

STANFORD UNIVERSITY

THE STANFORD EMERGING TECHNOLOGY REVIEW 2026


A Report on Ten Key Technologies and Their Policy Implications

CO-CHAIRS Condoleezza Rice, Jennifer Widom, and Amy Zegart

DIRECTOR AND EDITOR IN CHIEF Herbert S. Lin | **MANAGING EDITOR** Martin Giles







CRYPTOGRAPHY AND COMPUTER SECURITY

KEY TAKEAWAYS

- Cryptography is essential for protecting information, but alone it cannot secure cyberspace against all threats; it must operate in concert with the broader field of computer security.
- Cryptography is the enabling technology of blockchain, which is the enabling technology of cryptocurrencies.
- Rather than pursue a central bank digital currency, the United States has adopted a policy preference for privately issued digital assets, promoting stablecoins and other cryptocurrencies as vehicles for financial innovation and resilience.

Overview

The word *cryptography* originates from Greek words that mean “secret writing.” Once limited to simple codes and ciphers, it now relies on advanced mathematics to protect data from unauthorized access or tampering.¹ Though largely invisible, cryptography secures many everyday interactions, from online shopping to cell phone calls.

Cryptography is essential for internet activity—from messaging and banking to everyday browsing—but it cannot, on its own, guarantee the confidentiality, integrity, or availability of information. Various vulnerabilities ensure that cybersecurity will remain an ongoing challenge. These include technical vulnerabilities in the digital systems that humans operate and use; human vulnerabilities, such as the tendency to bypass security mechanisms because using them is considered inconvenient; and strong incentives for attackers.

Cryptography is essential for internet activity— from messaging and banking to everyday browsing.

Cryptography Basics: Public Keys, Private Keys, and Hashes

Here's an example of how cryptography works: Drew wants to send a private message to Taylor. She scrambles (encrypts) the message using an encryption algorithm and sends the scrambled (ciphertext) version. When Taylor receives it, he unscrambles (decrypts) it to recover the original (plaintext) message. Ellen, an eavesdropper, tries to intercept the message and must find a way to break the cryptographic protection to see the plaintext.

One example of an encryption algorithm is the shift cipher, where each letter is replaced by one N positions later in the alphabet. If $N = 2$, A in the plaintext becomes C in the ciphertext, B becomes D, and so on. If $N = 3$, A becomes D. To decrypt the ciphertext, Taylor must know that the algorithm was the shift cipher and the key N —so if he sees C and knows $N = 2$, he writes down A. (Modern encryption is far more secure and complex than this example but also harder to explain.)

Both Drew and Taylor must share a secret: N , the cryptographic key—a string of digits used for both encryption and decryption. They must also know that the algorithm is the shift cipher. If Ellen learns both the algorithm and key, she can decrypt the message. This type of encryption, where the same key is used by both parties, is known as symmetric or secret-key cryptography. It requires secure key distribution—a way to share keys with intended recipients while keeping them from others.

Symmetric cryptography poses a practical problem: Parties have to meet in person to exchange secret keys before communicating securely. Imagine having

to meet every phone contact face to face before speaking. In the 1970s, Stanford professor Martin Hellman and Whitfield Diffie introduced asymmetric, or public-key, cryptography. This uses two keys: a public key, which anyone can use to encrypt a message and can be distributed over insecure channels, and a private key, known only to the intended recipient (see figure 3.1), which is needed to decrypt it. Although the keys are mathematically linked, deriving the private key from the public key would take longer than the age of the universe (unless quantum computing changes that, as discussed later in this chapter; for an in-depth discussion of quantum computing, see chapter 7, on quantum technologies).

Cryptography also enables the creation of secure hashes. A hash accepts a message of any length and computes a unique fixed-length string of numbers—called the hash value—corresponding to that message. Hashes have two key properties: It is extremely difficult to find another message with the same hash value, and it is infeasible to recover the original message from the hash value alone.

Using a secure hash function, the sender can use public-key cryptography to ensure integrity (protection against tampering) and identity (the message originated from the stated sender).

To illustrate, Alice (the sender) first computes the hash value of her message. Next, she encrypts the hash value with her private key, a process analogous to signing a document, generating a digital signature of the message's hash.² Alice then sends the message and its digital signature to Bob (the receiver).

