Modeling Propagation of Financial Fraud: A Case Study of Cryptocurrency Scam in Social Networks

Naglaa Mostafa
Information System Department, Faculty of Computers & Information, Menoufia University

Hatem Abdelkader
Information System Department, Faculty of Computers & Information, Menoufia University

Asmaa Ali
Information System Department, Faculty of Computers & Information, Menoufia University

ABSTRACT
Information diffusion in social networks enables rapid knowledge sharing but also facilitates the viral spread of misinformation. This duality becomes particularly dangerous in the context of financial schemes. The OneCoin cryptocurrency scam exploited social media platforms and personal networks to spread deceptive narratives about its legitimacy. Utilizing emotional triggers such as the fear of missing out (FOMO), trust in influencers, and fabricated blockchain claims, the scam reached millions globally, exploiting the structure and dynamics of social networks, including echo chambers. Influencer hubs played a crucial role in the speed and scale of this misinformation cascade. This study aims to investigate how misinformation related to financial fraud propagates through social networks. It focuses on the OneCoin case to understand the mechanisms of influence, diffusion patterns, and the role of social structures in the sustainability of misinformation. By analyzing user impact, engagement behavior, and viral spread patterns, the objective is to propose data-driven strategies to detect, contain, and ultimately prevent the future dissemination of fraudulent content. We employed a multi-method analytical approach that combines quantitative and structural techniques. Data was sourced from YouTube and social media posts related to OneCoin and Ruja Ignatova. Metrics, including average views, engagement rates, and influencer activity, were analyzed over time. We integrated network analysis models to identify key propagation nodes and cascades, and applied sentiment and hashtag economic analysis to evaluate the virality of information. Findings reveal that the One Coin misinformation campaign achieved broad reach through early influencer amplification, repeated emotional appeals, and minimal counter-narratives. The average engagement rate was 0.76%, with significant spikes during orchestrated events. These results underscore the urgency of early detection systems grounded in network science.

General Terms
Computer Science, Social Science, and Financial Fraud.

Keywords
Information diffusion, Misinformation, Social Network Analysis, Cryptocurrency Scam

1. INTRODUCTION
Social networks have transformed the way information is shared, enabling the rapid diffusion of both accurate knowledge and malicious misinformation. OneCoin, a fraudulent cryptocurrency scheme, leveraged these dynamics to defraud investors worldwide.

In recent years, information diffusion within online social networks has emerged as a powerful force shaping public opinion, behaviors, and financial decisions. While such diffusion enables beneficial knowledge sharing, it also poses substantial risks when it is exploited to disseminate misinformation. Particularly in the domain of cryptocurrencies, the decentralized nature and lack of regulatory oversight have created an ideal environment for the spread of false narratives and fraudulent schemes.

Misinformation in financial contexts often exploits network dynamics, emotional heuristics, and information asymmetry. Research has shown that misinformation spreads significantly faster, deeper, and more broadly than factual information, particularly when fueled by emotionally resonant content, such as promises of wealth or appeals to urgency [1]. This phenomenon becomes particularly perilous in financial fraud cases, where trust and the speed of information dissemination can determine whether a scam succeeds or fails.

The OneCoin cryptocurrency scheme, orchestrated by Ruja Ignatova and launched in 2014, exemplifies this risk. Marketed as a "Bitcoin killer," OneCoin defrauded over $4 billion from investors across 175 countries through a sophisticated blend of multi-level marketing (MLM), fabricated blockchain claims, and online propaganda [2][3]. Despite lacking a public ledger or verifiable blockchain technology, OneCoin proliferated rapidly, mainly through network-based trust mechanisms and the viral nature of social media amplification.

Key elements of the OneCoin scam's virality include:

1. Influencer amplification: Early adopters and local influencers unknowingly promoted the scheme, serving as hubs in diffusion networks

2. Echo chambers: Online communities silence dissent, reinforcing legitimacy and stifling counterinformation [4].

3. Platform acceleration: YouTube, Facebook, and LinkedIn hosted numerous promotional events, seminars, and testimonials that mimicked authentic investment channels [5].

