Capozzi *et al. Applied Network Science* (2025) 10:28
https://doi.org/10.1007/s41109-025-00711-0

Applied Network Science

RESEARCH Open Access

FlowSeries: flow analysis on financial networks

Arthur Capozzi[1*], Salvatore Vilella[2], Dario Moncalvo[3], Marco Fornasiero[3], Valeria Ricci[3], Silvia Ronchiadin[3] and Giancarlo Ruffo[2]

*Correspondence:
Arthur Capozzi
arthur.capozzi@gess.ethz.ch
[1]Computational Social Science, ETH Zürich, Zurich, Switzerland
[2]DISIT, Università degli Studi del Piemonte Orientale "A. Avogadro", Alessandria, Italy
[3]Anti Financial Crime Digital Hub, Turin, Italy

Abstract
The digitalization and automation of anti-financial crime (AFC) investigations has made significant progress in recent years. However, key challenges remain—in particular, the need for interpretability in the output of AI models and the limited availability of labeled data for training. Criminal activity in transaction networks often involves complex, evolving patterns specifically designed to evade detection. We introduce FlowSeries, a top-down flow analysis methodology to explore transaction data and analyze complex interaction patterns over time. Rather than relying on predefined patterns or labeled training data, our approach scales to large transaction volumes and provides interpretable insights into anomalous behaviors, aiding AFC analysts in their investigations. We evaluate the effectiveness of this method using a dataset provided by the bank Intesa Sanpaolo (ISP), comprising 80 million cross-border transactions over a 15-month period. In collaboration with ISP's AFC experts, our analysis focuses on detecting anomalous transactions and identifying suspicious actors in the context of the economic sanctions imposed on Russia following its invasion of Ukraine on February 24th, 2022.

Keywords Network analysis, Anti financial crime, Anti money laundering, Temporal networks, Graph search algorithm

Introduction

The development of effective financial crime detection models presents significant challenges, as these models need to be: (i) explicable, (ii) scalable, and (iii) highly adaptable. In particular, AFC investigative tools must provide a high level of interpretability—an essential requirement for regulators to initiate formal investigations. As a result, many black-box models, including some deep learning approaches, despite their potential effectiveness, fall short of the transparency standards required for real-world application.

In 2020, the Single Euro Payments Area (SEPA) supported financial transactions for more than 523 million European citizens across 36 countries, processing more than 25.1 billion credit transfers and 23.2 billion direct debits (Bank 2021). Given this enormous volume of transactions, even small to medium-sized financial institutions in Europe need to deploy analytical models that can efficiently scale to handle large datasets and detect complex transaction patterns.

In addition, the AFC domain faces a significant challenge due to the scarcity of labeled data, largely due to privacy regulations and data protection requirements. To mitigate this limitation, some approaches rely on generic, pre-defined transaction patterns that are commonly flagged as suspicious due to their potential to mask criminal or illegal activity. However, these patterns must be continuously validated and updated to remain effective against evolving crime tactics.

Against this background, we present FlowSeries, a top-down search methodology designed to act as a "magnifying glass" for AFC analysts, aiding in the detection of potentially illicit transactions and non-compliant entities. We apply FlowSeries to a dataset of 80 million anonymised cross-border banking transactions over a 15-month period. This dataset, described in Sect. 4, was provided by Intesa Sanpaolo and complies with all relevant legal privacy and security regulations. Data supporting the findings of this study are available upon request from Intesa Sanpaolo's AFC Digital Hub[1]. The use cases presented in Sect. 5, developed in close collaboration with Intesa Sanpaolo's AFC experts, focus on detecting attempts to circumvent the economic sanctions imposed on Russia following its military invasion of Ukraine on February 24th, 2022.

This manuscript is an extended version of the work presented in the International Conference on Complex Networks and Their Applications, held in Istanbul in December 2024 (Capozzi et al. 2024).

Related work

Financial markets are inherently complex, adaptive and dynamic systems in which different actors—including hedge funds, individual investors and banks—interact in ways that influence overall market stability. Complex network theory provides a powerful framework for modeling various facets of economics and finance, and considerable research has focused on identifying the key drivers behind both stabilizing and destabilizing market dynamics (Lillo et al. 2008; Mu et al. 2010).

A wide range of studies have examined different aspects of stock markets, including correlations between stock prices (Vandewalle et al. 2001) and the structure of shareholder networks (Caldarelli et al. 2004). Other research has analyzed the topological properties of financial market networks, with particular attention to their resilience to destabilizing shocks (Kauê Dal'Maso Peron et al. 2012; Yan et al. 2014). Rauch (2001) has studied international trade, highlighting its complex interdependence with the social networks of countries shaped by linguistic, cultural and religious ties. Network-based analytical frameworks have also proved valuable in the study of financial crises, particularly the 2008 global financial crisis (Faggini et al. 2019; Lee et al. 2011; Leila 2011).

Beyond stock and trade markets, complex network theory has been applied to a variety of economic and financial domains. For example, López et al. (2003) used network-based approaches to model competitive dynamics in the World Wide Web market. Similarly, network analysis has been widely used to study cryptocurrency markets from different perspectives. Papadimitriou et al. (2020) conducted a comprehensive three-year study to identify dominant cryptocurrencies, while other research has focused on the structural properties of cryptocurrency transaction networks (Serena et al. 2022) and the temporal

[1] 2adh@pec.afcdigitalhub.com.

and multiplexing characteristics of connections within the Ethereum network (Lin et al. 2020).

Network analysis for AFC and AML

Network analysis has been widely used in the fight against financial crime, particularly in the areas of anti-financial crime (AFC) and anti-money laundering (AML), and more generally in the detection of anomalies within transaction networks. Complex networks provide a natural and effective framework for representing transaction data. Despite the increasing digitization of crime detection processes, human insight remains essential. Automating some or all of this complex workflow can significantly improve the efficiency of financial crime detection, while optimizing the allocation of time and resources.

In this research context, many scholars have emphasized the value of network analysis in strengthening the operational capacity of financial intelligence units. A notable example is the VISFAN system (Didimo et al. 2011), which supports the visual exploration of networks of financial activity. This system uses complex network metrics to help identify suspicious transactions or actors, even within large and dense datasets.

Network analysis, as outlined by García and Mateos (2021), has been widely applied in investigative processes within the Spanish Tax Agency's Tax Control Study for the period 2015 to 2020. In their work, the authors present case studies that highlight the practical implementation of network analysis in real-world scenarios. Graph-based pattern recognition algorithms enable the identification of criminal activities. In addition, García and Mateos (2021) use community detection techniques to refine the representation of the economic landscape. The importance of social network metrics in detecting money laundering practices associated with corporate entities is underlined by Fronzetti Colladon and Remondi (2017). The authors present a strategic mapping of four different types of relational graphs designed to capture various risk factors, such as economic sectors, geographical regions, transaction volumes and links between entities with common ownership or representatives. Through visual analysis of these graphs, the authors effectively identify clusters of companies involved in litigation. Similarly, the CoDetect framework(Huang et al. 2018) is presented as a tool that combines both network data and feature data—describing entities—to improve the detection of financial crime.

AMLSim[2] is a project presented in 2018 (Weber et al. 2018), which aims to provide a multi-agent-based simulator that generates synthetic bank transaction data along with a predefined set of recognizable money laundering patterns. In addition, preliminary results are presented showing that (i) graph learning for AML remains feasible even in the context of large, sparse networks with 1 billion nodes and 9 billion edges; (ii) memory-efficient graph representations based on the Ligra+ graph compressor (Shun et al. 2015) exhibit compression ratios of up to $2x$.

The approach proposed by Garcia-Bedoya et al. (2020) combines artificial intelligence with network analysis techniques. They highlight the limitations of many existing anti-money laundering (AML) methods, which often fail to detect money laundering because they rely on static analysis conducted days or months after financial transactions occur. The authors also emphasize that within a transactional network, the nodes involved in money laundering typically have complex connections that are deliberately designed to

[2] https://github.com/IBM/AMLSim.

conceal illicit financial activity. In particular, they identify three different types of interactions that are used to conceal money laundering attempts:

- Path: Money is sent from node x to node y through multiple intermediaries. The single transaction is hidden by a path of transactions.
- Cycle: The money starts from node x and after a path of transactions returns to x.
- Smurf: The single transaction is divided into multiple smaller transactions that reach a target node through both physical and legal intermediaries.

The study by Liu et al. (2020) suggests that also network cycles, generated by time-sequenced transactions, can also serve as indicators of potential crime in online transaction networks.

Alternative methods often rely on machine learning paradigms (Chen et al. 2018) or explore deep learning techniques, such as the use of graph neural networks (Kute et al. 2021). In the absence of annotated data suitable for training, many machine learning models turn to unsupervised anomaly detection approaches (Chen et al. 2018). Innovative strategies include zero-shot learning and its variants, such as one-shot learning or few-shot learning, as well as meta-learning. Pan (2022) advocates a deep-set algorithm that combines both meta-learning and zero-shot learning for money laundering detection. In the initial meta-learning phase, the model learns to make contrastive comparisons, enabling it to judge the membership of a query point against a set of positive and negative samples. The model is then further trained using zero-shot learning techniques to improve its accuracy.

Broader reviews of anomaly detection, particularly in areas such as crime detection (Bolton and Hand 2002; Phua et al. 2010), highlight the exploration of social networks—such as those derived from mobile phone communications or financial transactions—as a valid strategy for detecting illicit activity. Identifying anomalies within these network structures has proven effective in uncovering organized criminal behavior, as seen in cases such as insurance crime (Šubelj et al. 2011).