Once Bob receives it, he can recover the hash value for the message that Alice purportedly sent and

compare that value to his own computation of the hash value. If these match, Bob can be assured that the message has not been altered in transmission and that Alice sent it, since only Alice could have used her private key to create a digital signature of the message's hash.

Messages can also be digitally time stamped. A known authoritative time and date server—such as the Internet Time Service, operated by the National Institute of Standards and Technology—accepts a message, appends the current date and time, and then provides a digital signature for the stamped message.


Computer Security

Computer security traditionally focuses on safeguarding computer systems against unauthorized access and misuse. It emphasizes the core principles

of confidentiality, integrity, and availability—with all three also known collectively as the CIA triad. Confidentiality refers to the privacy of data (i.e., preventing unauthorized disclosure). Integrity refers to preserving data (i.e., guarding against unauthorized alterations). Availability refers to data and resources being accessible to authorized users, especially during critical times.

Historically, computer security focused on protecting individual machines from actions perpetrated by malicious actors, whether individuals or states. Over time, the focus has expanded—first to securing the infrastructure of increasingly networked systems and now to addressing vulnerabilities in machine learning (ML) models. Cryptography is one of many tools whose use can enhance computer security. However, the protections afforded by even perfect cryptography can often be circumvented by taking advantage of vulnerabilities in the computer systems on which

FIGURE 3.1 How public-key cryptography works



that cryptography is implemented. For example, if an intercepted encrypted message is too hard to decrypt, the attacker's focus will most likely be on exploiting vulnerabilities in computer security to obtain the message before the sender encrypts it or after the receiver decrypts it.

Thus, cryptography and security are inseparable, and using cryptography is by no means a guarantee of security.

Blockchain Technology

A blockchain is a chain of digital blocks, each containing a transaction and a cryptographic hash of the previous block. This links every block (except the first) to its predecessor. As new transactions occur, new blocks are added, extending the chain.

Blockchains are distributed across thousands of computers, ensuring they are highly decentralized. They enable multiple parties to coordinate transactions without a central trusted authority—a common need in financial settings. Transactions recorded on them cannot be altered retroactively without detection. Because blockchains are widely distributed across thousands of computers, they are always accessible: Anyone can deploy or interact with applications, and no one can block access to them. Data on blockchains cannot be erased; later transactions may correct errors, but the original record remains.

The distributed nature of blockchains also increases security. A new transaction on a blockchain is broadcast to every party in the network, each of which has a replica of the entire blockchain (see figure 3.2). Each party then tries to validate the new transaction. These replicas may not be fully synchronized; some might have received the new transaction, while others may have not. To ensure that all replicas are identical, blockchains use consensus mechanisms to agree on the correct information. Ethereum, for example, accepts transactions that have been validated by two-thirds of the participants. Blockchains are designed with economic incentives for replicas to behave honestly.

Applications that run on a blockchain are called smart contracts—computer programs that are always available and whose execution cannot be reversed. They can implement financial instruments, record ownership of digital assets, or support marketplaces for buying and selling. Smart contracts are also composable: One contract can use another, enabling a vibrant ecosystem where projects build on each other. Once deployed, the contracts remain available indefinitely. This is in contrast to cloud applications, which disappear when developers stop paying hosting fees.


Key Developments

A Host of Blockchain Applications

Blockchain technology was developed decades ago but has recently been used for a variety of applications. Many of these are operational today, though they are often at limited scale. (For a more comprehensive discussion of examples, see the chapter on cryptography in the *2025 Stanford Emerging Technology Review [SETR]*.) Some current examples include the following:

- **Time stamping and data provenance** Because data written to a blockchain cannot be modified or removed, blockchains provide a secure mechanism for data provenance and time stamping. For instance, creators can post cryptographic hashes of work to a blockchain to establish authorship or creation dates.
- **Identity management** Blockchains enable secure storage and selective disclosure of personal records (e.g., diplomas, birth certificates, financial records), allowing users to prove facts—such as being of a certain age—without revealing sensitive details (such as their actual age). One such application, SpruceID, is already being deployed.³
- **Supply chain management** Blockchains provide a transparent and secure way to track the

FIGURE 3.2 How a blockchain manages transactions



movement of goods and their origin and quantity in industries ranging from luxury goods to food labeling.

- **Transactional records** Storing contracts or sales records on blockchains can reduce fraud, simplify auditing, and streamline operations.
- **Cryptocurrencies** These are digital instruments that many people use as a medium of exchange. Well-known ones include Bitcoin, Ethereum, Avalanche, and Polygon, each of which has its own unique features and applications. Because they are not issued by any central authority, they

are not subject to the same national regulatory regimes that govern traditional currencies (i.e., so-called fiat currencies).

Cryptocurrencies use a blockchain structure to ensure the integrity and immutability of transaction data, making it resistant to fraud and counterfeiting and reducing its susceptibility to government interference or manipulation. Contrary to a common belief, cryptocurrencies can, but do not have to, support private or secret transactions. Indeed, the most popular cryptocurrencies deliberately do not hide the details of their transactions. Those who transact in cryptocurrencies often wish to exchange

their instruments for fiat currency (e.g., real dollars) and generally use a cryptocurrency exchange to do so.

Secure Computation

The field of cryptography has also expanded in scope to include secure computation, which enables multiple parties to jointly compute functions where the inputs from each party are kept secret from the others. Secure computation enables data privacy during computation, ensuring that no party learns more information about the other parties' inputs than what can be inferred from the result alone. It also allows users to prove they possess knowledge of a statement without having to disclose the actual content of that statement. (For a more detailed explanation of secure computation with illustrative examples and explanations, see the chapter on cryptography in *SETR 2025*.) A few representative applications include the following:

- **Private statistics** Stanford's Prio system lets users contribute data, such as COVID-19 exposure status, to an aggregate total without disclosing individual responses.⁶
- **Financial privacy** Banks can collaborate to detect fraud patterns across institutions without revealing individual customer records.
- **Privacy-preserving auctions** These can determine a winner without exposing losing bids, maintaining fairness while protecting private financial information.

Zero-Knowledge Proofs

Zero-knowledge proofs are cryptographic protocols that allow one party (the prover) to convince another (the verifier) that a statement is true without revealing why it is true. For example, someone can prove they know a password or have enough funds for a purchase without disclosing the password or the amount of money. This privacy-preserving technique

has moved from theory into real-world applications, such as the following:

- **Banking** The cryptocurrency Zcash uses zero-knowledge proofs to let users prove they can afford a transaction without having to reveal their account balance.⁴
- **Provenance for digital images** The Coalition for Content Provenance and Authenticity employs zero-knowledge proofs to ensure that an image was captured by a verified camera and underwent only permitted edits—without trusting the editing software itself.⁵
- **Cooperative tracking and verification of numbers of tactical nuclear warheads** Experimental systems have used zero-knowledge proofs to track changes in warhead status while concealing sensitive military information. Though the use has not yet been adopted in formal treaties, its feasibility in principle has been demonstrated.⁶

A more detailed introduction to zero-knowledge proofs and their use cases is available in the chapter on cryptography in *SETR 2025*.


Over the Horizon

Impact of Cryptography

The applications described above suggest a broad range of possibilities for cryptographically enabled data management services. Whether we will see their widespread deployment depends on complicated decisions about economic feasibility, costs, regulations, and ease of use.

Misaligned incentives can affect how fast innovations are deployed. Some of the applications described above provide significant benefits for the parties whose data can be better protected and kept more private. But existing companies, having built their

FIGURE 3.3 How adversarial noise can fool image classifiers



Adding carefully crafted noise (center) to an image of a cat (left) produces an altered version (right) that looks identical to humans but can cause a model to misclassify it—for example, as guacamole.

Source: Neil Alexander Perry

business models on legacy systems that ingest all their customers' data, have no incentive to change their practices. They are the ones who would have to pay for these privacy-protecting capabilities, yet they would not benefit from their adoption.

Widespread deployment will also require confidence that proposed innovations will work as advertised (i.e., would-be users of these innovations must have confidence in them). But concepts such as secure computation and zero-knowledge proofs are math heavy and counterintuitive to most people. Getting policymakers, consumers, and regulators to place their trust in these applications will be challenging.