These dynamics can be analyzed using information diffusion models, such as cascade models, in which nodes adopt beliefs based on their neighbors' choices. Also, Influence propagation models identify key individuals who trigger widespread adoption. Finally, Trust propagation models explain how legitimacy spreads through perceived authority and social proximity. Understanding the OneCoin case offers critical insights into how misinformation can exploit the structural and behavioral properties of social networks and how early detection of diffusion patterns can inform countermeasures to such digital financial crimes.

2. RELATED WORK
In this section, we cover the literature on Graph and Diffusion Models for Fake News and Fraud Detection in Table 1. Pierri

et al. [6]analyze Twitter diffusion network topologies, finding that misleading content spreads more deeply and faster than factual news. They focus only on diffusion patterns without integrating trustworthiness or sentiment dynamics. Phan et al.[7] Apply GNN-based trust and community health metrics to predict the spreaders of false information with an accuracy of over 90%. Primarily, the research focuses on spreaders in conventional news rather than financially motivated scams. For fraud detection in cryptocurrency networks, Wu et al.[8] The study introduces a hybrid GNN and data augmentation approach for Ethereum scam detection, achieving strong accuracy. The problem is that the focus was on blockchain transaction structures, which do not incorporate content or information that cascades on social platforms. Patel et al. [9] examine the transaction patterns of fraudsters using SHAP interpretability to flag suspicious activity. The analysis is limited to on-chain behavior, ignoring off-chain propagation and trust dynamics. Zhu et al. [10] demonstrate that poisoning attacks can subvert trust-based link prediction in signed social networks by examining adversarial vulnerabilities, but not misinformation diffusion, in the context of financial fraud. Sentiment Analysis in Cryptocurrency Diffusion was covered by Kajol et al.[11], who linked sentiment (optimism and trust) to cryptocurrency adoption patterns using social network analysis (SNA) techniques. Stitini et al.[12]propose a trust-enhanced semi-supervised recommendation model that combines trust networks with fake news detection.

Table 1: Summary of Related Work

Paper	Key Methodology	Limitation
Pierri et al. [6].	Diffusion topology analysis	No trust, sentiment, or financial context
Phan et al.[7]	GNN + trust metrics	News media only
Wu et al.[8]	Hybrid GNN for fraud	On-chain focus w/o social cascades
Patel et al. [9]	SHAP interpretability for fraud transactions	Ignores social content dynamics
Zhu et al. [10]	Poisoning attacks in signed trust networks	Adversarial focus only
Kajol et al. [11]	Sentiment-driven SNA for crypto adoption	Generic adoptions only
Stitini et al. [12]	Trust-aware fake news detection	Not crypto-focused

3. METHODOLOGY

This study employs a multi-layered analytical framework that integrates quantitative social media analytics and network theory to analyze the OneCoin cryptocurrency scam. The approach includes data collection, exploratory analytics, sentiment assessment, diffusion modeling, and cross-scenario validation to ensure methodological completeness and reproducibility as clarified in the flow chart (fig.1) :

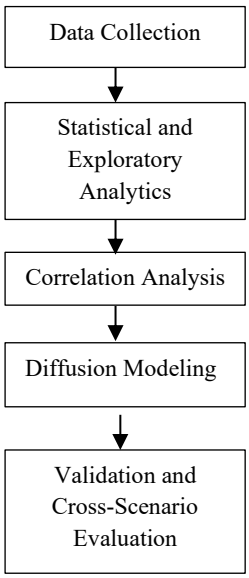

Fig 1: The proposed Framework

3.1. Data Collection

Data were collected from multiple social platforms associated with the OneCoin scheme, including YouTube, TikTok, LinkedIn, and Twitter (X). Third-party monitoring tools such as Brand24 and TweetBinder were used to retrieve historical hashtag frequencies, sentiment patterns, and engagement metrics. The core dataset consists of 51 videos from the official OneCoin YouTube channel (2015–2018). Extracted variables include view count, like count, comment count, timestamp, and metadata. Derived metrics include engagement rate, hashtag distributions, and sentiment polarity scores. Data were normalized and inspected for inconsistencies.