Our contribution: FlowSeries

Detecting illicit activity within transaction networks requires highly flexible and adaptive analytical models. Financial crime often involves complex and dynamically evolving transaction patterns that are specifically designed to evade detection. Unlike traditional methods that rely on predefined patterns, such as some presented in the previous section, FlowSeries operates without such constraints. As shown in Fig. 1, these patterns can involve multiple payment paths routed through an undefined number of intermediaries. We define this sequence of payment paths from a source x to a destination y, potentially involving intermediaries, as a *transaction flow*.

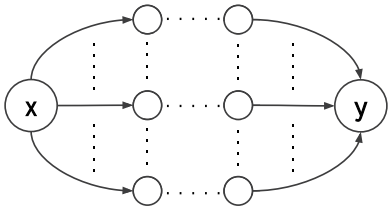

Fig. 1 Node x sends money to node y through multiple payment lines involving multiple intermediaries. We define this scheme as transaction flow from x to y

Section A of the appendix presents an interview with the Anti-Financial Crime Digital Hub, offering insights into the investigative processes of an AFC team and an evaluation of the implementation choices behind FlowSeries.

The application of FlowSeries begins with the definition of the appropriate level of structural granularity for the transaction network. Using the data provided by Intesa Sanpaolo Bank, we construct a transaction network in which each node can represent a *ISO code*, a *BIC* or a *IBAN*. ISO 3166 defines country codes and is published by the International Organisation for Standardisation (ISO). The BIC (Bank Identifier Code) is an $11 or 8$ digit code used in international payments to identify the beneficiary's bank. The International Bank Account Number (IBAN) is a globally recognizee system for identifying bank accounts across national borders.

The next step is to create multiple weighted directed temporal networks, each corresponding to a specific time period. This is done by aggregating transfers over time, with typical aggregation periods including weeks, months, quarters or years. Once a node of interest is selected, an algorithm computes all possible paths from that node for each temporal aggregation. At this stage, the AFC analyst can explore the paths, looking for trends, recurring patterns or anomalies. The analyst can also aggregate all paths—regardless of length—between two nodes and estimate the maximum amount of money transferred from one node to another through multiple payment lines and intermediaries.

The application tests of FlowSeries, presented in Sect. 5, yield the following observations:

1. FlowSeries is effective in supporting the AFC analysts in identifying anomalous transaction flows.
2. FlowSeries can be applied on databases of millions of transactions.
3. Unlike black box techniques, FlowSeries's interpretability allows the analyst to initiate a formal investigation.

The implementation of the FlowSeries algorithm, customized for synthetic data experiments, is publicly available in an open-access GitHub repository (Capozzi 2025). In the following section, we formalize the concepts of transaction networks and transaction flows. We also describe the algorithm for identifying transaction flows and propose a method for weighting the transaction flow between two nodes. In Sect. 4, we present the data provided by Intesa Sanpaolo and the AFC Digital Hub to support the experiments, while in Sect. 5 we present real-world examples showing how the FlowSeries pipeline assists AFC investigations.

Methodology

Transaction networks as weighted directed temporal graphs

A transaction network can be represented as a weighted directed temporal graph where nodes are the actors and edges are the transactions. A weighted directed temporal network can be defined as $G = (V, L)$, where $V = \{id_1, id_2, ..., id_n\}$ is the set of nodes, and $L = \{(e(i_1, j_1), t, a), ..., (e(i_n, j_n), t, a)\}$ is the set of edges such that $i, j \in V$, t is the timestamp, and a is the weight of the edge representing the total amount of money transferred from node i to node j. So, an edge represents a transaction, a node represents an bank account identified by an IBAN code. Timestamps domain is $\tau = [t_0, t_\omega]$, where

t_0 is the time of the first transaction and t_ω is the time of the last transaction. It is possible to aggregate the timestamps in time intervals T within τ and defined by an interval $[t_x, t_y[$ where $t_x, t_y \in \tau$ and $t_x \leq t_y$, to extract a temporal layer G_T from G. If we consider a time interval T equal to τ, it would be like removing the time variable from the definition of transaction network because τ is the largest possible time interval, so that $G_T = G$ in this case.

Each node has the properties of the IBAN it represents. These include the BIC, the bank branch ID, and the country to which it belongs. The transaction network can then be aggregated at different levels of granularity by exploiting these node properties. In this paper, we focus on the transaction network between countries and between BICs.

In the case of aggregation by country, the transaction network at interval T is defined as $G_T^C = (V^C, L_T^C, \phi)$. V^C is the set of countries, and the edge $e(i, j)$ with $e(i,j) \in L_T^C$ and $i, j \in V^C$ represents all the transactions of the IBANs of country i towards the IBANs of country j. The weight $w(e(i, j))$ is computed by applying the aggregation function ϕ to all the transactions during the time interval T between node i and node j. In this study, we use SUM as the aggregation functions: the weight of the edges represents the total amount of transactions from node i to node j during the time interval T. Different types of analysis may involve COUNT as the aggregation function. The weight of the edges represents the number of transactions from node i to node j during the time interval T.

Equation (1) defines a path in G_T as an ordered finite collection of n distinct edges in such a way that they connect the vertices i and j at time interval T.

$$P_{i,j}^T = \{e(v_1, v_2), e(v_2, v_3), ..., e(v_{n-1}, v_n)\}, \quad \text{with } v_i \in V \text{ and } e(v_i, v_{i+1}) \in L_T \text{ for all } i \in \{1, ..., n-1\} \tag{1}$$

The weight of the path $P_{i,j}^T$ is given by $\sum_{i=1}^{n-1} w(e(i, i+1))$. We can write $W(P_i)$ as $\{w(e_1, e_2), ..., e(v_{n-1}, v_n)\}$, so the set containing the weights of the edges of the path P_i. $\text{Paths}_{i,j}(G_T)$ is the set of all possible paths from i to j on the graph G at the interval T.

Weight a transaction flow

In many real-world scenarios, to evade detection, an agent may route payments through a series of intermediaries before they reach the intended recipient. As illustrated in Fig. 1, node x avoids sending money directly to node y by using an unknown number of intermediaries. It is not possible to determine in advance either the total number of paths connecting x to y, or the number of intermediaries involved in each individual path (i.e., the length of each path).

A transaction from one actor to another, routed through n intermediaries in a transaction network, corresponds to a path of length $n + 1$ between the two nodes. The core idea behind the FlowSeries pipeline is to identify transaction flows originating from a given input node x and to determine the maximum amount of money that could be transferred from x to any other node within a maximum distance of n.

Equation (2) defines $Flow^n(x, y)$ as the set of all paths of maximum length n from node x to node y.

$$Flow^n(x, y) = \{P_1, ..., P_m\}, \text{ with } P_i \in \text{Paths}_{i,j}(G_T) \tag{2}$$

In a transaction network, the maximum hypothetical amount that a node x could try to send to node y through a group of intermediaries is the minimum weight of the edges involved in the path between x and y. In the more complex case where there are multiple groups of intermediaries and the money transaction is therefore hidden by multiple paths, we consider the weight of the money flow between x and y as the sum of the minimum weights of each path between x and y. Thus, we define the weight of a flow of maximum length n from x to y as follows:

$$w(\text{Flow}^n(x,y)) = \sum_{i=1}^{m} \min(W(P_i)), \text{ with } P_i \in \text{Paths}_{i,j}(G_T) \tag{3}$$

Figure 2a presents a simple transaction network. To assess the hypothetical amount of money transferred from node x to node y, it is insufficient to consider only the direct edge weight $e_{x,y}$. Nodes such as h, k, and z may act as intermediaries through which additional flows from x to y can occur. Table 2b lists the relevant paths and the corresponding minimum weight for each. The total flow weight is computed as the sum of the minimum weights across all paths from x to y. In this example, the total flow from x to y amounts to 1, 850.

Note that the temporal aggregation of a network for the interval T is an approximation. In fact, each transaction has a timestamp t and therefore all edges involved in a valid path should satisfy the constraint $t_{e_i} < t_{e_{i+1}}$.

Pseudocode

The pseudocode underlying the FlowSeries algorithm is presented in Algorithm 1. It is a variant of the Depth-First Search (DFS) graph traversal algorithm (Tarjan 1971), and performs a recursive search over a network G to identify all paths of maximum length n starting from a given input node u. The algorithm explores nodes in a depth-first approach, starting at node u and terminating each recursive branch when the path reaches the specified maximum length n. The output of the Algorithm 1 is a list of lists, where each sublist corresponds to a distinct path and contains the ordered sequence of nodes traversed.

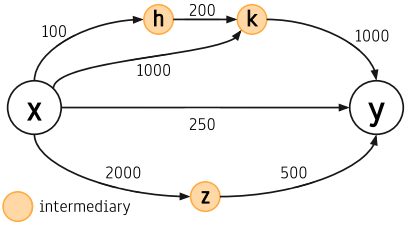

(a) Figure

P_i	$min(W(P_i))$
$e(x,h), e(h,k), e(k,y)$	100
$e(x,k), e(k,y)$	1000
$e(x,z), e(z,y)$	500
$e(x,y)$	250

(b) Table

Fig. 2 In the transaction network in (**a**), node x has an edge of weight 250 towards node y. The paths between nodes x and y through nodes h, k and z could be an attempt to hide a direct edge of higher weight. Table (**b**) lists all paths of maximum distance 3 from x to y and their respective minimum weights

```
1: procedure FindPaths(G, u, n, n_recursive)
2:     paths← [ ]
3:     if n_recursive == 0 then
4:         return [[u]]
5:     end if
6:     if n_recursive < n then
7:         paths.append([u])
8:     end if
9:     for neighbor in G.neighbors(u) do      ▷ loop through all neighbours of u
10:        for path in FindPaths(G, neighbor, n, n_recursive - 1) do
11:            if u not in path then
12:                path.insert(0, u)      ▷ append node u on the head of path list
13:                paths.append(path)
14:            end if
15:        end for
16:    end for
17:    return paths
18: end procedure
```

Algorithm 1 Flow Algo

The time complexity of the DFS algorithm computed with an adjacency list is $O(|V| + |E|)$. The DFS recursion can be called on all nodes, as there may be disconnected components. In the Algorithm 1, the recursion only starts at node *u*. This means that $|V|$ does not play a role in the time complexity. However, edges can be visited multiple times, since an edge can be part of multiple paths. Therefore the time complexity is greater than $O(|E|)$. The complexity of the algorithm is determined by the number of times an element is added to a *path* (line 13 in algorithm1). So the complexity depends on the number of nodes contained in all paths.