Machine Learning Security: Adversarial Risks and Systemic Vulnerabilities

As ML systems move into high-stakes settings—including autonomous vehicles, financial platforms, and healthcare diagnostics—their security under adversarial conditions is becoming a critical concern (for more, see chapter 1, on AI). In an ML system, small, malicious changes to inputs can cause large, unexpected model failures. These brittle responses undermine trust in a system's ability to operate safely in environments where reliability is paramount.

This fragility stems from a core asymmetry. While ML performs well on the inputs that most average

users would give it, it often fails on inputs that are crafted by deliberately malicious adversaries. In other words: ML systems are great for random data, but they often perform poorly when confronted with deliberately crafted adversarial data.

ATTACKS ON ML SYSTEMS

Researchers have identified attacks targeting every stage of the ML pipeline:

- Training-time attacks corrupt models during learning. Carefully altering even a single image in a dataset of ten thousand can lead to persistent misclassifications—such as labeling images of dogs as fish.⁷ These techniques, once confined to research, now appear in the real world. The tool Nightshade deliberately allows artists to corrupt images before posting them online, sabotaging unauthorized AI training on scraped content.⁸
- Inference-time attacks occur after deployment. An attacker may introduce “noise”—tiny, imperceptible modifications to data that cause the model to produce incorrect outputs.⁹ For example, making small alterations to the pixels in a cat image can make the model label the image as guacamole (figure 3.3). The key point is that such attacks—repeatedly shown to be possible over the past several years—demonstrate that a

model can sometimes be tricked into errors that would be obvious to humans. Imagine the risk if an adversary could make a military reconnaissance system mistake tanks for school buses or induce an airport security scanner to mistake a gun for a notebook.

These threats extend beyond images. For example, because large language models (LLMs) cannot distinguish between inputs intended as data versus inputs intended as commands, they may interpret a phrase embedded in data as a command, which may be hostile or malicious rather than benign. For example, this strategy—which is an instance of a well-known hacker technique known as prompt injection—could result in getting LLMs to leak confidential information, ignore safety constraints, or perform unintended actions.¹⁰ As these models become linked to tools like email and payments, such attacks carry more risk.¹¹ As AI agents—autonomous AI software programs that have access to important data or controls—become increasingly popular, there is a growing risk that these could be “tricked” by malicious content on the internet.

EMERGING RISKS AND THE SECURITY GAP

The rapid deployment of AI systems has outpaced available security solutions. Traditional techniques from computer security and common defenses like input filtering offer only partial protection and often shift the vulnerability elsewhere. For example, if inputs are digitally filtered for safety, attackers may instead target the software filters themselves, which are often susceptible to similar exploits. Some researchers are now exploring new defenses for inference-time threats, such as isolation to protect sensitive components and data from being compromised by malicious or untrusted inputs.¹² Another defense involves using stricter control flows that explicitly manage how decisions, loops, branches, and data interactions occur to help ensure a system’s predictable, secure, and reliable operation.¹³

To secure the training process, other efforts focus on hardware safeguards, such as trusted execution

environments—secure zones within a host system that preserve data confidentiality and computational integrity even if the system is compromised. Stanford researchers are developing auditable training pipelines that log each intermediate step in training. This enables users to verify the model’s training process (e.g., to ensure that the data on which it was trained was not compromised in some way) and trace certain security issues back to their origin when problems arise.¹⁴

All of these defenses remain in their early stages. No current approach offers broad protection across all tasks, data types, or adversarial techniques. The field remains in an arms race: New attacks emerge rapidly, while robust, scalable defenses continue to lag behind. In this landscape, any claim to deploy ML to solve a problem should prompt an immediate question: What have you done about adversarial inputs and attacks?

DUAL-USE CAPABILITIES AND MODEL INTEGRITY

LLMs raise classic dual-use concerns. Their ability to identify software vulnerabilities can assist defenders in fixing systems—or can arm attackers to more easily exploit vulnerabilities. Studies show LLM-based agents can already solve many standardized cybersecurity tasks, rivaling novice human hackers.¹⁵ Whether they will ultimately favor offense or defense remains uncertain. Their use in software development can accelerate productivity but also create new vulnerabilities. For example, LLMs often generate insecure or outdated code, especially when they are used by nonexperts who lack awareness of best practices.¹⁶

Even after they are deployed, models remain vulnerable to extraction attacks. These involve adversaries reconstructing similar models through repeated queries of a target model, enabling them to gather training data that the original model uses.¹⁷ This threatens both intellectual property and the safeguards meant to prevent misuse of the target model. This issue is further compounded

Bitcoin mining uses more energy than the Netherlands.

by the phenomenon of transferability, where an attack on one model often works on similar models. Transferability means that attackers don't need internal access to the original model to succeed. It also means that similar models can be constructed to aid in the development of attacks on the original model, regardless of the protections and safeguards embedded in the original.

Policy Issues

Research Infrastructure

Although cryptography is fundamentally a mathematical discipline, it requires both human talent and substantial computing resources to examine the efficiency of new techniques, write computationally expensive software such as zero-knowledge provers, and conduct comprehensive scans of the internet. Progress also relies on interdisciplinary centers that bring together faculty from different fields to share problem sets and understand the potential benefits of cryptographically enabled techniques and approaches.

Research is funded by both the US government and private industry, but funding from the US government is subject to many requirements that increase the difficulty of proposal submission manyfold (as much as by a factor of sixty). Thus, research faculty often prefer arrangements with the private sector, which tend to be much simpler. On the other hand, only the US government is able to fund research that may not pay off for many years (as in the case of quantum computing).

Exceptional Access

Exceptional access regulations would require communications carriers and technology vendors to provide US law enforcement agencies access to encrypted information (both data storage and communications) under specific legal conditions. Opponents of exceptional access argue that implementing this capability inevitably weakens the security afforded by encryption to everyone. Supporters of exceptional access do not debate this technical assessment: It is true that exceptional access, by definition, weakens encryption. However, they argue that even if lower security is the result of implementing exceptional access, that price is worth the benefits to law enforcement.¹⁸

Energy Consumption

Bitcoin, an older cryptocurrency and today the dominant one, consumes an enormous amount of energy; Bitcoin mining uses more energy than the Netherlands.¹⁹ For this reason, newer blockchains—notably Ethereum—are designed to use far less energy; today Ethereum's annual energy use is less than a ten-thousandth of YouTube's annual consumption. But Ethereum's market capitalization is less than half that of Bitcoin, and it remains to be seen whether any less energy-intensive cryptocurrency will displace the latter.

Quantum Computing and Cryptography

Current public-key cryptography is based on the extraordinarily long times—ones comparable to the age of the universe—today's computers require to derive a private key from its public-key counterpart. When realized, quantum computing (discussed

more fully in chapter 7, on quantum technologies) will pose a significant threat to today's public-key algorithms. Experts disagree on how long it will take to build quantum computers that are capable of this, but under the May 2022 National Security Memorandum 10, Promoting US Leadership in Quantum Computing While Mitigating Risks to Vulnerable Cryptographic Systems, the US government has initiated the transition to quantum-resistant public-key algorithms. Many experts in the field expect quantum-resistant algorithms will be widely available by the time quantum computing comes online.

At the intersection of quantum computing and cryptography are two important issues: (1) that support for the transition to a quantum-resistant encryption environment should continue with urgency and focus, and (2) that messages protected by pre-quantum cryptography will be vulnerable in a post-quantum world. If those messages have been saved by adversaries (which is likely in the case of parties like Russia), those bad actors will be able to read a host of old messages. Containing secrets from the past, they may reveal embarrassments and dangers with potentially detrimental policy implications.²⁰

Cryptocurrencies and the Emerging US Policy Approach

While many countries are pursuing central bank digital currencies (CBDCs) to modernize their financial systems, the United States is taking a markedly different path. CBDCs are cryptography-based currencies issued by central banks, with legal tender status and value tied to a nation's traditional currency. They promise fast, low-cost payments with centralized oversight. Advocates cite benefits such as greater financial inclusion, lower cross-border costs, and preserving the dollar's global role—especially as rivals like China advance their own CBDCs. Critics, however, warn of privacy risks, centralized surveillance, and excessive government control. (For a full discussion of CBDCs, see the chapter on cryptography in SETR 2025.)

Departing from earlier policy, the Trump administration has signaled support for privately issued cryptocurrencies over a CBDC. In January 2025, it issued an Executive Order (EO) on Digital Assets that revoked President Biden's 2022 EO 14067, which had promoted CBDC exploration, consumer protection, and anti-illicit finance measures. The new order explicitly prohibits the development or promotion of a US CBDC and establishes an interagency working group to coordinate digital asset policy and regulation.

These policy shifts have occurred amid ongoing debate over how to classify and regulate digital assets. The 2023 collapse of cryptocurrency exchange FTX and the subsequent conviction of its founder Sam Bankman-Fried intensified scrutiny over whether cryptocurrencies should be treated as securities or currency. Despite some legislative progress, many investors, consumers, and entrepreneurs remain uncertain about their regulatory status.

A March 2025 EO established a US Strategic Bitcoin Reserve and digital asset stockpile, using Bitcoin seized from illegal activities and directing agencies to pursue additional budget-neutral acquisitions of the cryptocurrency. Supporters see it as a way to diversify national reserves, hedge against inflation, and promote US leadership in digital asset innovation. Critics point to Bitcoin's price volatility, its limited utility in crises and cybersecurity risks, and the risk of potential conflicts of interest that could undermine public trust, given that the policymakers themselves may have significant cryptocurrency holdings.

In addition to the above activity, the Guiding and Establishing National Innovation for US Stablecoins (GENIUS) Act, signed in July 2025, created a federal framework for issuing payment stablecoins—cryptocurrencies that are designed to have stable prices. By combining the speed and programmability of cryptocurrencies with the familiarity of fiat currency (i.e., ordinary money), stablecoins allow users to transact without the price volatility typically associated with many cryptocurrencies and other digital

assets. However, the law requires stablecoins to be backed by reserve assets such as Treasury bills and precious metals rather than algorithmic mechanisms to enhance their price stability. This may limit innovation via experimentation with alternative cryptocurrency designs.

The Trump administration is also backing the Digital Asset Market Clarity Act,²¹ or the Clarity Act, which proposes making the Commodity Futures Trading Commission the primary regulator of digital commodities and their intermediaries while maintaining certain Securities and Exchange Commission (SEC) powers over initial crypto sales. The act would also introduce a special exemption that eases SEC registration requirements for fundraising purposes. Under this proposed legislation,²² digital commodities would be defined as digital assets whose worth is “intrinsically linked” to their activity on a blockchain. This group includes nearly all cryptocurrencies in use today. However, the definition of the term “digital commodity” would exclude securities, derivatives, and stablecoins, even if they are based on blockchains.

Against this backdrop, cryptocurrencies pose complex and evolving policy challenges. These include the following:

- **Regulatory clarity and market integrity** Regulatory ambiguity and weak oversight continue to be a challenge facing digital asset markets. The decentralized, cross-border nature of cryptocurrencies complicates efforts to classify and supervise them. Inherent volatility of cryptocurrencies, combined with limited transparency across many exchanges, exposes users to fraud, manipulation, and financial risk. As use of these assets grows, so do calls for clearer rules, better disclosures, and stronger consumer protections. Tax reporting remains another challenge: The pseudonymous nature of cryptocurrencies complicates enforcement, and users often lack a clear understanding of their tax-reporting obligations. The GENIUS and Clarity Acts are first steps

toward regulatory clarity, but as the use of cryptocurrencies expands, the surfacing of other issues requiring further legislative and executive branch attention is inevitable.

- **Financial crime and illicit activities** The pseudonymous nature of cryptocurrencies and their cross-border use also enables or facilitates money laundering, tax evasion, and sanctions evasion, creating major enforcement challenges. Authorities are expanding international cooperation, tightening anti-money laundering and know-your-customer protocols related to cryptocurrencies, and working to close regulatory gaps.
- **Economic and monetary policy risks** Because cryptocurrencies bypass traditional financial systems, their widespread use weakens central banks’ ability to control the money supply and set interest rates across the economy. If cryptocurrencies are integrated into mainstream finance (e.g., through retirement funds, banking systems, or national reserves), a collapse in cryptocurrency valuations could trigger a financial crisis, impacting savings and investments across the economy. Additionally, as more people use cryptocurrencies instead of fiat money, confidence in government-issued currencies may erode.
- **Conflicts of interest and governance transparency** Government actions can significantly affect cryptocurrency prices (as is true of any other investment asset), raising concerns about personal financial gain among policymakers and regulators. Industry influence over the regulatory process also prompts political and ethical scrutiny.
- **Cybersecurity risks** The decentralized architecture of cryptocurrencies creates novel opportunities for cyberattacks across digital wallets, crypto exchanges, and smart contracts. Hacks, phishing, and other exploits may cause substantial losses and erode trust. As crypto assets intertwine with traditional finance, their vulnerabilities may trigger broader economic fallout.

Strong cybersecurity standards, incident reporting, and federal coordination are essential to limit systemic risk to the national and global financial system.

- **Privacy and surveillance concerns** Digital asset regulation increasingly intersects with debates over financial privacy and civil liberties. Expanding anti-money laundering and know-your-customer rules may lead to calls for digital identity systems, raising concerns about surveillance and state overreach. At the same time, technologies that enhance privacy may face increased scrutiny. Policymakers will need to carefully balance law enforcement needs with privacy concerns.

NOTES

1. "Cryptography," National Institute of Standards and Technology, US Department of Commerce, n.d., <https://www.nist.gov/cryptography>.
2. In this context, encrypting the hash value simply means running the encryption algorithm using a string of numbers that just happen to be Alice's private key as the input. In most cases involving public-key cryptography, the private key is used only for decryption purposes, but nothing stops a user from using it in other ways.
3. "SpruceID," Spruce Systems, accessed October 13, 2024, <https://spruceid.com/>.
4. "What Are Zero-Knowledge Proofs?," Zcash, accessed September 7, 2025, <https://z.cash/learn/what-are-zero-knowledge-proofs>.
5. Trisha Datta and Dan Boneh, "Using ZK Proofs to Fight Disinformation," Medium, September 29, 2022, <https://medium.com/@boneh/using-zk-proofs-to-fight-disinformation-17e7d57fe52f>.
6. Miles A. Pomper, William Alberque, Marshall L. Brown Jr., William M. Moon, and Nikolai Sokov, "Everything Counts: Building a Control Regime for Nonstrategic Nuclear Warheads in Europe," 2022, CNS Occasional Paper #55, Monterey, CA: James Martin Center for Nonproliferation Studies, <https://nonproliferation.org/op55-everything-counts-building-a-control-regime-for-nonstrategic-nuclear-warheads-in-europe>.
7. Pang Wei Koh and Percy Liang, "Understanding Black-Box Predictions via Influence Functions," *Proceedings of the 34th International Conference on Machine Learning*, in *Proceedings of Machine Learning Research* 70 (2017): 1885–94, <https://proceedings.mlr.press/v70/koh17a.html>.
8. Shawn Shan, Wenxin Ding, Josephine Passananti, Stanley Wu, Haitao Zheng, and Ben Y. Zhao, "Nightshade: Prompt-Specific Poisoning Attacks on Text-to-Image Generative Models," *2024 IEEE Symposium on Security and Privacy*, 2024, 807–25, <https://doi.org/10.1109/SP54263.2024.00207>.
9. Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy, "Explaining and Harnessing Adversarial Examples," preprint, arXiv, December 20, 2014, <https://arxiv.org/abs/1412.6572>; Anish Athalye, Logan Engstrom, Andrew Ilyas, and Kevin Kwok, "Synthesizing Robust Adversarial Examples," *Proceedings of the 35th International Conference on Machine Learning*, 2018, 284–93, <https://proceedings.mlr.press/v80/athalye18b.html>.
10. Hezekiah J. Branch, Jonathan Rodriguez Cefalu, Jeremy McHugh, Leyla Hujer, Aditya Bahl, Daniel del Castillo Iglesias, Ron Heichman, and Ramesh Darwishi, "Evaluating the Susceptibility of Pre-Trained Language Models via Handcrafted Adversarial Examples," preprint, arXiv, September 5, 2022, <https://arxiv.org/abs/2209.02128>.
11. Kai Greshake, Sahar Abdelnabi, Shailesh Mishra, Christoph Endres, Thorsten Holz, and Mario Fritz, "Not What You've Signed Up for: Compromising Real-World LLM-Integrated Applications with Indirect Prompt Injection," *Proceedings of the 16th ACM Workshop on Artificial Intelligence and Security*, November 26, 2023, 79–90.
12. Luca Beurer-Kellner, Beat Buesser, Ana-Maria Crețu, Edoardo Debenedetti, Daniel Dobos, Daniel Fabian, Marc Fischer, et al., "Design Patterns for Securing LLM Agents Against Prompt Injections," preprint, arXiv, last updated June 27, 2025, <https://arxiv.org/abs/2506.08837>.
13. Edoardo Debenedetti, Iliia Shumailov, Tianqi Fan, Jamie Hayes, Nicholas Carlini, Daniel Fabian, Christoph Kern, Chongyang Shi, Andreas Terzis, and Florian Tramèr, "Defeating Prompt Injections by Design," preprint, arXiv, last updated June 24, 2025, <https://arxiv.org/abs/2503.18813>.
14. Megha Srivastava, Simran Arora, and Dan Boneh, "Optimistic Verifiable Training by Controlling Hardware Nondeterminism," *Proceedings of the 38th International Conference on Neural Information Processing Systems*, 2024 (Curran Associates, Inc., 2025), 95639–61, <https://dl.acm.org/doi/10.5555/3737916.3740946>.
15. Andy K. Zhang, Neil Perry, Riya Dulepet, Joey Ji, Celeste Menders, Justin W. Lin, Eliot Jones, et al., "Cybench: A Framework for Evaluating Cybersecurity Capabilities and Risks of Language Models," preprint, arXiv, 2024, <https://arxiv.org/abs/2408.08926>.
16. Neil Perry, Megha Srivastava, Deepak Kumar, and Dan Boneh, "Do Users Write More Insecure Code with AI Assistants?," *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*, 2023 (Association for Computing Machinery, 2023), 2785–99, <https://doi.org/10.1145/3576915.3623157>.
17. Florian Tramèr, Fan Zhang, Ari Juels, Michael K. Reiter, and Thomas Ristenpart, "Stealing Machine Learning Models via Prediction APIs," *25th USENIX Security Symposium*, 2016, preprint, arXiv, 2016, 601–18, <https://arxiv.org/abs/1609.02943>.
18. "Attorney General William P. Barr Delivers Keynote Address at the International Conference on Cybersecurity," Office of Public Affairs, US Department of Justice, July 23, 2019, <https://www.justice.gov/opa/speech/attorney-general-william-p-barr-delivers-keynote-address-international-conference-cyber>.
19. "Bitcoin Energy Consumption Index," Digiconomist, accessed September 7, 2025, <https://digiconomist.net/bitcoin-energy-consumption>.
20. Herbert Lin, "A Retrospective Post-Quantum Policy Problem," *Lawfare*, 2022, <https://www.lawfaremedia.org/article/retrospective-post-quantum-policy-problem>.
21. Although the vast majority of cryptocurrencies are indeed based on blockchain, it is possible in principle to build a cryptocurrency on a different underlying technology.
22. Paul Tierno, *Crypto Legislation: An Overview of H.R. 3633, the CLARITY Act*, Congressional Research Service, Library of Congress, July 7, 2025, <https://www.congress.gov/crs-product/IN12583>.

STANFORD EXPERT CONTRIBUTORS

Dr. Dan Boneh

SETR Faculty Council and Professor of Computer Science and of Electrical Engineering

Dr. David Tse

Thomas Kailath and Guanghan Xu Professor of Engineering

Neil Perry

SETR Fellow and PhD Student in Computer Science

Copyright © 2026 by the Board of Trustees of the Leland Stanford Junior University

This publication reflects updates through December 2025

32 31 30 29 28 27 26 7 6 5 4 3 2 1

Designer: Howie Severson

Typesetter: Maureen Forys

Image credits: Linda A. Cicero/Stanford News and iStock.com/PTC-KICKCAT92 (cover); iStock.com/mofuku (p. 22); iStock.com/wacomka (p. 38); iStock.com/FeelPic (p. 56); iStock.com/JONGHO SHIN (p. 70); iStock.com/Chartchai San-saneeyashewin (p. 88); iStock.com/ArtemisDiana (p. 102); iStock.com/PhonlamaiPhoto (p. 116); iStock.com /imaginima (p. 142); iStock.com/Floriana (p. 156); iStock.com/dima_zel (p. 170); Tim Griffith (p. 225)