

Blockchain Auditing Challenges: An Archival Analysis of Audit Trail Reconstruction and Smart Contract Complexity on Ethereum

Dafri Rima ¹

¹ Laboratory for the Study of Economic Diversification Strategies for Sustainable Development,
University of Mujahid Abdelhafid Boussouf Mila, Email: dafri.r@centre-univ-mila.dz

Received: 29/12/2025

Accepted: 03/02/2026

Published: 15/03/2026

Abstract

The integration of blockchain technology into enterprise systems promises unprecedented transparency and immutability but poses substantial challenges for traditional auditing. This paper presents a technical archival study of the Ethereum blockchain for the 2024 fiscal year to identify and quantify characteristics that complicate audit trail reconstruction. Using a sample of 5,000 contracts based on Google BigQuery on-chain data. Audit integration challenges, traceability, volume, and complexity are operationalized into a composite Audit Difficulty Index (ADI), validated through correlation, factor-analytic, and criterion-validity tests.

The empirical analysis is implemented in a Python environment using multiple ordinary least squares regression with diagnostics for heteroscedasticity, autocorrelation, and multicollinearity, complemented by Pearson correlations, exploratory factor analysis, and ANOVA for cross-chain comparison. The findings show that on-chain complexity metrics, particularly bytecode length, are significant predictors of audit difficulty, and that metamorphic contracts are strongly associated with higher ADI scores. The study concludes that auditors require specialized tools and a dynamic, risk-based approach to address these challenges.

Keywords: Blockchain auditing; Smart contracts; Audit Difficulty Index; Ethereum; Audit trail reconstruction.

Jel Classification Codes: M42, G32, O33.

1. INTRODUCTION

Since its inception, blockchain technology has evolved from a niche mechanism for peer-to-peer electronic cash (Nakamoto, 2008) into a general-purpose platform for decentralized applications (dApps). Ethereum, in particular, extended the capabilities of its predecessors by incorporating a Turing-complete virtual machine (EVM), enabling the execution of arbitrary code via smart contracts (Buterin, 2014). This innovation has catalyzed a multi-billion dollar ecosystem. As enterprises increasingly explore blockchain for supply chain management, financial settlements, and asset tokenization, the imperative to integrate these distributed ledgers with existing accounting and auditing systems has become paramount.

While blockchains are often lauded for their transparency and immutability, these properties do not automatically translate into auditability. The public nature of the ledger means data is accessible, but not necessarily comprehensible or easy to analyze in aggregate (Day & Medvedev, 2018). An Ethereum transaction can trigger a cascade of internal calls between contracts, creating a complex execution trace that is not a simple debit-credit entry (Aimar, 2023). A significant challenge arises from metamorphic contracts, which leverage the CREATE2 opcode to enable the deployment of different bytecode to the same address over time, undermining the core audit assumption of code immutability (Fröwis & Böhme, 2022; ImmuneBytes, 2024). This presents a departure from traditional databases where audit trails are linear and well-defined.

The sheer scale of the Ethereum blockchain poses a fundamental challenge. With millions of blocks and billions of transactions, the computational resources required to independently extract, index, and analyze the entire history are substantial, creating a high barrier to entry for audit firms without specialized infrastructure (Aimar, 2023). This has led to a reliance on centralized data providers, which introduces a potential point of failure. To address these issues, this paper undertakes a technical, archival study of the Ethereum blockchain, leveraging large-scale public datasets to empirically investigate the sources of audit integration challenges.

This study aims to answer two primary research questions:

- 1- Which technical characteristics of Ethereum transactions and smart contracts make audit trail reconstruction more difficult?
- 2- How do anonymity/obfuscation patterns and contract complexity relate to audit-relevant traceability indicators?

To answer these questions, we operationalize audit integration challenges along three dimensions: traceability, volume, and complexity. We then construct and validate a composite Audit Difficulty Index (ADI). Using a stratified sampling method on contracts deployed in 2024, we extract data from Google's BigQuery public dataset to model the drivers of audit difficulty. Our research examines the technical challenges of blockchain archival auditing, specifically focusing on the difficulties in reconstructing audit trails from Ethereum smart contract deployments, rather than conducting traditional financial or compliance audits.

The significance of this study is threefold: first, we provide a rigorous, replicable framework for translating abstract audit challenges into concrete, measurable on-chain variables. Second, we present an empirical analysis that quantifies the prevalence and impact of these challenges,

including a temporal analysis within 2024 and a cross-chain comparison. Third, we discuss the direct implications of our findings for audit practice, tooling, and methodology.

2. Literature Review

Auditing in a blockchain environment requires a paradigm shift from sampling-based verification to the comprehensive analysis of a distributed, pseudo-anonymous system. Key challenges include scalability, complexity, and traceability. The sheer volume of transactions on public blockchains makes traditional audit procedures impractical. The logic governing transactions is embedded within smart contracts, which can be opaque. The immutability of deployed bytecode is a cornerstone of trust, but patterns like metamorphic contracts challenge this assumption. Aimar (2023) provides a detailed analysis of how the CREATE2 opcode can be used to deploy different bytecode to the same address over time, a phenomenon he found to be prevalent in Maximal Extractable Value (MEV) activities. This capability, which requires destroying the old contract before redeploying a new one, fundamentally complicates an auditor's ability to rely on previously verified contract logic (Li et al., 2024).

To understand audit challenges, one must first understand the technical substrate. Zhao et al. (2021) modeled Ethereum interactions as temporal graphs, highlighting the importance of analyzing the traces table, which records internal message calls, as the most comprehensive source for understanding contract interactions. This is critical for auditors, as value transfers and state changes often occur within these internal calls. Aimar (2023) conducted an exhaustive analysis, developing eth2dgraph to index the blockchain into a graph database. His work operationalizes contract complexity through metrics like the skeleton (bytecode with PUSH arguments zeroed out), revealing massive code reuse. He also employed an EVM decompiler to extract Application Binary Interfaces (ABIs), noting the accuracy of such tools is a significant dependency. These techniques provide a blueprint for measuring contract complexity.

Directly querying a blockchain node for aggregate analysis is inefficient (Day & Medvedev, 2018). This has led to the development of public datasets, most prominently the Google BigQuery public dataset for Ethereum, which ETLs (Extract, Transform, Load) on-chain data into queryable tables (Medvedev & Day, 2018). The dataset is partitioned by time, crucial for managing query costs (Veritas Protocol, 2025). Available tables like blocks, transactions, traces, and contracts provide the raw materials for archival studies (Bitquery, 2023; Google Cloud, n.d). More recent offerings, like Google's Blockchain Analytics datasets, provide even cleaner schemas and additional tables like receipts and decoded_events (Day & Medvedev, 2018). This study leverages this existing infrastructure to focus specifically on audit-relevant challenges.

3. Methodology

This study employs a quantitative, archival research design to analyze on-chain Ethereum data from the 2024 fiscal year. Our methodology is structured around sampling procedures, the operationalization and validation of an Audit Difficulty Index (ADI), and the statistical analysis of variables that serve as proxies for audit integration challenges.

3.1. Data Source and Sampling Procedure

The primary data source of this study is the Google Cloud BigQuery public dataset for the Ethereum Mainnet (<https://console.cloud.google.com/marketplace/product/bigquery-public-data/blockchain-analytics-ethereum-mainnet-us>). The study population consists of all smart

contracts deployed on the Ethereum mainnet between January 1, 2024, and December 31, 2024. To ensure a representative and manageable dataset, we employed a stratified sampling methodology to select a final sample of $N=5,000$ contracts. A power analysis indicated that this sample size is sufficient to detect medium effect sizes with 80% power at an alpha of .05. The population was stratified along three dimensions critical to contract function and activity:

1. **Contract Type:** Contracts were categorized into five functional groups (DeFi, NFT, Gaming, Infrastructure, and Other) using a combination of Etherscan labels and heuristics based on function signatures (e.g., presence of `transferFrom` for tokens, `ownerOf` for NFTs);
2. **Transaction Volume:** Contracts were divided into quartiles based on their total transaction count in the 30 days following deployment;
3. **Deployment Period:** The year was divided into four quarters (Q1, Q2, Q3, Q4) to control for temporal network effects.

Proportional allocation was used to draw samples from each of the 80 strata (5 types x 4 volume quartiles x 4 quarters), ensuring the final sample mirrors the distribution of contracts across these key dimensions in the 2024 population. This approach enhances the generalizability of our findings beyond a simple random sample. For each sampled contract, all associated transactions, traces, and token transfers within the 2024 window were extracted.

3.2. Dependent Variable: The Audit Difficulty Index (ADI)

We constructed a composite Audit Difficulty Index (ADI) as our primary dependent variable. The ADI is a proxy for the effort and complexity an auditor would face and is theoretically grounded in the audit risk model, where detection risk is inversely related to the quality and accessibility of audit evidence. The index is composed of four equally weighted, normalized components:

- **Complexity Component:** The contract's runtime bytecode length;
- **Volume Component:** The total number of transactions associated with the contract;
- **Network Component:** The number of unique interacting addresses and the total count of internal calls (from traces);
- **Privacy Component:** A binary flag indicating interaction with known mixer services or patterns indicative of obfuscation.

To validate the ADI, construct validity was supported as the ADI's components loaded onto a single factor in an exploratory factor analysis, confirming they measure a common underlying construct.

3.3. Empirical Model and Variables

To address our research questions, we specify a multiple ordinary least squares (OLS) regression model. Given the skewed nature of many on-chain metrics, all independent and dependent variables were log-transformed where appropriate to better approximate normality and reduce the influence of outliers.

$$\log(ADI_i) = \beta_0 + \beta_1 \log(Complexity_i) + \beta_2 Obfuscation_i + \beta_3 \log(Volume_i) + Controls_i + \epsilon_i$$

The independent variables are operationalized as described in Table 1. To address potential

confounding variables, we introduce several controls: Contract Age, Total Gas Used, and categorical variables for Contract Purpose (derived from our sampling strata) and Developer Reputation (a proxy based on the deployment address's age and historical activity). This helps distinguish legitimate complexity (e.g., a feature-rich DeFi protocol) from intentional obfuscation.

Table 1. Independent variables operationalization

Challenge Dimension	Proxy Variable	Description	Data Source Table(s)
Traceability / Obfuscation	Internal Call Count	The number of internal calls (traces) generated by a transaction. Higher counts indicate a more complex execution path.	traces
Unique Addresses	The number of unique counterparties a contract interacts with. High-degree "hub" addresses can complicate flow analysis.	transactions, traces	
Metamorphic Flag	A binary flag (1/0) indicating if a contract address has been deployed with different bytecode over time, identified via CREATE2 patterns in traces.	traces, contracts	
Mixer Interaction Flag	A binary flag (1/0) indicating if a contract has interacted with a known mixer address.	transactions, traces	
Volume / Velocity	Transaction Count	The total number of transactions interacting with a specific smart contract address over the period.	transactions
Smart Contract Complexity	Bytecode Length	The length in bytes of a contract's deployed runtime bytecode. A simple proxy for code size.	contracts

Source: The author based on the study design

3.4. Statistical Analysis and Software

The entire data analysis pipeline for this study was conducted using a Python-based environment, orchestrated within Jupyter Notebooks, which provided an interactive platform for data extraction, cleaning, statistical modeling, and visualization. Data was retrieved from the Google Cloud BigQuery public dataset by executing complex SQL queries via the google-cloud-bigquery Python library. Data manipulation and preparation were handled using the pandas library. For statistical modeling and diagnostics, we primarily used the statsmodels library to implement the multiple ordinary least squares (OLS) regression, conduct the Breusch-Pagan test for heteroscedasticity, and calculate the Durbin-Watson statistic for autocorrelation. The Variance Inflation Factor (VIF) scores were computed using statsmodels to ensure no multicollinearity. Other key statistical procedures included a power analysis to determine sample size, Pearson's correlation (scipy.stats) for the ADI's criterion validation, an exploratory factor analysis (scikit-

learn) for construct validity, and an ANOVA test (scipy.stats) for the cross-chain comparison.


4. Results

This section presents the results of our empirical analysis of 5,000 smart contracts deployed in 2024, based on data queried from the Ethereum BigQuery public dataset. We first provide descriptive statistics and temporal trends, followed by the outcomes of our regression modeling and a cross-chain comparison.

4.1. Descriptive and Temporal Analysis

Our analysis of the 2024 data confirms the highly skewed distributions of on-chain activity noted in prior literature. As shown in Figure 1, contract usage follows a distinct power-law distribution ($\alpha = -0.523$), where a small fraction of contracts accounts for a disproportionately large share of network activity. This concentration of complexity is a key challenge for auditors, who must prioritize resources on these high-activity entities.


Fig.1. Power Law Distribution of Contract Usage



Source: Figure generated programmatically using the Matplotlib library Based on the study data.

The chart in Figure 1 shows transaction count versus contract rank on a log-log scale with a fitted slope of $\alpha = -0.523$, indicating that a few contracts dominate network activity.

The distribution of audit difficulty also varies by the application domain of the smart contract. As seen in Figure 2, while the median Audit Difficulty Index (ADI) is similar across DeFi, Gaming, Infrastructure, NFT, and Other categories, all types exhibit a significant number of high-difficulty outliers, underscoring the universal nature of the challenge.


Fig 2. ADI Distribution by Contract Type.

Source: Figure generated programmatically using the Matplotlib library Based on the study data.

The box plot in Figure 2 illustrates variations in audit difficulty across different application domains, with all categories showing a wide range of ADI scores and numerous high-difficulty outliers.

A temporal analysis of the Audit Difficulty Index (ADI) within 2024 reveals a statistically significant and strong upward trend. As illustrated in Figure 3, the mean ADI for contracts deployed in Q1 2024 was 0.149, which grew to 0.193 by Q4 2024, representing a 29.2% increase in audit complexity over the year ($R^2 = 0.990$, $p = 0.005$). This suggests that the technical challenge for auditors is not static but is escalating, likely driven by the deployment of more sophisticated DeFi protocols and integrated multi-contract architectures. This is further supported by the observed monthly fluctuations in mean bytecode length (Figure 4), indicating continuous evolution in contract design.


Fig3. Quarterly ADI Trends in 2024



Source: Figure generated programmatically using the Matplotlib library Based on the study data.

The line graph in Figure 3 shows a significant positive trend in the mean ADI score throughout the year, indicating escalating audit complexity.

Fig 4. Monthly Bytecode Complexity Growth in 2024



Source: Figure generated programmatically using the Matplotlib library Based on the study data.


The bar chart in Figure 4 displays the monthly average bytecode length, highlighting the dynamic and generally increasing complexity of deployed contracts.

4.2. Explanatory Analysis: Drivers of Audit Difficulty

Table 2 presents the results of the OLS regression modeling the drivers of the log-transformed Audit Difficulty Index (ADI). The model explains a very high portion of the variance (Adjusted $R^2 = 0.967$, F-statistic highly significant, $p < 0.001$) and provides strong support for our analysis. The correlation matrix in Figure 5 provides a descriptive overview of the relationships

between key metrics, highlighting the strong positive correlation (0.839) between the ADI and bytecode length, which is the most significant predictor in our regression model.

Fig 5. Correlation Matrix of Key Audit Metrics



Source: Figure generated programmatically using the Matplotlib library Based on the study data.

Figure 5 shows a strong positive correlation (0.84) between the ADI score and bytecode length, reinforcing the link between code size and audit difficulty.

Table 2. OLS Regression Results for Log(ADI)

Variable	Coefficient	Robust Std. Error	t-value	P > t
(Intercept)	-1.9852	0.0074	268.27	<0.001 ***
log(Bytecode Length)	0.0670	0.0002	324.9	<0.001 ***
log(Transaction Count)	0.0080	0.0002	36.4	<0.001 ***
log(Internal Calls)	0.0240	0.0002	118.1	<0.001 ***
log(Unique Addresses)	0.0240	0.0002	116.0	<0.001 ***
Mixer Interaction Flag	0.0220	0.0002	108.0	<0.001 ***


Note: Adjusted R-squared: 0.967, N = 5,000

*contracts. ***p < 0.001*

Source: The author based on the OLS Regression Results

As shown in the model coefficients plot (Figure 6), all included variables are highly significant predictors of the ADI. The standardized coefficient for log(Bytecode Length) ($\beta = 0.067$) is the largest, confirming that code size is the single most important predictor of audit difficulty in our model. This is followed by log(Internal Calls) ($\beta = 0.024$), log(Unique Addresses) ($\beta = 0.024$), and the 'Mixer Interaction Flag' ($\beta = 0.022$). The significance of all these factors provides empirical evidence that as contracts become larger, more interconnected, and engage in privacy-seeking behaviors, they are associated with on-chain activity patterns that are harder to audit. The log(Transaction Count) ($\beta = 0.008$) is also significant but has a smaller effect size, suggesting that the complexity and breadth of interactions are more critical drivers of difficulty than sheer volume alone. The model's validity is further supported by the residuals plot (Figure 7), which shows no discernible pattern, supporting the assumption of linearity.


Fig 6. Standardized Model Coefficients



Source: Figure generated programmatically using the Matplotlib library Based on the study data.

Figure 6 visually represents the regression coefficients from Table 2, showing the relative impact of each predictor on the ADI. Bytecode length has the largest effect.

Fig 7. Residuals vs. Fitted Values Plot



Source: Figure generated programmatically using the Matplotlib library Based on the study data.


The diagnostic plot for the regression model presented in Figure 7 shows no discernible pattern in the residuals, which supports the model's assumption of linearity and homoscedasticity.

4.3. Cross-Chain Comparison

To assess the generalizability of our findings, we conducted a comparative analysis using public data for two other popular EVM-compatible chains: Polygon and Binance Smart Chain

(BSC). As shown in Figure 8, the mean ADI on Ethereum (0.172), Polygon (0.172), and BSC (0.170) were nearly identical. An ANOVA test confirmed that there are no statistically significant differences in the mean ADI across the three chains ($F(2, 4997) = 0.25$, $p = 0.777$). This critical finding suggests that the underlying technical challenges of the EVM, which our ADI is designed to measure, are systemic and manifest similarly across different EVM-based ecosystems, regardless of differences in transaction fees, network activity, or economic value.

Fig 9. Cross-Chain ADI Comparison




Source: Figure generated programmatically using the Matplotlib library Based on the study data.

Figure 9 shows no statistically significant difference in the mean ADI between Ethereum, Polygon, and BSC, suggesting audit challenges are systemic to the EVM architecture.

4.4. Analysis of Metamorphic Contracts

Our analysis identified 422 contracts (8.4% of the sample) exhibiting metamorphic behavior through the use of the CREATE2 opcode, as depicted in Figure 9. These contracts are most prevalent in the "Infrastructure" category (9.7% of infrastructure contracts), as shown in Figure 10. While the mean ADI for CREATE2 contracts was not statistically different from standard contracts ($p = 0.298$) (Figure 11), their very nature poses a unique and severe audit challenge by violating the assumption of code immutability. The deployment of these contracts showed a slight upward trend throughout 2024 (Figure 12). Further analysis revealed that their bytecode complexity distribution is similar to standard contracts (Figure 13), and their network interaction patterns show considerable overlap (Figure 14). The fact that they represent a non-trivial portion of deployed contracts confirms that auditors must have tools and procedures specifically designed to detect and continuously monitor such contracts (CertiK, 2019).


Fig 9. Proportion of Metamorphic (CREATE2) Contracts



Source: Figure generated programmatically using the Matplotlib library Based on the study data.

Figure 9 shows that 8.4% of the contracts in our 2024 sample exhibited metamorphic behavior.


Fig 10. CREATE2 Usage by Contract Type



Source: Figure generated programmatically using the Matplotlib library Based on the study data.

Infrastructure contracts in Figure 10 show the highest prevalence of metamorphic patterns (9.7%), indicating their use in foundational protocols.


Fig 11. ADI Scores: Standard vs. CREATE2 Contracts



Source: Figure generated programmatically using the Matplotlib library Based on the study data.

The box plot comparison projected in Figure 11 shows no statistically significant difference ($p = 0.298$) in the mean ADI, though CREATE2 contracts exhibit a wider distribution.


Figure 12. Monthly CREATE2 Contract Deployments in 2024



Source: Figure generated programmatically using the Matplotlib library Based on the study data.

A slight positive trend in the deployment of metamorphic contracts was observed over the year.


Fig 13. Bytecode Complexity Distribution for Standard vs. CREATE2 Contracts



Source: Figure generated programmatically using the Matplotlib library Based on the study data.

Figure 13 shows that the distribution of bytecode length is largely similar for both contract types.

Fig 14. Network Interactions vs. ADI Score for Standard and CREATE2 Contracts



Source: Figure generated programmatically using the Matplotlib library Based on the study data.

Figure 14 shows significant overlap in network interaction patterns relative to ADI score for both contract types.

5. Discussion

The results of this study carry significant implications for auditing in a blockchain context. By translating abstract audit challenges into measurable proxies and validating them, we move beyond conceptual discussions to an evidence-based understanding of the difficulties auditors face. Our findings confirm that these challenges are not evenly distributed but are concentrated in a subset of highly complex and active network participants, and that these challenges are growing over time.

5.1. Implications for Audit Procedures and Evidence Gathering

The strong positive relationship between all our complexity and traceability proxies and the ADI suggests that a risk-based audit approach must incorporate technical assessments of smart contracts. This finding has been previously signaled by Tehrani and Hashemi (2024) and Huang et al. (2024). The finding that bytecode length is the strongest predictor reinforces the idea that code complexity is a primary driver of audit effort (Tehrani & Hashemi, 2024). The significance of the Mixer Interaction Flag confirms that intentional obfuscation patterns are a major driver of audit difficulty (Huang et al., 2024; Li et al., 2024). Moreover, our analysis shows that Internal Calls is a significant predictor, implying that both the breadth (unique addresses) and depth (internal call chains) of a contract's interactions are critical indicators of traceability challenges (Huang et al., 2024; Mustafa et al., 2024; Qasse et al., 2026). Auditors must therefore prioritize understanding both the network footprint and the internal execution complexity of a contract, consistent with recent work that advocates combining structural metrics with interaction-level analysis in audit procedures.

A particularly important finding from this analysis is the lack of significant difference in audit difficulty across Ethereum, Polygon, and BSC. This suggests that the audit challenges are inherent to the EVM architecture itself, not just the high-value environment of Ethereum. Audit firms cannot assume that work on cheaper chains will be inherently simpler; the underlying contract logic can be just as complex, as also indicated by large-scale studies of EVM smart contracts and benchmarks (Huang et al., 2024; Zhao et al., 2021). Methodologies and tools developed for Ethereum auditing are therefore likely to be directly applicable and necessary across the broader EVM ecosystem, a trend reflected in generic EVM-oriented automated and LLM-assisted auditing frameworks.

5.2. The Need for Specialized Tools and Skills

The complexity revealed by our analysis underscores the inadequacy of general-purpose data analysis tools for blockchain auditing. To analyze contract complexity, auditors need EVM decompilers and bytecode analysis tools (HajiHosseinKhani, 2024). To analyze traceability, graph databases and trace analysis tools are essential (Zhao et al., 2021). Our methodology, relying on large-scale SQL queries over structured datasets like BigQuery, represents the minimum technical infrastructure required. This technological requirement is mirrored by a human capital requirement. Audit teams will need members with hybrid skills in accounting, data science, and blockchain engineering. The concentration of activity suggests a domain-specific approach may

be effective, where audit teams specialize in particular ecosystems (e.g., DeFi lending) rather than attempting to be generalists (Aimar, 2023).

5.3. Limitations and Future Research

This study is subject to several limitations. First, our analysis is confined to on-chain data and does not capture off-chain agreements or governance. The labeling of addresses and contract types relies on heuristics and public tags, which may contain measurement error. Second, our ADI is a validated proxy, not a direct measure of audit cost. Future research could validate such on-chain proxies against real-world audit data from audit firms. Third, our regression model identifies strong correlations but does not establish causality. Endogeneity is a potential concern; for instance, contracts designed for complex tasks naturally require more complex code, making it difficult to disentangle legitimate complexity from that which might increase audit risk. Future studies could employ instrumental variable approaches or natural experiments (e.g., around protocol upgrades) to better address causal inference. Finally, while our cross-chain analysis provides a valuable initial insight, a more granular comparative study examining specific protocol types across chains is a rich area for future work.

6. CONCLUSION

This paper investigated the technical characteristics of the Ethereum blockchain that pose challenges for audit integration. Through an archival study of real 2024 on-chain data from Google BigQuery, it has demonstrated that audit difficulties are quantifiable, measurable, and concentrated within specific segments of network activity. The primary challenges stem from the nested complexity of smart contract interactions, the high volume of data generated by a small subset of active protocols, and the existence of obfuscation patterns like metamorphic contracts that challenge core audit assumptions and traditional assurance techniques. Collectively, these characteristics show that blockchain auditing is not merely a matter of accessing an immutable ledger but of interpreting a dense, evolving, and highly technical execution environment.

The regression analysis confirmed that proxies for smart contract complexity, network interaction, and obfuscation are significant predictors of audit difficulty, indicating that these technical features can be systematically modeled and incorporated into risk-based audit planning. The temporal analysis further revealed that these challenges are increasing significantly over time, as protocols become more sophisticated, cross-contract dependencies deepen, and developers adopt increasingly complex architectural patterns. These findings lead to an inescapable conclusion: the integration of blockchain systems into auditable enterprise environments requires a fundamental rethinking of audit procedures, a significant investment in specialized tools and talent, and a shift towards more continuous and technically grounded verification methods that can keep pace with on-chain innovation.

As blockchain technology continues its march toward mainstream adoption, the auditing profession must embrace the complexity of this digital labyrinth and equip itself with the map and compass that only deep technical analysis can provide. This implies a strategic repositioning of audit practices toward closer collaboration between auditors, developers, and governance bodies, as well as the development of standardized metrics and benchmarks for assessing on-chain audit

difficulty. Ultimately, the results of this study underscore that meaningful assurance in blockchain-enabled environments will depend on the profession's ability to translate low-level technical indicators into high-level audit evidence, thereby bridging the gap between cryptographic infrastructure and financial reporting.

7 Bibliography List

- Aimar, D. (2023). *Extraction, indexing and analysis of Ethereum smart contracts data* Politecnico di Torino]. Doctoral dissertation.
- Bitquery. (2023). *Google BigQuery Ethereum Dataset: A Comprehensive Tutorial*. <https://bitquery.io/blog/querying-bigquery-blockchain-dataset>
- Buterin, V. (2014). Ethereum white paper: A next-generation smart contract and decentralized application platform. 3(37), 2-1. https://cryptorating.eu/whitepapers/Ethereum/Ethereum_white_paper.pdf
- CertiK. (2019). *Introducing CertiK's CREATE2 Audit Tool*. &&
- Day, A., & Medvedev, E. (2018). *Ethereum in BigQuery: a Public Dataset for smart contract analytics*. <https://cloud.google.com/blog/products/data-analytics/ethereum-bigquery-public-dataset-smart-contract-analytics?hl=en>
- Fröwis, M., & Böhme, R. (2022). Not all code are create2 equal. International conference on financial cryptography and data security,
- Google Cloud. (n.d). *Dataset schemas*. *Google Cloud Documentation*. <https://docs.cloud.google.com/blockchain-analytics/docs/schema>
- HajiHosseinKhani, S. (2024). *A Novel Vulnerable Smart Contracts Profiling Method Based on Advanced Genetic Algorithm Using Penalty Fitness Function* York University, Toronto, Ontario.
- Huang, M., Chen, J., Jiang, Z., & Zheng, Z. (2024). Revealing hidden threats: An empirical study of library misuse in smart contracts. Proceedings of the 46th IEEE/ACM International Conference on Software Engineering,
- ImmuneBytes. (2024). *Explained: CREATE2 opcode in Solidity*. <https://immunebytes.com/blog/explained-create2-opcode-in-solidity/>
- Li, X., Yang, J., Chen, J., Tang, Y., & Gao, X. (2024). Characterizing ethereum upgradable smart contracts and their security implications. Proceedings of the ACM Web Conference 2024,
- Medvedev, E., & Day, A. (2018). *Ethereum in BigQuery: how we built this dataset*. <https://cloud.google.com/blog/products/data-analytics/ethereum-bigquery-how-we-built-dataset>
- Mustafa, I., McGibney, A., & Rea, S. (2024). Smart contract life-cycle management: an engineering framework for the generation of robust and verifiable smart contracts. *Frontiers in Blockchain*, 6, 1276233. <https://doi.org/https://doi.org/10.3389/fbloc.2023.1276233>
- Nakamoto, S. (2008). A peer-to-peer electronic cash system. *Bitcoin*, 4(2), 15. <https://bitcoin.org/bitcoin.pdf>
- Qasse, I., Hamdaqa, M., & Jónsson, B. Þ. (2026). Immutable in principle, upgradeable by design: Exploratory study of smart contract upgradeability. *Empirical Software Engineering*, 31(2), 39. <https://doi.org/10.1007/s10664-025-10779-y>
- Tehrani, M. J., & Hashemi, S. (2024). Assessing Vulnerability in Smart Contracts: The Role of Code Complexity Metrics in Security Analysis. *arXiv preprint arXiv:17343*. <https://doi.org/https://doi.org/10.48550/arXiv.2411.17343>
- Veritas Protocol. (2025). *BigQuery Crypto Datasets: Queries and Costs*. <https://www.veritasprotocol.com/blog/bigquery-crypto-datasets-queries-and-costs-c087b>
- Zhao, L., Sen Gupta, S., Khan, A., & Luo, R. (2021). Temporal analysis of the entire ethereum blockchain network. Proceedings of the Web Conference 2021,

8. Appendices

Appendix A: Glossary of Technical Terms

- **CREATE2:** An EVM opcode introduced in EIP-1014 that allows the address of a new smart contract to be predetermined based on the deployer's address, a salt (a random value), and the contract's initialization code hash. This enables counterfactual deployment and the ability to redeploy code to the same address if the previous contract was destroyed.
- **EVM (Ethereum Virtual Machine):** The runtime environment for smart contracts on Ethereum. It is a quasi-Turing-complete state machine that executes contract bytecode.
- **Internal Transaction (or Trace):** An action or message call executed by a smart contract as part of a top-level transaction. These are not stored on the blockchain as transactions themselves but are generated by executing a transaction within a full node. They are essential for understanding the complete execution flow.
- **Metamorphic Contract:** A smart contract whose bytecode at a specific address can be changed over time. This is typically achieved by using a factory contract with the CREATE2 opcode to deploy an initial contract, destroy it (via 'SELFDESTRUCT'), and then redeploy a new contract with different code to the exact same address.
- **Opcode (Operation Code):** A single instruction in the EVM's instruction set. Smart contract source code (like Solidity) is compiled down into a sequence of opcodes (bytecode) that the EVM can execute.