To determine the complexity of the algorithm, we first consider the toy network in Fig. 3a. The network has 16 nodes and 20 edges. Table 3b shows the number of paths from node *S* and the total number of nodes (i.e., the sum of the length of each path) as *n* varies. At each "crossing" of the network (and at the starting node), there are two possible paths, resulting in a total of $2^{n/2}$ paths of length *n*. In addition to these, we must also account for the paths of length less than *n*. Therefore, for this type of network, 2^n serves as an upper bound on the total number of paths. To compute the total number of nodes, we consider that each path can have at most *n* edges, and hence $n + 1$ nodes. In the network in Fig. 3a, the total number of nodes can thus be at most $(n + 1) * 2^n$.

Now we consider a complete graph. In this case, the number of paths is independent of the starting node or the intermediate nodes traversed, since all nodes have the same

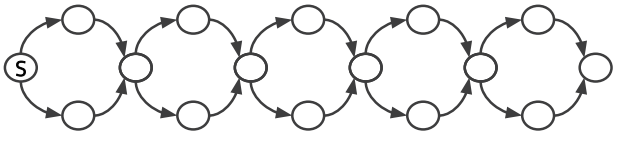

(a) Figure

n	Paths	Nodes
2	4	10
3	8	26
4	12	46
5	20	94
6	28	150
7	44	278

(b) Table

Fig. 3 **a** A particular transaction network configuration. The paths of maximum length *n* are computed from the node *S*. The total number of paths found and the total number of nodes (i.e. the sum of the length of each path) as *n* varies are shown in Table **b**

out-degree (i.e., the number of edges leaving a vertex in a directed graph). Therefore, if d_{out} is the out-degree of any node in the network, the total number of nodes included in paths of maximum length n is given by $(n+1) * d_{out}^n$.

In the general case of a non-complete graph, we can consider the maximum out-degree of the network, $max(d_{out})$, and assert that the total number of nodes in the paths of maximum length n (and hence the time complexity) is at most $(n+1) * max(d_{out})^n$.

In real-world scenarios, it is important to consider that the degree distribution of wire transfer networks follows a power law distribution (Semeraro et al. 2020), where most nodes have very low out-degree, while a few nodes–referred to as hubs–have very high out-degree (Fig. 5). As a result, the benchmarks reported in Section D demonstrate that, especially in the BICs wire transfer network, the choice of the starting node significantly impacts the execution time of Algorithm 1.

Time series of flows

As shown in Fig. 4, transaction networks exhibit structural stability over time. As a result, anomalies typically arise not from the appearance of new edges between nodes, but rather from statistically significant changes in the weights of existing edges. Therefore, path level anomaly detection generally focuses on identifying fluctuations in edge weights rather than the discovery of new paths. In addition, edge weights in transaction networks often follow temporal patterns that reflect recurring financial activities, such as salary payments, service fees, or intercompany transfers. By analyzing a node's transaction history, these patterns can be modeled to detect deviations and highlight anomalous transactions.

In an AFC investigation, the application of FlowSeries involves analyzing the time series of $w(Flow)$. It is important to note that a flow is a complex structure composed of multiple paths, and as such, an anomaly at the level of a single edge may not have a significant impact on the overall weight of the flow.

Figure 7a shows the time series of the weight of a transaction flow between two nodes. The actual flow weight $w(Flow(x, y))$ is represented in orange, while the weighted moving average (WMA) and the exponentially weighted moving average (EWMA) are shown in blue and green, respectively. To identify anomalous flows, the deviation Δ_w at each time t is calculated by comparing the actual observed value $w_t(Flow(x, y))$ with the expected value derived from the WMA. This deviation is defined as

$$\Delta_{w,t} = \frac{|w_t(Flow(x, y)) - \mu_t|}{max(w(Flow(x, y)))} \tag{4}$$

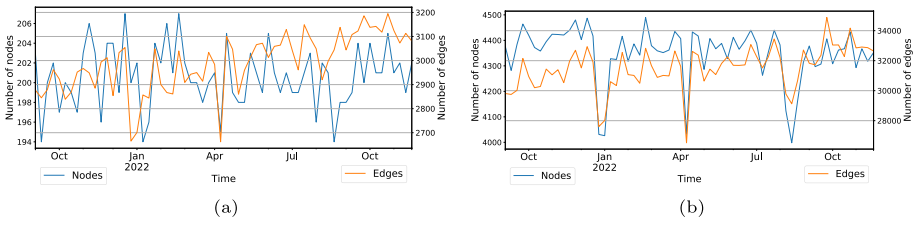

Fig. 4 Number of nodes and edges of transaction networks with weekly temporal aggregation. In **a** the nodes are aggregated by country, in **b** the nodes are aggregated by BIC

where $w_t(Flow(x, y))$ is the weight of the flow *Flow*(*x*, *y*) at time *t*, μ_t is the moving average computed over *w*(*Flow*(*x*, *y*)) at time *t*, and *max*(*w*(*Flow*(*x*, *y*))) is the maximum observed value of *w*(*Flow*(*x*, *y*)) over the entire time series.

Objective is to identify transaction flows that show a significant increase in flow weight over time, accompanied by a concomitant decrease in the amount transferred directly between the end nodes of the flow, i.e., without intermediaries. To achieve this, we implement a simple but effective strategy. First, for each flow, we compute the maximum deviation $\Delta_{w,t}$ over the time series. We then calculate $\Delta_{e,t}$, which quantifies the change in the volume of direct transactions between the initial and final nodes of the flow. Unlike Eq. (4), the numerator here does not use the absolute value, as we are particularly interested in highlighting reductions in direct transfers.

Finally, we rank all flows by their maximum $\Delta_{w,t}$, sorting them in descending order with respect to $\Delta_{e,t}$. This ranking allows us to prioritize flows that show the largest increase in routed (i.e. indirect) transaction activity, while at the same time registering the largest decrease in direct transactions between their endpoints.

The investigation of anomalies in the time series of *w*(*Flow*(*x*, *y*)) can benefit from more advanced analytical techniques, including forecasting models such as the Autoregressive Integrated Moving Average (ARIMA) model (Box et al. 2015) or its seasonal variant, the Seasonal Autoregressive Integrated Moving Average (SARIMA) model (Kumar Dubey et al. 2021). However, due to their high computational demands, these methods are typically only applicable to a limited number of nodes that have already been flagged as suspicious. As discussed in Sect. 5, the number of flows originating from a given country can be substantial, especially if the starting node has a large out-degree. However, AFC analysts often have additional domain-specific knowledge that allows them to focus on a smaller, more targeted subset of nodes and intermediaries. In such cases, advanced anomaly detection techniques or forecasting models could function as an effective early warning system.

Apply FlowSeries

In this section we formalize the concepts outlined above within a pipeline that can integrate the application of FlowSeries to the analysis of network flows. The pipeline is designed to analyze all potential flows of money between entities in a financial graph. This transaction network can be constructed at various levels of spatial and temporal aggregation. While this methodology is generalizable to any suitable graph, we focus specifically on financial graphs, as FlowSeries was initially developed to support AFC investigations. The application of FlowSeries within a pipeline can be described by the following steps:

1. Choose a structural aggregation of the nodes (e.g., BIC or country) and a temporal aggregation *T*.
2. For each temporal aggregation, create a weighted temporal transaction network G_T, as described in Sect. 3.1.
3. Select a node *x* from which to start the investigation, and, for each network G_T, execute Algorithm 1.
4. Store the result of the previous step in a table. Each row of the table represents a path in a time period *t*. For each path at time *t*, we have the list of the weights (SUM and

COUNT) of the edges involved in that path. Additional information of the nodes involved in the path can be included, such as the country of the BIC or the country risk.

5. The AFC analyst can examine the table by filtering and sorting the paths. Since we are interested in flows that may represent money transfers through intermediaries, we discard all flows for which a direct edge between the endpoints has never existed. This step is particularly relevant in the context of the examples presented in Sect. 5, where a previously stable network undergoes a significant perturbation (such as the introduction of economic sanctions against one or more of the nodes in the networks).
6. For each flow and each time t, we compute the weighted moving average, the $\Delta_{w,t}$(Eq. (4)), and the $\Delta_{e,t}$ between the endpoints, as explained in Sect. 3.4.
7. We select the maximum $\Delta_{w,t}$ for each flow and sort the resulting values according to $\Delta_{e,t}$. This ordering allows us to rank flows on the basis of both the increase in flow weight and the reduction in the amount directly exchanged between the endpoints.

Data

We evaluate the potential of FlowSeries by applying it to a dataset of 80 million cross-border transactions over 15 months provided by Intesa Sanpaolo (ISP), the largest Italian bank with operations across Europe. The dataset, described in more detail in Vilella et al. (2025), includes all cross-border transactions involving ISP customers, either as senders or recipients. It also includes transactions where ISP clients are not directly involved, but where the bank acts as an intermediary for financial partners. The dataset covers the period from September 2021 to November 2022. The data was provided to the research team in a fully anonymised form, in accordance with the strictest privacy and security regulations. The data supporting the findings of this study is available from ISP upon request to the AFC Digital Hub [3]

Several details are available for each transaction in the dataset. These include the transaction timestamp and the BIC codes of both the payer and the beneficiary. The ISO codes of the countries of residence of both the payer and the beneficiary and the ISO code of the country of the bank involved in the transaction are also recorded. Additional information includes the amount of the transaction, its currency, and the data stream of the transaction, which includes approximately 7,000 SEPA and 2,000 SWIFT transactions. In total, the dataset contains 8,000 different BICs from over 200 countries.