3.2 Statistical and Exploratory Analytics

Descriptive statistics were computed to quantify engagement behavior, including mean, median, standard deviation, and variance for views and interactions. Time-series visualization techniques were employed to examine temporal trends. Boxplots were generated to identify outliers and skewness in engagement distributions. Pearson correlation coefficients were calculated to evaluate relationships among key metrics, including likes, comments, views, and engagement rate. A heatmap visualization was developed to illustrate correlation patterns**.**

Statistical Analysis:

- **Descriptive metrics** such as average views (46,372), maximum views (438,728), and average engagement (275 interactions).
- **Boxplots** to visualize outlier behavior in view counts, likes, and comments.
- **Trend plots** to measure engagement and view performance over time.

Engagement Rate Tracking: The average engagement rate is **0.76%**.A time-series graph is plotted to illustrate the changing efficiency of engagement over time. The peak in 2016 indicates the effectiveness of strong promotional campaigns and viral

content.

3.3 Correlation & Heatmap Analysis

A correlation Matrix was created to understand the relationships between Views and engagement (high: r = 0.95), Likes and engagement (very high: r = 1.00), engagement rate, and other metrics (weak correlation: r = 0.40). The Engagement Rate appears independent of reach, suggesting that micro-influencer efficiency and content virality quality are key factors.

3.4 Diffusion Modeling and Interpretation

Information Diffusion Theory Applied: Cascade Models, where adoption patterns are observed as early influencers gain traction. Trust Propagation Models: Ruja Ignatova and regional leaders exploited authority bias. Engagement as Proxy for Spread Velocity: Higher engagement in fewer-viewed videos suggests depth over breadth.

3.5 Validation and Cross-Scenario Evaluation

Validation was conducted through multi-scenario comparison, examining variations in engagement across years, platforms, and content categories. Additional robustness checks were performed to determine whether the findings generalize beyond the YouTube dataset. Cross-platform consistency was evaluated using comparative distributions of metrics and sentiment trends.

4. RESULTS

4.1 YouTube Video Analysis Results

4.1.1 Reach and Audience Interaction

A total of 51 videos were analyzed from OneCoin's official YouTube channel. **Metrics Analysed in Table 2:**

- **Average Views**: This represents the typical number of views per video, indicating how widely the content was reached.
- **Average Engagement (Likes + Comments)**: This metric reflects the average level of viewer interaction.
- **Maximum Views**: Identifies the video with the highest reach.
- **Maximum Engagement**: Highlights the video that received the most audience interaction.

Table 2: One Coin YouTube Account Statistics

Number of subscribers	28000
Number of videos	51
The average number of views per video	46,372
The average engagement	275
maximum views recorded for a video	438,728
video with the highest engagement	1,757

4.1.2 Engagement Rate and Diffusion

The mean engagement rate across all videos in Fig. 2 was 0.76%, indicating moderately effective audience engagement. Engagement rate trends rose from 2015 to 2016, then declined in 2017, and then rose sharply again in 2018. This suggests that despite reduced video volume in later years, targeted or emotionally driven content yielded high per-view interaction, potentially due to stronger calls to action or more persuasive messaging.

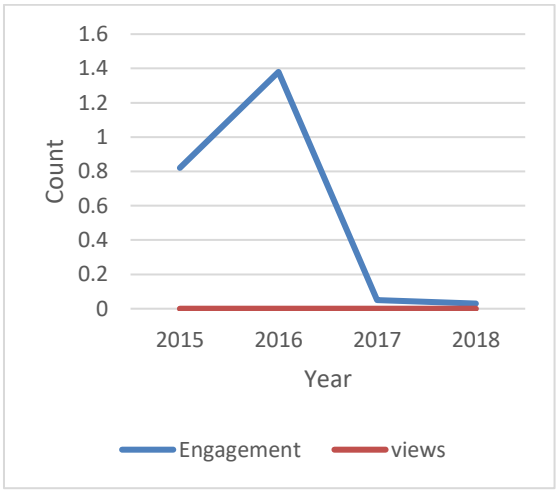

Fig 2: Yearly Trends in Views and Engagement

The graph illustrates the evolution of views and engagement metrics over time. Peaks in views or engagement indicate periods of high audience interest, potentially due to compelling content or increased promotion. Declines might suggest revisiting content strategies. This diagram highlights the years with the highest audience interaction, which may be due to seasonal or strategic factors affecting content performance. Diffusion is measured using the **Engagement Rate**, defined as:

$$Engagement\ Rate = \frac{Like+Comments}{Views} * 100 \quad (1)$$

Years with higher engagement rates, as shown in Fig. 3, indicate more efficient audience interaction per view. A declining trend could indicate less compelling content or audience engagement saturation. The Average Engagement Rate was **0.76%**

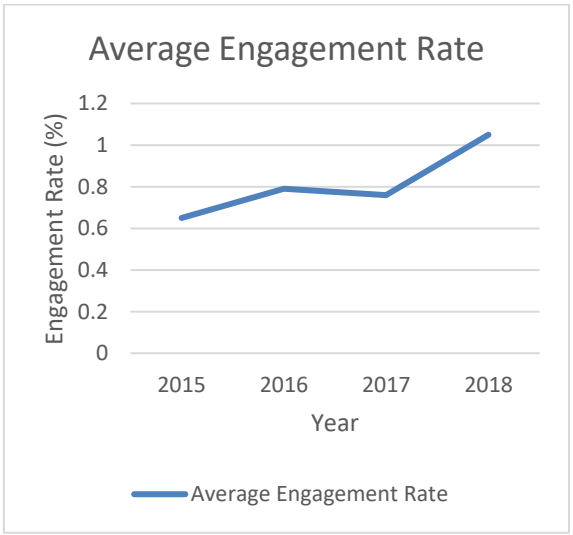

Fig 3. Diffusion Over Time(Yearly Average Engagement Rate)

4.2 Temporal Patterns in Viewership and Engagement

4.2.1 View Trends by Year

The highest average engagement and viewership occurred in 2016–2017 Fig 4, aligning with peak promotional efforts and the global expansion of the scheme.

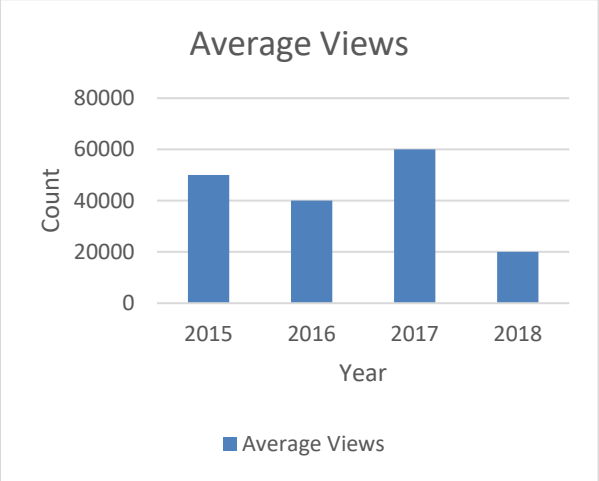

Fig 4: Average View Count by Year

4.2.2 Box Plot Analysis

Outliers in views, likes, and comments confirmed the presence of viral content, with some videos gaining 10× the average engagement, consistent with the cascade effect in diffusion models. Box plots also revealed right-skewed distributions, indicating that a few highly viral videos significantly inflated the mean, as shown in Figs 5, 6, and 7.

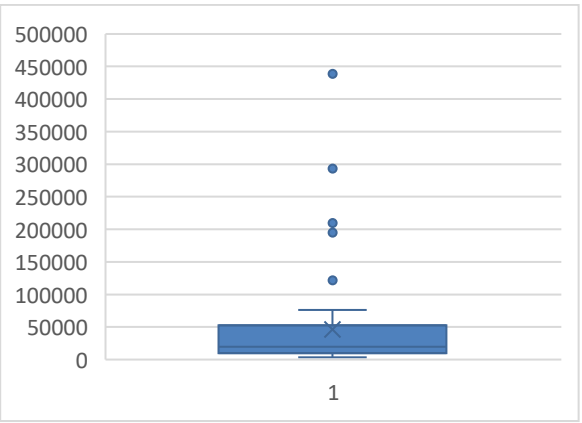

Fig 5: Box Plot for View_Count of OneCoine Official YouTube Page

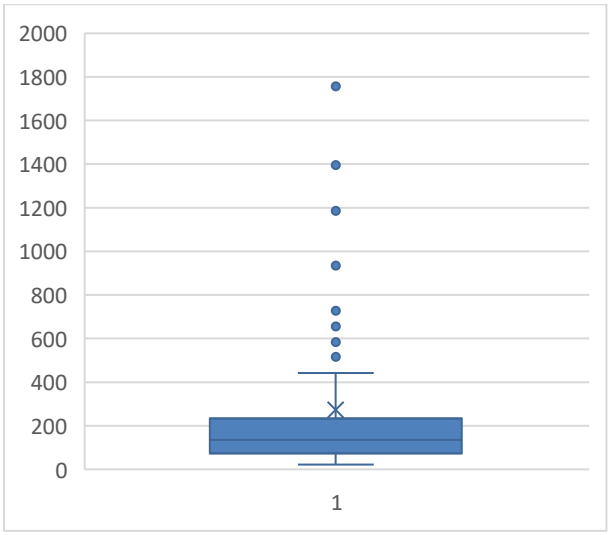

Fig 6: Box Plot for Like_Count of OneCoine Official YouTube Page

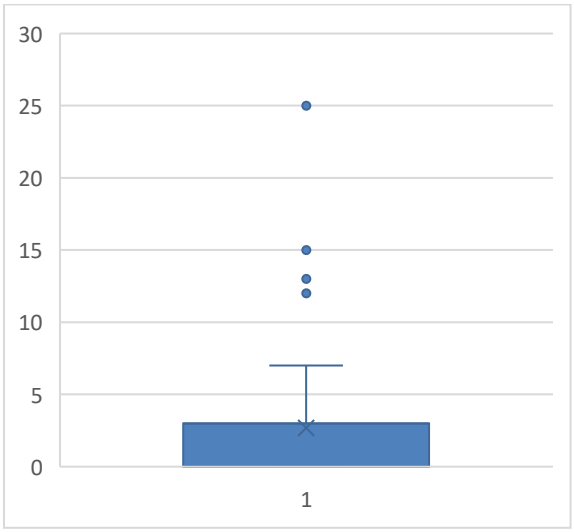

Fig 7: Box Plot for Comment_Count of OneCoin's Official YouTube Page

4.3 Correlation & Heatmap Analysis

Views and Engagement: There is a strong positive correlation between views and engagement (likes + comments). Videos with more views tend to have higher interactions.

Likes and Engagement: Likes have a very high correlation with overall engagement, as they are a significant component of the metric.

Engagement Rate: Weak correlation between engagement rate and other metrics (e.g., views). This suggests that the engagement rate is relatively independent of total views, reflecting how efficiently videos engage their audience.

.

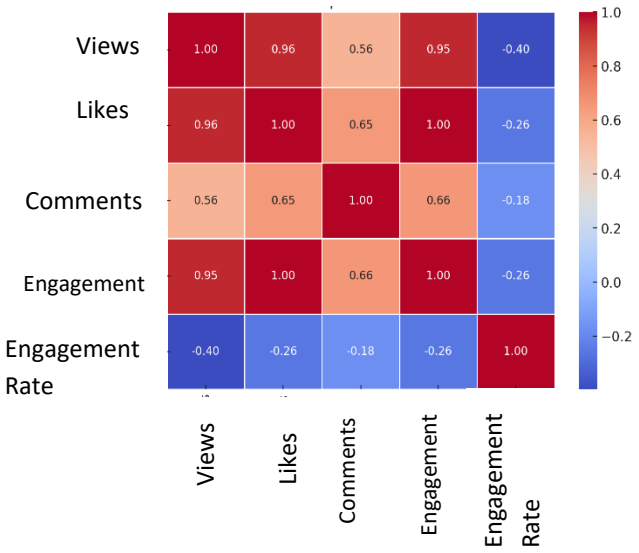

Fig 7: Collaboration Heat map between numerical data

We would expect to see a strong positive correlation in Fig. 7, close to 1, between view count and like count. This means that videos with higher view counts tend to have higher like counts. This is quite common in social media data. There would also be a moderate positive correlation between view_count and comment_count, as well as between like_count and comment_count. This suggests that videos with higher view counts tend to receive more comments, and videos with more likes also tend to have more comments.

5. DISCUSSION

The findings highlight the interplay between influencer amplification, emotional appeal, and platform structure in driving the viral spread of financial misinformation. Despite modest engagement rates, strategically crafted content achieved substantial reach due to trust-based cascades. Weak correlations between engagement rate and reach reveal that depth of interaction defines propagation efficiency. These insights suggest that early-stage detection of trust-driven cascades is essential for preventing large-scale fraud dissemination. The consistency of diffusion patterns across platforms underscores the trans-channel nature of misinformation.

6. CONCLUSIONS

The analysis of the OneCoin case reveals how social media and network-based trust dynamics can be leveraged to facilitate the rapid and large-scale dissemination of financial misinformation. By employing a multi-method analytical approach—integrating engagement metrics, sentiment analysis, correlation modeling, and information diffusion theories—this study provides a comprehensive view of how deceptive narratives can achieve viral propagation within digital environments.

Key findings indicate that the OneCoin campaign featured relatively modest average engagement rates (0.76%) but employed a strategically targeted dissemination approach, enabling it to reach significant global audiences. Influencer amplification and emotionally charged messaging emerged as critical factors in boosting trust and accelerating diffusion. The cascade behavior observed in user interactions, particularly during promotional peaks, underscores the influence of both network topology and content design on diffusion efficiency.

The study also demonstrated that traditional volume metrics (views) do not always align with audience efficiency (engagement rate), suggesting that depth of influence, not just breadth, is central to misinformation spread. Furthermore, the weak correlation between engagement rate and total views highlights the importance of micro-influencer networks and targeted trust-based diffusion.

These insights underscore the importance of early-warning systems and real-time monitoring models that incorporate not only content metrics but also behavioral and structural aspects of social propagation. By analyzing the intersection of trust, sentiment, and information flow, future systems can be designed to detect, flag, and counter emerging financial misinformation campaigns before they reach critical mass.

7. REFERENCES

[1] S. Vosoughi, D. Roy, and S. Aral, "The spread of true and false news online," science, vol. 359, no. 6380, pp. 1146–1151, 2018.

[2] News, in The Cryptoqueen who disappeared with billions, 2023.

[3] Fbi, in Ten Most Wanted: Ruja Ignatova, 2023.

[4] W. Chung, Y. Zhang, and J. Pan, "A theory-based deep-learning approach to detecting disinformation in financial social media," Information Systems Frontiers, vol. 25, no. 2, pp. 473–492, 2023.

[5] Brand24, Ed., "Hashtag and sentiment analysis for cryptocurrency scams." 2024.

[6] F. Pierri, C. Piccardi, and S. Ceri, "Topology comparison of Twitter diffusion networks effectively reveals misleading information," *Scientific reports*, vol. 10, no. 1, p. 1372, 2020.

[7] H. T. Phan, N. T. Nguyen, and D. Hwang, "Fake news detection: A survey of graph neural network methods," *Applied Soft Computing*, vol. 139, p. 110235, 2023.

[8] X. Wu, "A trust-based detection scheme to explore anomaly prevention in social networks," *Knowledge and Information Systems*, vol. 60, pp. 1565–1586, 2019.

[9] O. Patel, "ANOMALY DETECTION IN CRYPTOCURRENCY TRANSACTIONS USING MACHINE LEARNING."

[10] Y. Zhu, T. Michalak, X. Luo, X. Zhang, and K. Zhou, "Toward secrecy-aware attacks against trust prediction in signed social networks," *IEEE Transactions on Information Forensics and Security*, vol. 19, pp. 3567–3580, 2024.

[11] K. Kajol, S. Devarakonda, R. Singh, and H. K. Baker, "Drivers influencing the adoption of cryptocurrency: a social network analysis approach," *Financial Innovation*, vol. 11, no. 1, pp. 1–25, 2025.

[12] F. Mlika, W. Karoui, and L. Ben Romdhane, "Trustworthy decentralization based on blockchain tools for social network architectures," *Social Network Analysis and Mining*, vol. 14, no. 1, p. 95, 2024.