Journal for Social Science Archives

Online ISSN: 3006-3310
Print ISSN: 3006-3302
Volume 4, Number 1, 2026, Pages 2730 – 2739
Journal Home Page
https://jssarchives.com/index.php/Journal/about

Ethical and Governance Dimensions of Artificial Intelligence in Cryptocurrency Crime Prevention: Balancing Privacy, Surveillance, and Civil Liberties

Akbar Ali Datoo[1], Dr. Naima Saeed[2], Syed Razi Hasnain[3] & Waleed Ahmed Khan[4]

[1]*M.Phil. in Criminology, AI, Digital Criminology, Intelligence & Security, Department of Criminology, University of Karachi, Email: akbarali.datoo@gmail.com*
[2]*Chairperson, Department of Criminology, University of Karachi, Email: nsaeed@uok.edu.pk*
[3]*Master in Criminology, Visiting lecturer, Department of Criminology, University of Karachi, Email: syedrazishah2001@gmail.com*
[4]*M.Phil. Scholar, Teaching Assistant, Department of Criminology, University of Karachi, Email: waleedahmedkhanofficial@gmail.com*

ARTICLE INFO

Article History:

Received:	*February*	*15, 2026*
Revised:	*March*	*06, 2026*
Accepted:	*March*	*20, 2026*
Available Online:	*March*	*30, 2026*

Keywords:

Artificial Intelligence Ethics, Cryptocurrency Surveillance, Privacy Rights, Algorithmic Bias, Civil Liberties, Blockchain Forensics, Predictive Policing, Surveillance Governance, Financial Privacy

Corresponding Author:
Akbar Ali Datoo
Email:
akbarali.datoo@gmail.com

ABSTRACT

The deployment of artificial intelligence for cryptocurrency crime prevention has generated significant ethical tensions between security imperatives and individual rights. This article examines the ethical, legal, and social implications of AI-based surveillance systems in cryptocurrency ecosystems, Analyzing how these technologies implicate privacy rights, algorithmic bias, transparency obligations, and civil liberties protections. Drawing upon case studies including the Tornado Cash sanctions controversy, the application of blockchain forensics in criminal investigations, and the emerging use of AI for predictive risk assessment, the research identifies fundamental tensions between the pseudonymous ideals of cryptocurrency networks and the transparency requirements of effective law enforcement. The analysis reveals that AI-based monitoring systems risk creating surveillance infrastructures that extend far beyond legitimate crime prevention objectives, potentially chilling lawful financial activity and undermining the privacy protections that attract many users to cryptocurrency. Algorithmic bias represents a particular concern, as machine learning models trained on historical transaction data may perpetuate or amplify existing patterns of discriminatory enforcement, disproportionately affecting users from marginalized communities or regions with legitimate risk factors. The lack of transparency in many AI systems complicates accountability, as individuals whose transactions are flagged or accounts frozen may have no meaningful opportunity to understand or contest automated decisions. The article proposes a governance framework for responsible AI deployment in crypto crime prevention, incorporating privacy-by-design principles, algorithmic impact assessments, meaningful transparency and appeal mechanisms, and democratic oversight of surveillance systems. The analysis concludes that while AI offers powerful capabilities for crime prevention, realizing these benefits without undermining fundamental rights requires deliberate attention to ethical design, regulatory safeguards, and ongoing stakeholder engagement.

1. Introduction

The rapid integration of artificial intelligence into cryptocurrency crime prevention systems has generated profound ethical questions about the appropriate balance between security and liberty in digital financial networks. On one hand, the scale and sophistication of cryptocurrency-enabled crime, including money laundering, ransomware attacks, and sanctions evasion, demand robust preventive capabilities that only AI can provide. On the other hand, the surveillance infrastructures required to detect and prevent these crimes risk encroaching on fundamental privacy rights, perpetuating algorithmic bias, and operating without adequate transparency or accountability mechanisms (Ferguson, 2017).

These tensions are not merely theoretical. The OFAC sanctions against Tornado Cash, a decentralized cryptocurrency mixer, raised fundamental questions about whether privacy-enhancing software code can legitimately be sanctioned and whether such actions represent appropriate enforcement or overreach that criminalizes innovation. Law enforcement applications of blockchain forensics have enabled remarkable successes in tracing stolen funds and identifying criminal actors, but they have also created surveillance capabilities that could be applied to lawful financial activity. Predictive AI systems that assess transaction risk profiles may embed biases that disproportionately flag activity from certain geographic regions or demographic groups, creating discriminatory enforcement patterns (O'Neil, 2016).

The pseudonymous nature of cryptocurrency networks creates particular challenges for balancing security and privacy. Unlike traditional financial systems, which are built around identified customer relationships, cryptocurrency protocols were designed to enable peer-to-peer transactions without intermediaries or identity verification. While this design offers important privacy benefits and financial inclusion opportunities, it also creates challenges for AML compliance and law enforcement investigation. AI systems that attempt to de-anonymize blockchain transactions by linking addresses to real-world identities may undermine the privacy protections that attract legitimate users to cryptocurrency (Narayanan et al., 2016).

This article addresses several ethical and governance questions: What privacy rights should cryptocurrency users reasonably expect, and how do AI-based surveillance systems implicate those rights? How can algorithmic bias in transaction monitoring systems be identified and mitigated? What transparency and accountability mechanisms should govern AI-based crime prevention systems? How can democratic oversight ensure that surveillance capabilities are deployed appropriately and not expanded beyond legitimate crime prevention objectives? And finally, what governance frameworks can balance security imperatives with civil liberties protections?

The theoretical framework for this analysis draws upon surveillance studies literature, particularly the concept of the surveillance society and the notion that data collection systems tend to expand beyond their original purposes (Lyon, 2001). This framework suggests that AI-based transaction monitoring systems, once established for crime prevention, may be repurposed for broader surveillance objectives, creating chilling effects on financial activity and political expression. The analysis also incorporates procedural justice theory, which emphasizes the importance of fair processes, transparency, and meaningful voice in interactions with legal and regulatory systems (Tyler, 2006). AI systems that flag transactions or freeze accounts without adequate explanation or appeal mechanisms violate procedural justice principles, even if they accurately identify illicit activity.

Methodologically, this article employs normative ethical analysis, examining the moral principles that should govern AI deployment in cryptocurrency crime prevention. The analysis draws upon

case studies of actual AI applications, including the Tornado Cash sanctions, blockchain forensics in criminal investigations, and predictive risk assessment systems, to illustrate ethical tensions and identify governance gaps. The article does not provide empirical evaluation of AI system performance but rather examines the ethical frameworks that should guide their design and deployment.

2. Privacy Rights and Surveillance Tensions

2.1 The Privacy Promise of Cryptocurrency

Cryptocurrency networks were designed with privacy considerations embedded in their core architectures. Bitcoin, the first successful cryptocurrency, offers pseudonymity: transactions are recorded on a public ledger, but wallet addresses are not directly linked to real-world identities. Users can generate unlimited addresses, creating the possibility of unlinkable transaction histories. While Bitcoin's pseudonymity has been characterized as weak privacy by modern standards, it represented a significant departure from traditional financial systems, where transactions are directly linked to identified accounts (Nakamoto, 2008).

Subsequent cryptocurrency projects have enhanced privacy protections. Monero uses ring signatures, stealth addresses, and confidential transactions to obscure sender identities, recipient addresses, and transaction amounts. Zcash offers shielded transactions that encrypt transaction details, revealing only that a valid transaction occurred. Privacy wallets and coinjoin implementations enable Bitcoin users to mix their transactions with others, breaking linkability. These privacy features attract users who value financial confidentiality, including political dissidents, whistleblowers, journalists, and ordinary citizens who object to comprehensive financial surveillance (Möser et al., 2018).

However, privacy features also attract criminal actors who seek to obscure illicit activity. This dual-use characteristic creates regulatory tension: measures that restrict privacy to combat crime also burden legitimate users, while privacy protections that benefit legitimate users also facilitate crime. Navigating this tension requires careful ethical analysis that recognizes the legitimate interests on both sides.

2.2 AI Surveillance Capabilities and Privacy Intrusions

AI-based blockchain forensics has substantially improved the ability to link cryptocurrency addresses to real-world identities. Machine learning algorithms can cluster addresses controlled by the same entity based on transaction patterns, identify exchange deposit addresses that require identity verification, and trace fund movements through mixing services. These capabilities enable law enforcement to de-anonymize cryptocurrency transactions, connecting on-chain activity to specific individuals (Meiklejohn et al., 2013).

The privacy intrusions created by these capabilities are substantial. AI systems can potentially reconstruct individuals' entire transaction histories, revealing financial relationships, spending patterns, and even political donations or other sensitive activities. This surveillance capability extends far beyond the specific transactions that initially attracted law enforcement attention, creating the possibility of comprehensive financial monitoring. Once address clusters are linked to identities, all associated transactions become visible, including those predating the investigation and those unrelated to any suspected crime.

Privacy risks are compounded by the permanence and public availability of blockchain data. Unlike traditional financial records, which may be deleted or sealed, blockchain transactions are

permanent and publicly accessible. Once an address is linked to an identity, that linkage persists indefinitely, and all historical transactions become retroactively de-anonymized. This permanence means that privacy compromises cannot be remedied; once lost, financial privacy cannot be recovered.

2.3 The Chilling Effects of Surveillance

Comprehensive financial surveillance creates chilling effects on lawful activity. Individuals who know their transactions may be monitored may self-censor lawful but controversial activities, including political donations, purchases of sensitive books or media, or transactions with disfavored but legal organizations. These chilling effects undermine democratic participation and individual autonomy, even in the absence of any enforcement action (Solove, 2004).

For cryptocurrency specifically, surveillance may discourage legitimate use by privacy-conscious individuals, driving them away from transparent blockchains and toward privacy-preserving alternatives. If AI-based surveillance effectively de-anonymizes Bitcoin and Ethereum transactions, users who value privacy may migrate to Monero or other privacy-focused cryptocurrencies. This migration could paradoxically reduce transparency, as criminal actors would also adopt more robust privacy protections, making law enforcement more difficult rather than less.

The chilling effects of surveillance are difficult to measure but potentially significant. Surveys of cryptocurrency users consistently identify privacy as a primary motivation for adoption, alongside investment returns and transactional efficiency. If AI surveillance substantially reduces the practical privacy of cryptocurrency transactions, it may undermine a core value proposition of the technology, reducing adoption and innovation.

3. Algorithmic Bias and Discrimination

3.1 Sources of Bias in AI Crime Prevention Systems

AI systems deployed for cryptocurrency crime prevention may perpetuate or amplify existing biases through several mechanisms. First, training data may reflect historical patterns of enforcement that themselves embed bias. If law enforcement has historically focused on certain geographic regions, transaction types, or demographic groups, machine learning models trained on this data will learn to predict crime in those contexts, potentially ignoring or under-detecting crime elsewhere (Barocas & Selbst, 2016).

Second, the features used to predict illicit activity may serve as proxies for protected characteristics. Transaction patterns that correlate with geographic location, economic status, or other demographic factors may lead to differential treatment even if the models do not explicitly include these characteristics. For example, remittance transactions from certain regions may share patterns with money laundering activity, leading AI systems to flag legitimate remittances at higher rates than other transaction types.

Third, the design choices made in developing AI systems, including which transactions to label as suspicious, what thresholds to set for alerts, and how to balance false positive and false negative rates, embed value judgments that may systematically favor or burden certain groups. Without deliberate attention to equity considerations, these design choices may produce discriminatory outcomes.

3.2 Evidence of Bias in Financial Surveillance

Research on traditional financial surveillance systems provides evidence of bias that may carry over to cryptocurrency applications. Studies of Suspicious Activity Report filings have found that banks file SARs at higher rates for transactions from minority communities and for cross-border transfers from certain regions, reflecting both legitimate risk factors and implicit bias (FATF, 2019). These patterns may be reinforced by AI systems trained on historical SAR data.

Algorithmic credit scoring systems have been shown to produce disparate outcomes across racial and ethnic groups, even when race is not explicitly included as a variable. Proxy variables, including geographic location and purchasing patterns, correlate sufficiently with race to produce discriminatory effects. Similar dynamics may occur in AI-based transaction monitoring, where location data, counterparty relationships, and transaction patterns may produce disparate flagging rates across groups.

The cryptocurrency context may introduce additional sources of bias. Users from countries with less developed financial infrastructure may have transaction patterns that resemble money laundering activity, simply because their financial systems operate differently. AI systems trained primarily on data from developed economies may misclassify legitimate transactions from emerging markets as suspicious, creating barriers to financial inclusion.

3.3 Mitigating Algorithmic Bias

Addressing algorithmic bias in cryptocurrency crime prevention requires multiple strategies. First, training data should be carefully audited for bias, with efforts to ensure representation across geographic regions, transaction types, and user populations. Synthetic data generation and data augmentation techniques can help address underrepresentation of certain groups.

Second, fairness constraints should be incorporated into model development, requiring that prediction errors be balanced across groups. Techniques including demographic parity, equalized odds, and individual fairness provide mathematical frameworks for operationalizing fairness, though each approach has limitations and trade-offs (Mitchell et al., 2021).

Third, ongoing monitoring for disparate impact is essential. AI systems should be regularly audited to assess whether flag rates, false positive rates, or other outcomes differ systematically across protected groups. When disparities are identified, root cause analysis should determine whether they reflect legitimate risk differences or problematic bias.

Fourth, transparency about model design, training data, and performance metrics enables external scrutiny and accountability. Publishing model cards that document intended use cases, performance characteristics, and known limitations helps users and regulators assess potential bias concerns.

4. Transparency, Accountability, and Procedural Justice

4.1 The Black Box Problem

Many AI systems deployed for cryptocurrency crime prevention operate as "black boxes," producing outputs without revealing the reasoning that led to those outputs. Deep learning models, in particular, involve millions of parameters and nonlinear transformations that resist human interpretation. While these models may achieve high predictive accuracy, their opacity creates accountability challenges (Pasquale, 2015).

When an AI system flags a transaction as suspicious or recommends freezing an account, affected individuals typically receive no explanation for the decision. They cannot know what transaction patterns triggered the flag, what weighting was given to different factors, or whether the system correctly applied relevant legal standards. This opacity violates procedural justice principles that require individuals to understand the bases for adverse decisions and to have meaningful opportunity to contest them.

The black box problem is particularly acute in cryptocurrency contexts because transactions cross jurisdictional boundaries and may involve complex relationships. A user whose transaction is flagged may have no relationship with the platform or law enforcement agency that initiated the flag, making it difficult to seek clarification or correction. The global nature of cryptocurrency networks compounds accountability gaps created by AI opacity.

4.2 Due Process and Appeal Mechanisms

Procedural justice requires that individuals subjected to adverse decisions have meaningful opportunity to contest those decisions before neutral decision-makers. For AI-based transaction monitoring systems, this implies several requirements. First, individuals should receive timely notice when their transactions are flagged or accounts affected. Second, notices should include sufficient explanation of the basis for the decision to enable meaningful challenge. Third, individuals should have access to a prompt appeal process before an impartial human decision-maker. Fourth, adverse decisions should be subject to periodic review to ensure continued accuracy and fairness (Citron, 2007).

Current systems fall far short of these requirements. When exchanges freeze accounts based on AI-generated flags, affected users often receive generic notices citing "suspicious activity" without specific explanation. Appeal processes, where they exist, are often opaque and slow, with users lacking meaningful opportunity to present evidence or challenge the system's determination. The automated nature of many interventions means that no human has reviewed the decision before it takes effect.

Regulatory guidance has begun to address these gaps. The Financial Action Task Force has recommended that countries ensure individuals subject to AML measures have access to effective appeal mechanisms. However, implementation remains uneven, and the specific challenges posed by AI-based monitoring have received limited attention. Developing standards for explainable AI in financial surveillance is an urgent priority.

4.3 The Role of Human Oversight

Meaningful human oversight of AI-based crime prevention systems is essential for accountability, though the form of oversight requires careful design. Simply inserting a human into the decision loop without addressing information asymmetries or cognitive biases may not improve outcomes. Humans reviewing AI recommendations may defer inappropriately to algorithmic outputs or may lack the information needed to assess their accuracy.

Effective human oversight requires several elements. First, human reviewers should have access to the same information as the AI system, plus any additional context needed to assess the decision. Second, reviewers should receive training on the AI system's capabilities and limitations, including known bias concerns and error patterns. Third, review processes should be structured to encourage independent judgment, avoiding the anchoring effects that cause humans to defer to algorithmic recommendations. Fourth, review decisions should be documented, enabling retrospective analysis of human override patterns and system performance.

In high-stakes contexts, including account freezes and law enforcement referrals, human review should be required before action rather than after. This pre-decision review ensures that automated systems do not impose irreversible harms before accountability mechanisms engage. While pre-decision review reduces efficiency gains from automation, the due process values at stake justify this trade-off in many contexts.

5. Governance Frameworks for Responsible AI Deployment

5.1 Privacy-by-Design Principles

Privacy-by-design approaches embed privacy protections into system architectures rather than adding them as afterthoughts. For AI-based cryptocurrency crime prevention, privacy-by-design implies several principles. Data minimization requires collecting only the transaction data necessary for legitimate crime prevention purposes, not comprehensive surveillance data that could enable broader monitoring. Purpose limitation restricts data use to specified crime prevention functions, prohibiting repurposing for unrelated surveillance. Data retention limits ensure that transaction data is deleted or anonymized after reasonable periods, recognizing that indefinite retention creates ongoing privacy risks (Cavoukian, 2009).

Privacy-by-design also implies technical protections that limit surveillance capabilities. Differential privacy techniques can enable aggregate analysis of transaction patterns without revealing individual transactions. Zero-knowledge proofs can verify compliance with AML requirements without exposing transaction details. Federated learning can train detection models across distributed data sources without centralizing sensitive information. These techniques can potentially achieve crime prevention objectives while preserving more privacy than current approaches.

Implementing privacy-by-design requires regulatory requirements that incentivize privacy protection and technical standards that enable interoperable privacy-preserving compliance. The development of such requirements and standards should involve diverse stakeholders, including privacy advocates, industry representatives, and affected communities.

5.2 Algorithmic Impact Assessments

Algorithmic impact assessments provide a structured process for identifying and mitigating the risks of AI systems before deployment. For cryptocurrency crime prevention, AI impact assessments should evaluate potential privacy intrusions, bias concerns, transparency gaps, and civil liberties implications. The assessment process should involve internal review, external expert evaluation, and stakeholder consultation, with results disclosed publicly to enable accountability (Reisman et al., 2018).

Impact assessments should be conducted at multiple stages: during system design, before deployment, after significant updates, and periodically during operation. Each assessment should update the risk evaluation based on observed system performance and emerging concerns. Assessment findings should inform system modifications, with high-risk applications subject to additional oversight or deployment restrictions.

The impact assessment process should include meaningful opportunities for affected communities to provide input. Users whose transactions may be monitored, civil liberties organizations, and privacy advocates should have channels to raise concerns and influence system design. This participatory approach recognizes that those most affected by surveillance systems have valuable perspectives on their risks and appropriate parameters.

5.3 Democratic Oversight and Accountability

Democratic oversight of AI-based crime prevention systems is essential to ensure that surveillance capabilities serve legitimate public purposes and do not expand beyond their intended scope. Several mechanisms can support democratic accountability. Legislative oversight committees should receive regular briefings on surveillance system capabilities, deployment decisions, and effectiveness metrics, with authority to investigate concerns and recommend changes.

Independent inspector general or ombudsperson functions should provide ongoing monitoring of surveillance system operations, investigating complaints from affected individuals and reporting findings publicly. These oversight bodies should have access to system design documents, training data, performance metrics, and other information needed to assess compliance with legal and ethical requirements.

Public reporting requirements should ensure transparency about surveillance system deployment, including the number of accounts flagged or frozen, the demographic characteristics of affected users, the accuracy of predictions, and the outcomes of appeal processes. While some operational details may require confidentiality, aggregate transparency enables public debate about appropriate surveillance parameters.

6. Balancing Security and Liberty

The central ethical challenge of AI-based cryptocurrency crime prevention is balancing legitimate security imperatives against fundamental liberty interests. On one side, the scale of cryptocurrency-enabled crime, including money laundering, ransomware attacks, and sanctions evasion, justifies robust preventive capabilities. Victims of these crimes suffer real harms, and effective prevention requires monitoring and intervention capabilities that intrude on privacy.

On the other side, comprehensive financial surveillance threatens privacy rights, creates chilling effects, perpetuates bias, and operates without adequate transparency or accountability. These concerns are not merely theoretical; they reflect demonstrated patterns of mission creep, discriminatory enforcement, and due process violations in analogous surveillance contexts (Lyon, 2018).

Navigating this tension requires rejecting both extremes. Absolute privacy that prohibits any monitoring or intervention would enable criminal activity to flourish, harming innocent victims and undermining trust in cryptocurrency systems. Absolute surveillance that monitors all transactions for any potential crime would destroy the privacy benefits that attract legitimate users and enable oppressive overreach.

The appropriate balance lies in proportionate, targeted, accountable surveillance. Monitoring should focus on patterns associated with high-risk activity, not comprehensive collection that enables fishing expeditions. Interventions should be calibrated to risk levels, with lower-risk activity receiving lighter-touch treatment. Automated decisions should be subject to meaningful human review and appeal. Transparency about surveillance practices should enable public debate and democratic control. And privacy-preserving techniques should be deployed wherever possible to minimize intrusions while achieving prevention objectives.

7. Conclusion

This article has examined the ethical and governance dimensions of artificial intelligence in cryptocurrency crime prevention, identifying fundamental tensions between security imperatives and civil liberties. AI-based surveillance systems offer powerful capabilities for detecting money

laundering, tracing sanctions evasion, and preventing ransomware attacks. However, these capabilities also create privacy intrusions, bias risks, transparency gaps, and accountability deficits that threaten fundamental rights.

Addressing these ethical challenges requires deliberate attention to system design, governance frameworks, and democratic oversight. Privacy-by-design principles should embed protections into system architectures from the outset. Algorithmic impact assessments should identify and mitigate risks before deployment. Transparency requirements should enable external scrutiny and accountability. Meaningful human oversight and appeal mechanisms should protect due process. And democratic debate should set appropriate parameters for surveillance deployment.

The cryptocurrency community, technology developers, regulators, and civil society organizations all have roles to play in developing governance frameworks that balance security and liberty. Technical innovation can develop privacy-preserving compliance tools that reduce surveillance burdens. Regulatory guidance can establish standards for explainability, fairness, and accountability. Democratic processes can determine appropriate trade-offs between competing values. And ongoing stakeholder engagement can ensure that governance frameworks evolve with technological capabilities and social values.

The future of cryptocurrency crime prevention need not choose between ineffective security and oppressive surveillance. With thoughtful design, robust governance, and democratic accountability, AI can enhance security while respecting fundamental rights. Achieving this vision requires sustained attention to ethical considerations, not as afterthoughts but as central design constraints that shape technology development from the outset.

References

1. Barocas, S., & Selbst, A. D. (2016). Big data's disparate impact. *California Law Review*, 104(3), 671-732.
2. Cavoukian, A. (2009). *Privacy by design: The 7 foundational principles*. Information and Privacy Commissioner of Ontario.
3. Citron, D. K. (2007). Technological due process. *Washington University Law Review*, 85(6), 1249-1313.
4. FATF. (2019). *Guidance on digital identity*. Financial Action Task Force.
5. Ferguson, A. G. (2017). *The rise of big data policing: Surveillance, race, and the future of law enforcement*. New York University Press.
6. Lyon, D. (2001). *Surveillance society: Monitoring everyday life*. Open University Press.
7. Lyon, D. (2018). *The culture of surveillance: Watching as a way of life*. Polity Press.
8. Meiklejohn, S., Pomarole, M., Jordan, G., Levchenko, K., McCoy, D., Voelker, G. M., & Savage, S. (2013). A fistful of Bitcoins: Characterizing payments among men with no names. *Proceedings of the 2013 Conference on Internet Measurement Conference*, 127-140.
9. Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., & Gebru, T. (2021). Model cards for model reporting. *Proceedings of the Conference on Fairness, Accountability, and Transparency*, 220-229.
10. Möser, M., Soska, K., Heilman, E., & Lee, K. (2018). An empirical analysis of traceability in the Monero blockchain. *Proceedings on Privacy Enhancing Technologies*, 2018(3), 143-163.
11. Nakamoto, S. (2008). *Bitcoin: A peer-to-peer electronic cash system*. Bitcoin.org.
12. Narayanan, A., Bonneau, J., Felten, E., Miller, A., & Goldfeder, S. (2016). *Bitcoin and cryptocurrency technologies*. Princeton University Press.

13. O'Neil, C. (2016). *Weapons of math destruction: How big data increases inequality and threatens democracy*. Crown Publishing.
14. Pasquale, F. (2015). *The black box society: The secret algorithms that control money and information*. Harvard University Press.
15. Reisman, D., Schultz, J., Crawford, K., & Whittaker, M. (2018). *Algorithmic impact assessments: A practical framework for public agency accountability*. AI Now Institute.
16. Solove, D. J. (2004). *The digital person: Technology and privacy in the information age*. New York University Press.
17. Tyler, T. R. (2006). *Why people obey the law*. Princeton University Press.
18. U.S. Treasury Department. (2022). *Treasury sanctions Tornado Cash, a virtual currency mixer used by malicious cyber actors*. OFAC Press Release.