For each node in the network, an AFC analyst may have additional information, such as the financial risk associated with the node. This risk may be associated with a country, a BIC or even an individual IBAN. The financial risk is usually defined by the AFC analyst group, the country's banking authority, or internationally recognized institutions. For example, on August 18th, 2023, the European Commission adopted a new regulation[4] (Commission Delegated Regulation (EU) 2023/2070) to identify third countries with strategic deficiencies in their AML/AFC regimes that pose a significant threat to the European Union's financial system. These countries are referred to as "high-risk third countries".

[3] Please note that access to the data is subject to restrictions. Researchers interested in obtaining the data for academic purposes will be required to sign a non-disclosure agreement. For further information please contact adh@pec.afc-digitalhub.com.

[4] https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32023R2070.

In the analyses presented in this article, transactions are temporally aggregated on a weekly or monthly basis and structurally aggregated by country (G^C) and by BIC (G^B). The number of nodes and edges in the G^C and G^B networks when using weekly temporal aggregation are shown in Fig. 4a, b respectively. In each country network there are approximately 200 nodes per week, representing almost all countries in the world. In the G^B network there are about 4500 different BICs per week.

The first step of the FlowSeries tool is to apply the Algorithm 1 to identify all paths of maximum length n starting from an input node u. This step can be computationally intensive. The number of nodes and edges in the network affects the performance of the flow analysis and must be carefully considered. However, in our case, the density of the network may be more informative. We define the network density as $\frac{m}{n(n-1)}$, where n is the number of nodes and m is the number of edges. The median network density computed on G^C is 0.075, while the median network density computed on G^B is 0.0017. Since the nodes in the G^C network are aggregated at a higher level, the density of G^C is much higher than that of G^B.

As shown in Sect. D, the execution time of Algorithm 1 is highly sensitive to the choice of the input node and its involvement in transaction paths. A node is more likely to participate in many paths if it has a large number of neighbors, i.e. if its out-degree is high. Figure 5a, b show the out-degree distributions of G^C and G^B respectively. These distributions have long-tail characteristics: while the majority of nodes have a low out-degree, a small number of nodes—known as hubs—have a very high out-degree. This long-tail behavior is more pronounced in G^B than in G^C. In Section D we further analyze how this distribution affects the performance of the Algorithm 1.

Results

This section presents two case studies that demonstrate the practical application of FlowSeries in financial crime investigations. The examples employ networks constructed with different node aggregation strategies, and highlight the effectiveness of FlowSeries in detecting complex transaction patterns that may be indicative of financial crime.

All analyses are conducted in strict compliance with applicable privacy and security regulations, in accordance with the legal framework and the guidelines of the AFC Digital Hub consortium. In particular, transaction amounts are normalized relative to the maximum value within each time series. BIC codes—originally provided in fully anonymised form—are further anonymised, and the names of the countries involved in the transactions are also anonymised to preserve confidentiality.

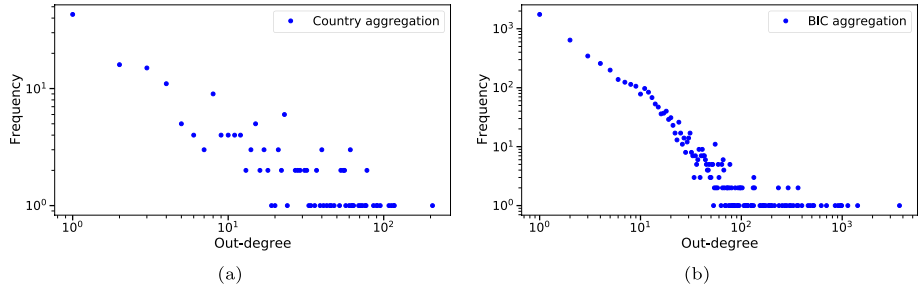

Fig. 5 Loglog frequency out-degree distribution of G^C (**a**) and of G^B (**b**). The time interval T is the first week of February 2022

Both case studies, based on real transaction data, focus on the broader European macroeconomic context, specifically analyzing changes in economic relations between nations following the imposition of sanctions on Russia following its invasion of Ukraine. Due to confidentiality requirements, we cannot disclose operational details of AFC Digital Hub investigations in which FlowSeries has been used. However, numerous journalistic investigations have documented recurring cases of sanctions evasion facilitated by the use of intermediaries. Countries such as Turkey, Georgia (AP 2023), Armenia (Times 2023), Kyrgyzstan, Kazakhstan, Tajikistan (Wagner 2023) and other Central Asian countries (Warrick 2023) have been publicly accused of enabling the flow of goods and capital between Europe and Russia. Despite pressure from the European Union and the United States to limit these financial corridors, their detection remains a significant challenge due to the high volume of transactions and the large number of potentially involved jurisdictions. The FlowSeries framework has been designed to address this challenge by providing a scalable and systematic approach to support financial crime investigations in complex international environments.

In the appendix, Section B, we present another case study that involves the creation of random financial networks and the introduction of appropriate perturbations to simulate financial crime through the use of intermediaries. The use of FlowSeries in this controlled environment allows a more precise evaluation of the tool's capabilities, highlighting both its strengths and limitations.

Case study 1: money flows between two countries

The first case study focuses on a macro-scale investigation of transaction flows between countries. Specifically, it illustrates how an AFC analyst can examine transaction flows between two countries. The reference period for this analysis spans the months immediately following the outbreak of war in Ukraine and precedes the full implementation of economic sanctions against Russia.

Figure 6 shows the weekly aggregated time series of direct transaction amounts from BICs in country C2 to BICs in country C1. The dashed gray line indicates the beginning of the war in Ukraine. Over the observed period, the trend remains relatively stable and the imposition of sanctions does not appear to have a significant impact on the volume of direct transfers from C2 to C1. Figure 7a shows the time series of $w(Flow^3(C2, C1))$, as

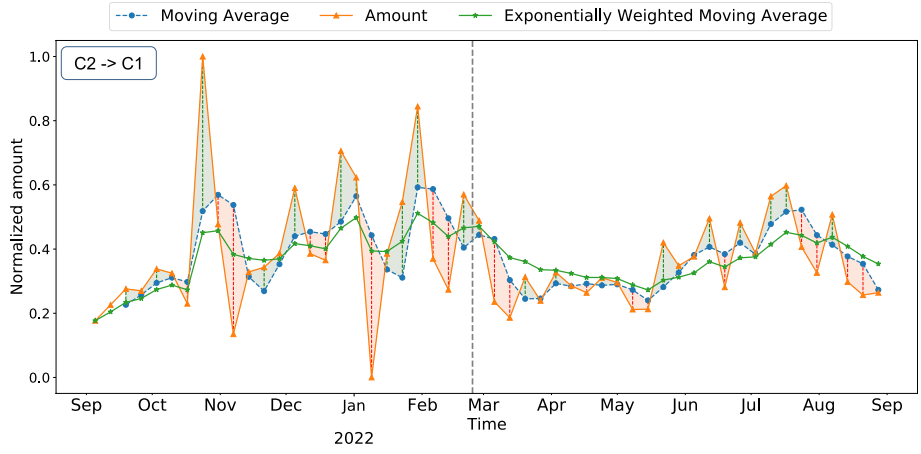

Fig. 6 Direct transactions from the BICs of country C1 to the BICs of country C2. The vertical gray dotted line represents the begin of the war in Ukraine (February 24th, 2022)

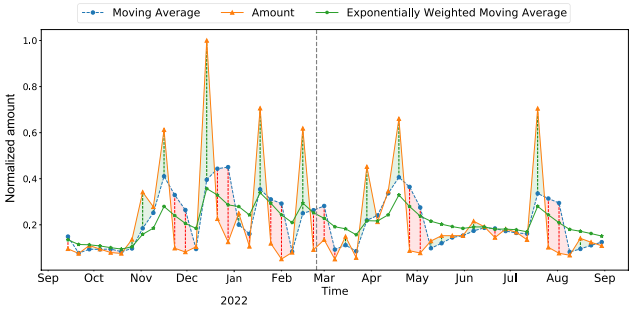

(a) Figure

Node	$\Delta_{w,t}$
IC1	0.427
IC2	0.436
IC3	0.493
IC4	0.535
IC6	0.636
IC5	0.665

(b) Table

Fig. 7 **a** The weight of the transaction flow $Flow^3(C2, C1)$, i.e., the hypothetical maximum amount of money sent from C2 to C1 through several payment lines, each with a maximum of one intermediary. Table **b** lists the intermediaries in $Flow^3(C2, C1)$ that show the largest deviation Δ_w, as defined in 4

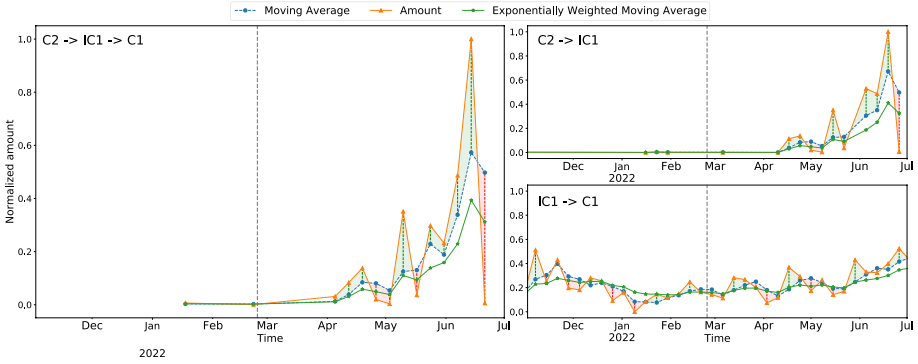

Fig. 8 Transaction flow from C2 to C1 through IC1

defined in Eq. (3), which estimates the hypothetical maximum amount transferred from C2 BICs to C1 BICs through transaction paths involving at most one intermediary. Although both the time series in Figs. 6 and 7a show fluctuations—particularly before the outbreak of the war—no clear trends or critical deviations are immediately apparent. At this point, the AFC analyst can further investigate the financial relationship between C2 and C1 by analyzing the main intermediaries involved in the transaction flow $Flow^3(C2, C1)$. There are a total of 86 transaction paths with a maximum length of 3 connecting C2 and C1. Table 7b identifies the intermediary entities that have the highest deviations Δ_w after February 24th, as defined in Eq.(4).

Figures 8 and 9 show the transaction flow from C2 to C1 through the intermediary countries IC1 and IC3, respectively. The transaction flows involving the other countries reported in Table 7b are shown in the appendix (Fig. 14).

On the left side of Figs. 8 and 9, the flow weights are computed based on weekly aggregated networks, while the right side displays the individual transaction edges involved in the flows, specifically, the links between C2 and IC1, and between IC1 and C1 in Fig. 8, as well as C2 to IC3 and IC3 to C1 in Fig. 9. In both figures, the transaction flow weight $w(Flow^3(C2, C1))$ shows a sharp and anomalous increase starting in May 2022. By analyzing the individual transaction edges within these flows, AFC analysts can gain a more detailed understanding of the mechanisms driving this increase. The analytical pipeline proposed in this study provides analysts with efficient tools to identify potentially correlated transaction patterns between groups of countries. However, establishing a causal

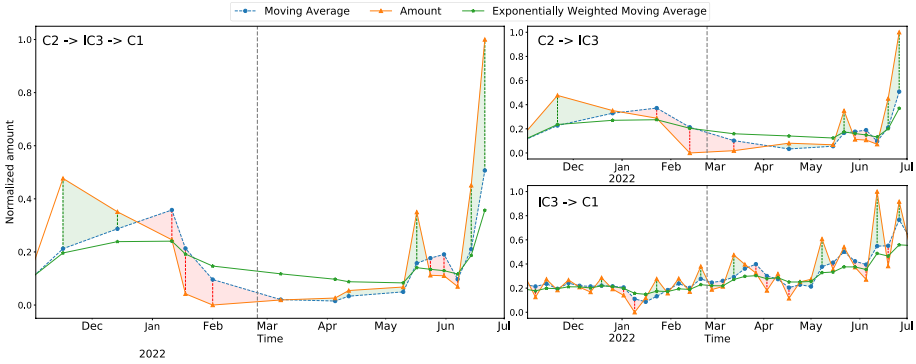

Fig. 9 Transaction flow from C2 to C1 through IC3

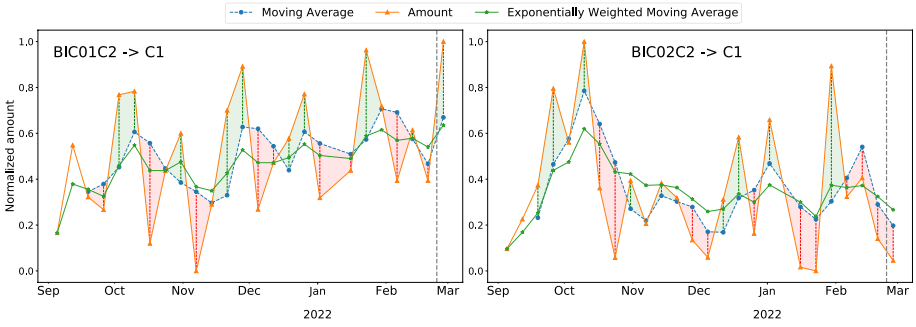

Fig. 10 The two figures show the amount of money sent from two C2 BICs (*BIC01C2* on the left, *BIC02C2* on the right) to C1. Both BICs stop sending money to country C1 the week after the outbreak of the war in Ukraine on February 24th, 2022

relationship requires an in-depth investigation and verification of individual transactions. It is important to contextualize these findings in the light of wider geopolitical developments. Following the imposition of economic sanctions on Russia, several journalistic and institutional investigations reported that entities in certain countries acted as intermediaries to circumvent restrictions, thereby facilitating the flow of goods and capital (AP 2023; Wagner 2023; Times 2023; Warrick 2023).

Case study 2: in-depth investigation of two malicious nodes

The second case study focuses on two BICs flagged as suspicious by the AFC Digital Hub group. These entities were identified using the anomaly detection pipeline introduced by Vilella et al. (2025). For confidentiality reasons, both BICs have been anonymised and are hereafter referred to as *BIC01C2* and *BIC02C2*.

Figure 10 shows the wire transactions from *BIC01C2* and *BIC02C2* to country C1. In a G^B network, each node represents a specific BIC and is associated with attributes such as the country of the bank branch and the corresponding financial risk level. Transaction flows in this network can be examined at different levels of aggregation: for example, while the source node may be a single BIC, the intermediate or destination nodes may represent entire countries or clusters of countries with common characteristics.

In the examples shown in Fig. 10, the source nodes are individual BICs, while the target is represented at the country level. In the week immediately following the outbreak of war in Ukraine, both *BIC01C2* and *BIC02C2* stopped direct transactions to C1. As in the previous case study, the AFC analyst can analyze the flows $Flow^3(\mathrm{BIC01C2}, \mathrm{C1})$

and $Flow^3(\text{BIC02C2}, \text{C1})$. The analysis shows that after the direct transfers stopped, both BICs began to redirect funds to C1 through a single intermediary—another BIC in the same country, identified as *BIC03C2*, which had not previously been identified in the investigation.

Figure 11 (left) provides a first exploratory analysis of the weight of transaction flow from *BIC03C2* to C1, specifically through countries classified as medium or high financial risk. Following the outbreak of war in Ukraine, the weight of transaction flow show a marked upward trend, peaking in June 2023. In the context of international economic sanctions against Russia, this increase may raise suspicions and warrant further investigation. As in previous examples, the FlowSeries algorithm allows analysts to identify and investigate the key intermediaries involved in these flows. Three such cases of interest are described below.

As in Use Case 1, the analyst can continue the investigation by focusing on the intermediaries with the largest deviation Δ_w. Figure 11 (right) shows the weight of the transaction flow from *BIC03C2* to BICs in C1 through intermediaries located in the country with the highest Δ_w, here referred to as IC7. While the overall weight of these flows remains relatively low, two distinct anomalies—sharp peaks in May and June 2022—stand out. These flows are routed through only five BICs acting as intermediaries. At this point, the AFC analyst can further investigate these specific BICs by delving into the details of the individual transactions that make up the flow to assess their legitimacy and uncover any potential indicators of sanctions evasion.

Figure 12 shows two additional examples of transaction flows from *BIC03C2* to BICs in C1 routed through other intermediary countries with the largest Δ_w. In both cases, the weight of the transaction flow shows an increasing trend after the outbreak of the war in Ukraine, peaking in May and June.

It is worth noting that an analyst could identify the individual edges forming the flow $Flow^3(BIC03C2, C1)$ as anomalies without using FlowSeries. However, these use cases emphasize how FlowSeries supports and facilitates the investigative process by enabling the simultaneous identification of multiple suspicious intermediaries that may be associated with the same financial crime activity.

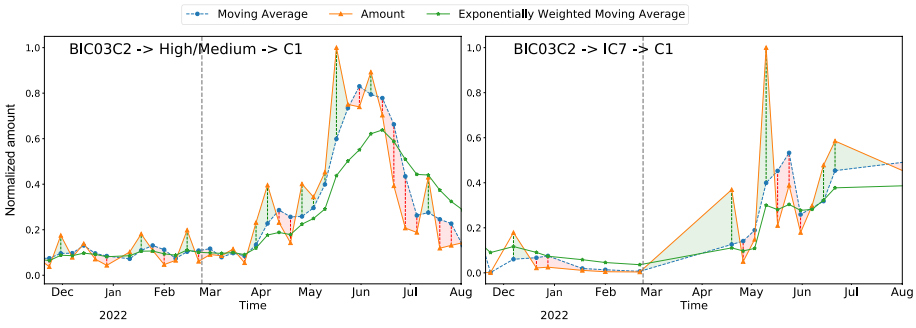

Fig. 11 Both figures illustrate transaction flows originating from BIC *BIC03C2* and directed towards C1 BICs. The figure on the left shows flows involving intermediaries located in medium or high risk countries, while the figure on the right shows flows involving only the anonymised intermediary IC7

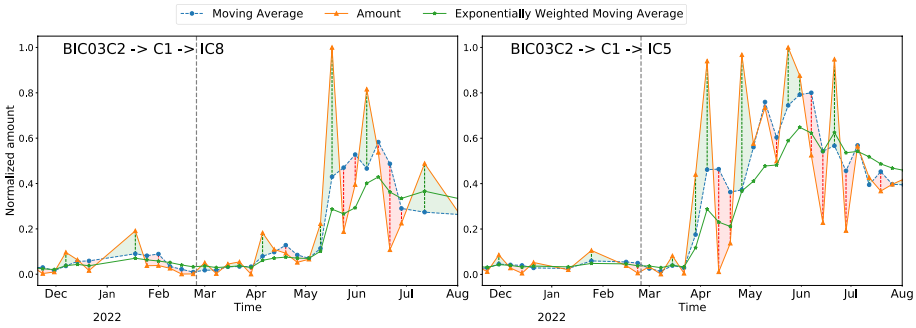

Fig. 12 Both figures illustrate a transaction flow from the BIC *BIC03C2* to the BICs of country C1, albeit through different intermediary countries

Limitations

One of the main limitations of FlowSeries is the inherent uncertainty in interpreting whether a detected transaction flow between two nodes represents an attempt to conceal illicit financial activity. The flow weight measures the maximum amount that can be transferred between nodes through various intermediaries and payment channels. While historical time series analysis of these weights can reveal anomalies and abrupt changes in behavior, definitive conclusions about financial crime require the judgment of an AFC analyst, supplemented by contextual knowledge and external data sources.

A second limitation is that FlowSeries requires the analyst to specify a root node as the starting point of the investigation. However, this limitation does not preclude the discovery of previously unknown suspicious nodes, even at different levels of aggregation, as illustrated in the second use case of Sect. 5. Once identified, these new nodes can serve as new starting points, allowing analysts to iteratively reapply FlowSeries and progressively expand the scope of the investigation.

Due to these limitations, the effective use of FlowSeries relies on the expertise of the analyst—both to define meaningful starting points and to accurately interpret the results. Importantly, FlowSeries is not a black box model for the automatic detection of financial crime or anomalies. Rather, it is a structured analytical tool designed to support and enhance established investigative workflows within AFC units. When used in conjunction with traditional investigative methods—as demonstrated in the real-world case studies—FlowSeries provides a systematic and scalable framework for analyzing complex transaction flows and uncovering potentially illicit patterns.

Conclusions and future work

In this paper, we present FlowSeries, an exploration method developed to support AFC investigations. FlowSeries enhances the analytical capabilities of AFC analysts by assisting in the identification of suspicious entities and anomalous transaction patterns. Crucially, it is not a fully autonomous or black box anomaly detection system, but a structured framework that relies on expert guidance and interpretation. By streamlining the analysis of complex financial networks, FlowSeries enables investigators to efficiently trace transaction flows and uncover hidden relationships that may suggest illicit activity.

Two key challenges hinder the adoption of fully automated systems in AFC: the need for interpretability and the scarcity of labeled data for model training. Interpretability is critical for the formal reporting of suspicious anomalies to regulatory authorities. The lack of labeled data is largely due to the highly confidential nature of financial

transactions, as well as the sheer scale involved—large European banks process millions of transactions per week. In addition, annotating such data is inherently difficult, and previous research (García and Mateos 2021; Garcia-Bedoya et al. 2020; Liu et al. 2020) has shown that criminal transaction patterns are highly complex and diverse, making systematic definition and labeling a challenge. While FlowSeries requires an analyst to select an initial point of investigation, the use cases in Sect. 5 show that an exploratory analysis can begin at the country level and, through iterative refinement, be progressively narrowed down to specific BICs or even individual suspicious transactions.

As reported in Section A of the appendix, FlowSeries has been successfully integrated into the investigative workflow of an AFC unit, with its effectiveness proven through real-world data analysis. However, several aspects could be further refined to increase its effectiveness. A key area for improvement is the refinement of the filtering mechanisms used to refine the paths identified by Algorithm 1. More sophisticated edge filtering techniques—possibly incorporating additional contextual data available to AFC analysts within financial institutions—could reduce the number of paths to investigate and improve the overall efficiency of the investigation. Unfortunately, such data was not available to our research team due to privacy and security constraints.

Another promising direction concerns anomaly detection in the time series of transaction flow weights, as discussed in Sect. 3.4. Since flows between two BICs can consist of thousands of individual paths, it is impractical to apply computationally intensive anomaly detection algorithms to each path. A more scalable approach would be to identify a smaller, high-risk subset of paths and monitor them periodically with advanced anomaly detection models.

Finally, upstream node filtering could further streamline the analysis by excluding nodes that are irrelevant to specific investigations. In particular, highly centralized nodes can increase computational complexity without significantly aiding the detection of suspicious behavior. Access to enriched node-specific metadata—not available to our research team—could enable more informed and targeted filtering, thereby improving investigative efficiency.

Appendix A: AFC investigative practice and experts' follow up on FlowSeries

To fully assess the implementation choices behind FlowSeries, it is essential to understand key aspects of investigative practices within the AFC units of large banking institutions. To this end, we posed the following three questions to the Anti-Financial Crime Digital Hub regarding AFC operations and the practical application of FlowSeries:

- Could you briefly describe the role of the AFC expert, including how cases are identified, to whom they are reported, and what kind of feedback is received from the authorities?
- What are the most critical aspects of AFC investigations where human in the loop remains essential, despite advancements in automation?
- Has the use of this method improved your investigative approach? If so, how?

The following is the response provided by the Anti-Financial Crime Digital Hub:

"FlowSeries is a solution designed to detect anomalous financial flows through a top-down search methodology. Unlike traditional bottom-up approaches that focus on identifying anomalies at the level of individual transactions or accounts, FlowSeries aggregates

transactions and analyzes them across higher-level units of analysis—such as countries or financial institution corridors. This innovative perspective allows the detection of suspicious patterns that may only emerge at a broader, systemic level.

While the bottom-up approach is effective at identifying anomalies at the transaction or account level, determining whether these anomalies are actually suspicious requires the integration of contextual information from multiple–and often unpredictable–domains. At present, this step necessitates a human-in-the-loop and human-in-command approach, relying on the expertise of AFC professionals to examine customer registries, account statements, media sources, and other data in order to complete the case investigation. Once an activity is confirmed to be suspicious, it is reported to the relevant national authority–typically the country's Financial Intelligence Unit (FIU)–via a Suspicious Activity Report (SAR), submitted in a standardized data format as required by the FATF guidelines adopted by most jurisdictions worldwide. At this point, the formal obligations of financial institutions conclude. However, depending on the jurisdiction and case, the FIU may request additional information or notify of case dismissal. In many instances, no feedback is provided to the reporting institutions, even when the reported activity later appears in public media coverage during legal proceedings.

AFC investigations are inherently unique, which limits the applicability of detailed procedures; thus, only general, high-level guidance can be provided to support investigators. Furthermore, criminal ingenuity continuously evolves, renewing the typologies of what were once well-defined risk categories. As a result, automation is currently essential to keep pace with adversaries. However, it still falls short of enabling a shift from a reactive posture to a fully prescriptive and proactive approach–one that must also account for individual subjectivity, which can be a source of either excellence or mediocrity in following the right investigative leads.

Some egregious international cases have demonstrated that a top-down analysis of financial flows is crucial to avoid overlooking emerging yet elusive maneuvers. At the same time, deploying such advanced analytics effectively requires skilled professionals–something particularly challenging in today's context of limited resources dedicated to anti-financial crime intelligence. In this light, FlowSeries represents a prototypical initiative aimed at exploring and validating a technological trajectory that will be industrialized within ISP in the near future."

Appendix B: Money flows in synthetic networks

This section explores the application of FlowSeries to synthetically generated data. This controlled experimental framework allows for the validation of the specific conditions under which an AFC analyst can derive significant benefits from the implementation FlowSeries. Additionally, the use of synthetic data ensures a more accurate and reliable assessment of the flow exploration model's performance. The code to reproduce the networks and the results presented in this section is available in a public GitHub repository (Capozzi 2025).

Our approach involves recreating networks similar to those analyzed in previous case studies and introducing perturbations that modify both the structure and the edge weights. Specifically, a perturbation targets a small group of randomly selected nodes (N_1), toward which another group of nodes (N_2) reduces the volume of direct transactions, simulating the imposition of economic sanctions. A third randomly selected group

of nodes (N_3) is introduced to model intermediary-mediated transaction flows from N_2 to N_1. This process is implemented as follows:

1. Initialization: Create a network that follows the Barabási-Albert model (given the out-degree distributions shown in Fig. 5). Add other edges between pairs of nodes with probability p to create a directed graph class that can store multiedges.
2. Stable Perturbation: In this phase from t_1 to t_g, the weight of each edge is updated as follows:

$$e_{i,j}^{t+1} = \text{rand.uniform}(0.8, 1) \cdot e_{i,j}^{t}, \tag{B1}$$

where $\text{rand.uniform}(0.8, 1)$ represents a random number generated from the uniform distribution over the interval $[0.8, 1]$.

3. Strong Perturbation: At time t_g, a set of random nodes N_1 and N_2 are selected, such that

$$\forall n_2 \in N_2, \quad \forall n_1 \in N_1, \quad e_{n_2,n_1} \in E, \tag{B2}$$

indicating the presence of edges between each node in N_2 and every node in N_1. Initially, the weight of all edges connecting nodes in N_2 to nodes in N_1 is reduced by half, and the average reduction for each edges is computed and recorded in μ_r:

$$e_{n_2,n_1}^{t_g} = \frac{1}{2} e_{n_2,n_1}^{t_{g-1}}, \tag{B3}$$

For instance, if at time t_{g-1}, two edges connect nodes i and j, with weights of 6 and 10, respectively, at time t_g, the weights of the edges will be adjusted to 3 and 5, with the average reduction value calculated as $\frac{3+5}{2}$. Subsequently, a third set of random nodes N_3 is selected such that each node in N_3 has at least one incoming edge from all nodes in N_2 and at least one outgoing edge to all nodes in N_1, as shown in Eq.(B4).

$$\begin{aligned} &\forall n_3 \in N_3, \quad \forall n_2 \in N_2, \quad \exists e_{n_2,n_3} \in E, \\ &\quad \forall n_1 \in N_1, \quad \exists e_{n_3,n_1} \in E \end{aligned} \tag{B4}$$

At each time t, the weight of the edges from N_2 to N_1 through N_3 is increased according to the following equation:

$$\begin{aligned} e_{n_2,n_1}^{t+1} = e_{n_2,n_1}^{t} + \text{random.uniform} \Bigg(\frac{e_{n_2,n_1}^{t}}{2}, \\ 2 \cdot e_{n_2,n_1}^{t} \Bigg) + \text{random.uniform}(0, \mu_r) \cdot 10 \end{aligned}$$

Where e_{n_2,n_1}^{t} represents the weight of the edge from node n_2 to node n_1 at time t. The first random term $\text{random.uniform}\left(\frac{e_{n_2,n_1}^{t}}{2}, 2 \cdot e_{n_2,n_1}^{t}\right)$ represents a random increase in the edge weight, bounded between half and double the current weight. The second random term $\text{random.uniform}(0, \mu_r) \cdot 10$ adds a random value based on the average reduction of the directed edge μ_r.The experiment is initialized by generating a network of 200 nodes using the Barabási-Albert model. The network underwent a stable perturbation up to the time step t_{13}, followed by a strong perturbation starting at t_{14}. We set the number of nodes to 200 because, as illustrated in Fig. 4a, the number of nodes in the real

networks aggregated by country varies between 194 and 207. In this experiment, sets N_1 and N_2 each contain a single node, while N_3 contains three nodes. Although there are multiple paths between N_1 and N_2, three were randomly selected for analysis. Consequently, the number of synthetically generated suspicious paths in the networks is three.

Following the application of the steps outlined in Sect. 3.5, FlowSeries identifies five potential paths exhibiting a negative $\Delta_{e,t}$, including all three paths traversing the N_3 nodes. Figure 13a shows the time series of the weight $w(Flow^3(N2, N1))$(along with its moving average with a sliding window of 4) of one of the suspicious flows from the node in N_2 to a node in N_1. This is one of the three synthetically generated suspect flows that are correctly identified. The values on the y axis are normalized with the maximum value in the time series. The vertical gray line marks the transition from a stable perturbation to a strong perturbation. The corresponding values of $\Delta_{w,t}$ are represented in Fig. 13b, showing a sharp increase immediately after the emergence of the strong perturbation. Subsequently, as the network stabilizes, the values return to a more gradual trend, reflecting the redistribution of weight from the direct edge to the intermediate-mediated paths. This example is particularly noteworthy, as it highlights the importance of examining the entire time series to characterize a state transition, which leads to a redistribution of weights from a direct edge to an indirect flow.

Due to inherent random variations in the weight update process during the perturbation phases, we may observe fluctuations that could be mistakenly interpreted as suspicious flow patterns. Within this controlled setting, the case shown in Fig. 13c, d represents a false positive. Specifically, the figure on the left shows the weight of the flow from the node in N_1 to other nodes in the network that do not belong to N_2 or N_3. The observed increase in weight is random noise, and not a consequence of the strong perturbation introduced at t_{14}. The corresponding $\Delta_{w,t}$ value is shown on the right, exhibiting a sharp peak followed by stabilization, similar to the pattern in Fig. 13b. However, in this instance, the increase does not coincide with the initiation of the strong perturbation.

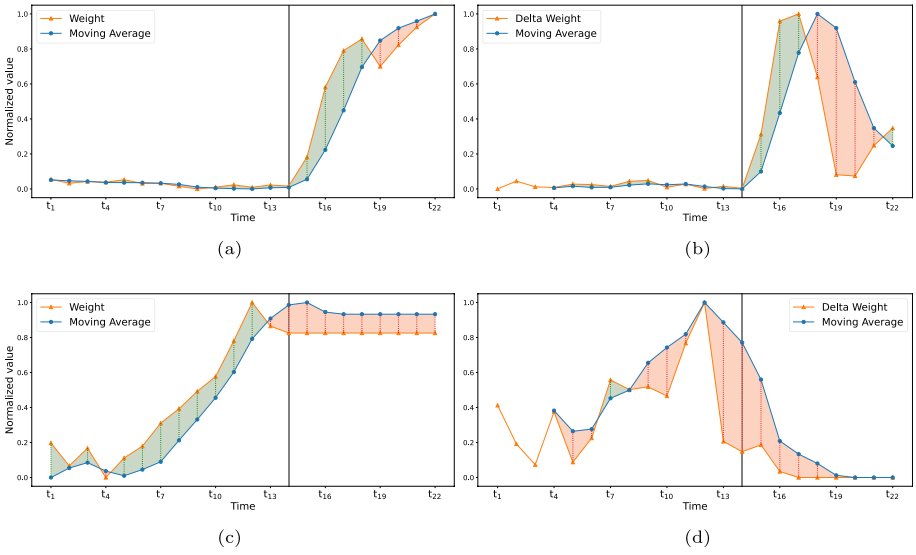

Fig. 13 The figures show the weight $w(Flow^3)$ and the Δ_w of a true positive suspicious flow (**a**, **b**), and a false positive suspicious flow (**c**, **d**). Gray vertical lines represent the introduction of the strong perturbation

Appendix C: Transaction flows from C2 to C1

See Fig. 14.

Appendix D: Benchmarks

The benchmarks of Algorithm 1 are presented in this section. The tests are performed on transaction networks aggregated by country and by BIC. The maximum path length is 3, and the networks are aggregated in time on a weekly basis. We implemented the algo-

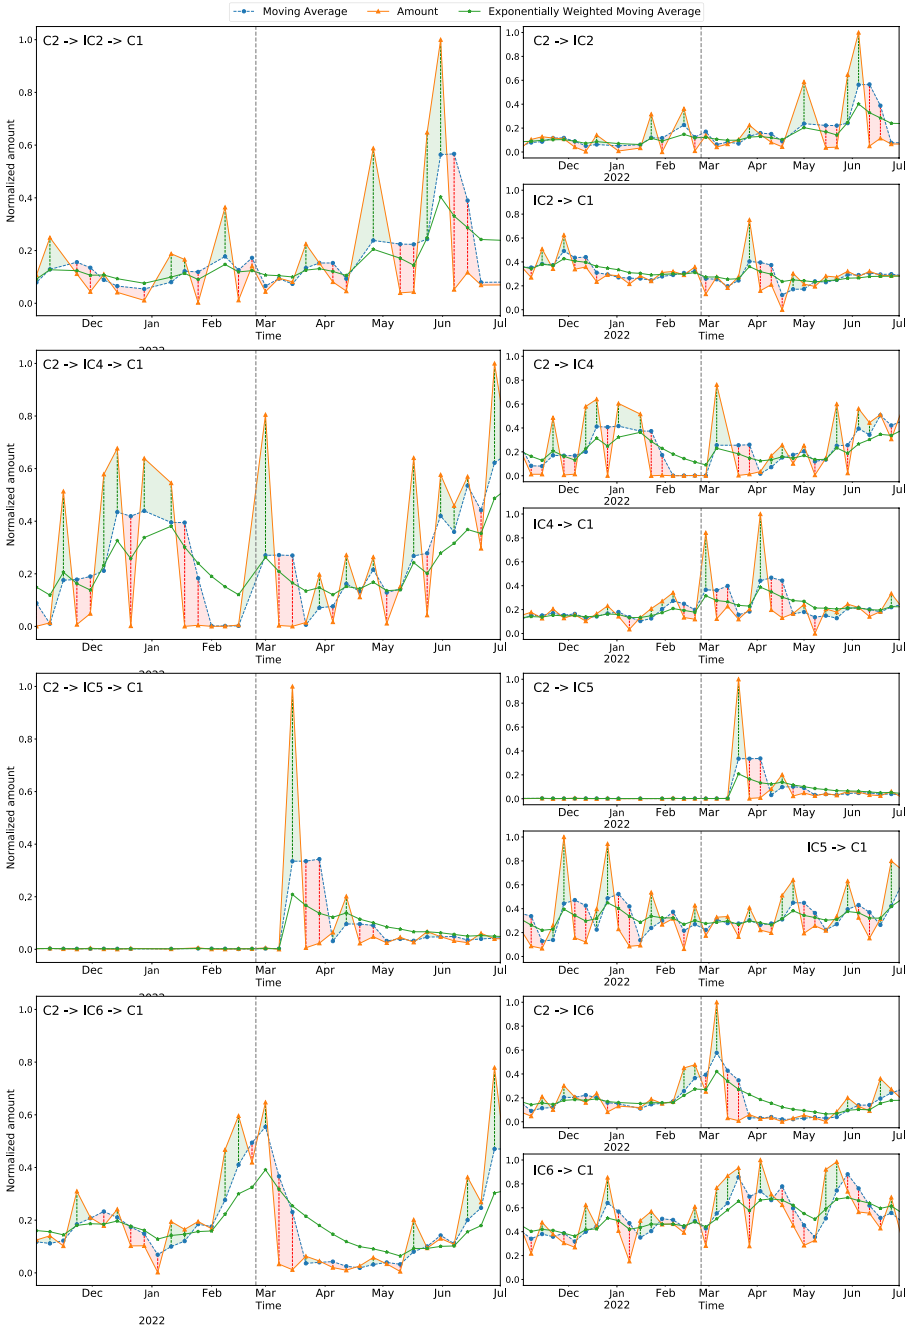

Fig. 14 Transaction flows involving intermediaries reported in Table 7b

rithm in Python and the tests are performed on a machine with a CPU Intel(R) Xeon(R) Gold 6252 2.10GHz and 128GB of RAM.

Since the execution time of Algorithm 1 is dominated by the input node, for each temporal aggregation T we select six nodes and run a benchmark for each of them. One of the six selected nodes is always the node with the largest out-degree in the interval T, the other 5 are randomly selected among the nodes with at least one outgoing edge.

The benchmark tests on the cross-country network are shown in Fig. 15a and those on the cross-BIC network are shown in Fig. 15b. Table 1 shows the average out-degree of the 5 randomly selected nodes, the largest out-degree of the network, the number of nodes, and the number of edges.

The average out-degree of randomly selected nodes is comparable between the network of BICs and the networks of countries. The out-degree of the cross-BIC network is orders of magnitude larger than that of the cross-country network. This is also observed by the longer tail of the out-degree distribution of G_B(Fig. 5b) compared to that of G_C(Fig. 5a). Consequently, as shown in Fig. 15, the maximum computation times for networks between countries are about 8 s, and for networks between BICs about 400 s.

Table 1 The table shows statistics about the transaction networks and the starting nodes used for the benchmarks of Algorithm 1

Date	Average out-degree		Top node out-degree		Number of nodes		Number of edges	
	Country	BIC	Country	BIC	Country	BIC	Country	BIC
2021-09-01	23.0	19.2	201	3779	213	5740	4385	55,870
2021-10-01	22.6	6.4	199	3711	213	5730	4383	56,224
2021-11-01	7.2	3.8	204	3802	214	5790	4560	58,170
2021-12-01	62.8	6.8	206	3810	214	5769	4542	59,363
2022-01-01	24.4	3.6	205	3680	212	5591	4319	55,509
2022-02-01	16.0	1.8	206	3699	212	5635	4309	56,271
2022-03-01	23.8	9.0	205	3755	210	5731	4518	58,630
2022-04-01	20.8	8.4	203	3704	210	5659	4509	56,912
2022-05-01	15.4	3.6	204	3753	211	5689	4479	57,595
2022-06-01	32.2	13.0	205	3720	212	5693	4651	58,835
2022-07-01	10.8	3.2	203	3746	212	5614	4631	58,772
2022-08-01	9.2	90.4	198	3701	209	5578	4649	58,049
2022-09-01	20.2	4.6	204	3785	211	5660	4750	60,266
2022-10-01	10.4	3.4	200	3774	212	5644	4697	60,353
2022-11-01	1.8	12.6	202	3767	212	5637	4723	61,142

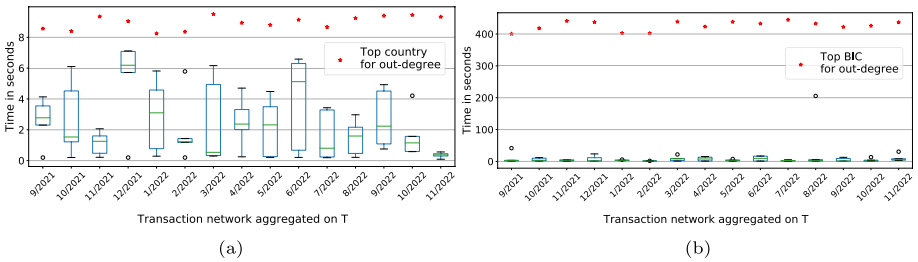

Fig. 15 Execution times in seconds of algorithm 1 computed on transaction networks between countries (on the left) and between BICs (on the right)

Author contributions
AC: Conceptualization, Methodology, Software, Validation, Visualization, Writing—Original Draft. SV: Methodology, Software, Validation, Writing—Review & Editing. DM, MF, VR, and SR: Resources, Data Curation, Validation. GR: Methodology, Supervision, Writing—Review & Editing.

Funding
Open access funding provided by Swiss Federal Institute of Technology Zurich. This research has been funded by AFC Digital Hub (Anti Financial Crime Digital Hub) consortium, whose members are Intesa Sanpaolo Innovation Center, University of Turin, Polytechnic University of Turin, and CENTAI. SV also acknowledges funding by PRIN 2022 (PNRR M4C2) the European Union—Next Generation EU, Mission 4 Component 2—CUP C53D23005810006.

Availability of data and materials
Real data of cross-country financial transactions is provided by Intesa Sanpaolo (ISP) and AFC Digital Hub. For more information, write to ***adh@pec.afcdigitalhub.com***.

Declarations

Competing interests
The author(s) declare(s) that they have no competing interests.

Received: 31 March 2025 / Accepted: 3 June 2025
Published online: 10 July 2025

References

AP (2023) New US sanctions target people and companies in Turkey, Georgia and Russia. Accessed 14 Sept 2023. https://www.euronews.com/business/2023/09/14/new-us-sanctions-target-people-and-companies-in-turkey-georgia-and-russia
Bank EC (2021) Payments statistics: 2020. Accessed 28 Mar 2025. https://www.ecb.europa.eu/press/stats/paysec/html/ecb.pis20205d0ea9dfa5.en.html
Bolton RJ, Hand DJ (2002) Statistical fraud detection: a review. Stat Sci 17(3):235–249
Box GE, Jenkins GM, Reinsel GC, Ljung GM (2015) Time series analysis: forecasting and control. Wiley, Hoboken
Caldarelli G, Battiston S, Garlaschelli D, Catanzaro M (2004) Emergence of complexity in financial networks, vol 650, pp 399–423
Capozzi A, Vilella S, Moncalvo D, Fornasiero M, Ricci V, Ronchiadin S, Ruffo G (2024) Flowseries: anomaly detection in financial transaction flows. In: International conference on complex networks and their applications. Springer, pp 29–40
Capozzi A (2025) FlowSeries: flow analysis on financial networks—GitHub Repository. Accessed 31 Mar 2025. https://github.com/PotenteOpossum/FlowSeries
Chen Z, Le DV-K, Teoh E, Nazir A, Karuppiah E, Lam K (2018) Machine learning techniques for anti-money laundering (AML) solutions in suspicious transaction detection: a review. Knowl Inf Syst. https://doi.org/10.1007/s10115-017-1144-z
Didimo W, Liotta G, Montecchiani F, Palladino P (2011) An advanced network visualization system for financial crime detection, pp. 203–210. https://doi.org/10.1109/PACIFICVIS.2011.5742391
Faggini M, Bruno B, Parziale A (2019) Crises in economic complex networks: black swans or dragon kings? Econ Anal Policy 62:105–115. https://doi.org/10.1016/j.eap.2019.01.009
Fronzetti Colladon A, Remondi E (2017) Using social network analysis to prevent money laundering. Expert Syst Appl 67:49–58. https://doi.org/10.1016/j.eswa.2016.09.029
Garcia-Bedoya O, Granados O, Burgos J (2020) Ai against money laundering networks: the Colombian case. J Money Laund Control Ahead-of-Print. https://doi.org/10.1108/JMLC-04-2020-0033
García IG, Mateos A (2021) Use of social network analysis for tax control in Spain. Hacienda Pública Española / Rev Public Econ 239(4):159–197
Huang D, Mu D, Yang L, Cai X (2018) Codetect: financial fraud detection with anomaly feature detection. IEEE Access 6:19161–19174. https://doi.org/10.1109/ACCESS.2018.2816564
Kauê Dal'Maso Peron T, Fontoura Costa L, Rodrigues FA (2012) The structure and resilience of financial market networks. Chaos Interdiscip J Nonlinear Sci 22(1):013117. https://doi.org/10.1063/1.3683467
Kumar Dubey A, Kumar A, García-Díaz V, Kumar Sharma A, Kanhaiya K (2021) Study and analysis of sarima and lstm in forecasting time series data. Sustain Energy Technol Assess 47:101474. https://doi.org/10.1016/j.seta.2021.101474
Kute DV, Pradhan B, Shukla N, Alamri A (2021) Deep learning and explainable artificial intelligence techniques applied for detecting money laundering—a critical review. IEEE Access 9:82300–82317. https://doi.org/10.1109/ACCESS.2021.3086230
Lee K-M, Yang J-S, Kim G, Lee J, Goh K-I, Kim I (2011) Impact of the topology of global macroeconomic network on the spreading of economic crises. PLoS ONE 6(3):1–11. https://doi.org/10.1371/journal.pone.0018443
Leila A (2011) Global crises: a network perspective on the economic integration. J Econ Integr 26(2):197–216x
Lillo F, Moro E, Vaglica G, Mantegna RN (2008) Specialization and herding behavior of trading firms in a financial market. New J Phys 10(4):043019. https://doi.org/10.1088/1367-2630/10/4/043019
Lin D, Wu J, Yuan Q, Zheng Z (2020) Modeling and understanding ethereum transaction records via a complex network approach. IEEE Trans Circuits Syst II Express Briefs 67(11):2737–2741. https://doi.org/10.1109/TCSII.2020.2968376
Liu Z, Zhou D, Zhu Y, Gu J, He J (2020) Towards fine-grained temporal network representation via time-reinforced random walk. In: Proceedings of the AAAI conference on artificial intelligence, vol 34, No. 04, pp 4973–4980. https://doi.org/10.1609/aaai.v34i04.5936
López L, Almendral AJ, Sanjuáin MAF (2003) Complex networks and the www market. Physica A Stat Mech Appl 324(3):754–758. https://doi.org/10.1016/S0378-4371(02)01867-8
Mu G-H, Zhou W-X, Chen W, Kertész J (2010) Order flow dynamics around extreme price changes on an emerging stock market. New J Phys 12(7):075037. https://doi.org/10.1088/1367-2630/12/7/075037

Pan J (2022) Deep set classifier for financial forensics: an application to detect money laundering. arXiv. https://doi.org/10.48550/ARXIV.2207.07863
Papadimitriou T, Gogas P, Gkatzoglou F (2020) The evolution of the cryptocurrencies market: a complex networks approach. J Comput Appl Math 376:112831. https://doi.org/10.1016/j.cam.2020.112831
Phua C, Lee VCS, Smith-Miles K, Gayler RW (2010) A comprehensive survey of data mining-based fraud detection research. CoRR arXiv:1009.6119
Rauch JE (2001) Business and social networks in international trade. J Econ Lit 39(4):1177–1203. https://doi.org/10.1257/jel.39.4.1177
Semeraro A, Tambuscio M, Ronchiadin S, