement Executives also came out in favor.

The same problem could soon arise with respect to the way the law treats AI developers. As increasingly autonomous AI tools perform functions that once required human intermediaries, the developers of those tools will need to advocate that they be treated appropriately under the law. The BRCA story illustrates an important lesson: even developer protections grounded in the seemingly straightforward distinction between building neutral technology and controlling its use can become the subject of significant opposition from law enforcement—and securing durable protections will require sustained engagement, education, and carefully negotiated limits that preserve authorities against actual bad actors without converting software development itself into regulated or criminal conduct.

California's Attempt to Expand AI Developer Liability

The fight over AI developer protections has actually already begun at the state level. In the absence of a federal AI regulatory framework, state legislatures have crafted ambitious legislation to regulate the AI industry, including by expanding developer liability.

For example, California SB 1047, which passed the legislature in 2024 before being vetoed by Governor Newsom, would have imposed a duty of reasonable care on developers of large models. Developers would have a duty to mitigate the risk of *foreseeable downstream misuse* or face enforcement by the attorney general, with penalties reaching 10 percent of the cost of the compute used to train the model and 30 percent for repeat violations. SB 1047 did not consider critical factors like whether the developer operated the system, controlled it at the moment of harm, or had any relationship to the person who caused the alleged harm.

As Anjey Midha, a prominent venture capitalist, explained in a blog Q&A for Andreessen Horowitz on SB1047, *"the bill requires developers to prove that their models can't possibly be used for any of the defined hazardous capabilities. But these definitions are way too vague, ambiguous, and subject to interpretation. How can a developer prove a negative, especially when the goalposts keep moving? It's an impossible standard to meet."*

Midha continued: *"It's like holding car manufacturers liable for every accident caused by a driver who's modified their car. The practical effect of these liability provisions will be to drive AI development underground or offshore. No rational startup founder or academic researcher is going to risk jail time or financial ruin just to advance the state of the art in AI. They'll simply move their operations to a jurisdiction with a more sensible regulatory environment, and the U.S. will lose out. Period."*

OpenAI, Google, and Meta opposed SB 1047, and Microsoft and other large AI developers pressed for substantial changes to the legislation. Nancy Pelosi, Ro Khanna, and other California Democrats also

spoke out against the legislation. Notably, the loudest objections to SB 1047 came from the open-source AI community; many community members pointed out that only companies with substantial legal departments and massive revenues could comply with such obligations, while research groups and other startups would suffer. While it was net-positive that Governor Newsom saw the danger of SB 1047 and vetoed it, the legislation is emblematic of attempts to expand liability for software developers in the AI space.

Consequences of Getting Developer Protections Wrong

The consequences of creating unworkable laws for emerging technology developers are severe for U.S. national security, economic dominance, and technological leadership. Developers could find themselves misclassified into regulated categories that don't apply to them or held criminally or civilly liable for the actions of unrelated individuals, and choose to build elsewhere. The impact is of misclassification and the novel application of third party liability to developers is evident:

- The U.S. share of open-source developers fell from 25% in 2021 to 18% in 2025, driven by a lack of clear, durable rules for software development.
- Electric Capital's developer census found that the American share of open-source crypto developers fell from roughly 38 percent in 2015 to about 19 percent in 2025. More than 80 percent of the people building this technology now live outside the United States, and in 2024, India onboarded more new crypto developers than the United States did.
- Stanford's 2026 AI Index found that the performance gap between the top American and Chinese models had narrowed to just 2.7 percent as of March 2026, after leading U.S. models held double-digit advantages on several major benchmarks in 2023, even as U.S. private AI investment reached $285.9 billion in 2025, more than twenty-three times China's $12.4 billion.
- Stanford's 2026 AI Index also showed that the number of AI scholars relocating to the United States has fallen 89 percent since 2017, with most of that decline in the past year alone.

This concerning trend demonstrates that the U.S. has to make an effort to keep and bring back emerging technology talent and innovation. There is no greater threat to national security than allowing major technologies to be built outside the U.S. Offshoring innovation shifts development into jurisdictions where the United States has less visibility, leverage, and access to information. Developers building outside of the U.S. are not required to comply with U.S. cybersecurity or anti-money laundering (AML) standards, and are not within U.S. law enforcement's jurisdiction, which hinders law enforcement's ability to monitor illicit activity, collect evidence, prevent transactions involving sanctioned entities, set security standards, and recover victim assets. Reversing the offshoring trend will require showing software developers that they can build in the U.S. without being treated unfairly under the law.

In conclusion: the AI era asks, "Are developers—or even ordinary people using AI to create software—liable for what autonomous systems or downstream users do?"

The reality is that policymakers, advocacy organizations, and legacy incumbents fearful of innovative competition will always want to expand liability for the developers of emerging technologies, like agentic

DEFI EDUCATION FUND
1155 F Street NW, Suite 1050
Washington, DC 20004, United States
defieducationfund.org

and decentralized systems, in an effort to hold *someone* responsible when something goes wrong. Of course, concerns about criminal misuse are valid and there must be safeguards against bad conduct. However, the correct path forward is to create laws tailored to the technology at issue and to the conduct that bad actors are actually responsible for.

The technologies will change, but the legal principle should not. Whether the tool is a blockchain protocol, an AI agent, or something that has not yet been invented, developers should be responsible for their own conduct—not for the independent actions of every person who uses the technology they create. Where a developer controls another person's assets, intentionally facilitates unlawful activity, or affirmatively participates in wrongdoing, existing law can and should hold them accountable. But writing code, publishing it, or building a tool capable of both lawful and unlawful uses should not itself become the basis for liability.

The debate over the BRCA was an early warning for the AI industry and, increasingly, for anyone who writes software. As AI makes software development accessible to millions more Americans, the line between "developer" and "user" will only continue to blur. The United States therefore has a choice about the legal architecture it builds for the next generation of technology: preserve the longstanding distinction between creating a tool and culpably misusing it, or create a regime in which innovators must anticipate and answer for every downstream misuse of what they build. If America wants the next generation of transformative technology developed here, the freedom to build must come with clear rules about when liability begins—and, just as importantly, where it ends.

If you have a perspective on what we've written, or you are an AI or blockchain developer keen to join us in the educational effort, please reach out. We'd love to hear from you: info@defieducationfund.org.

Amanda Tuminelli
CEO and CLO
DEFI EDUCATION FUND

Gavin Zavatone
Policy Lead
DEFI EDUCATION FUND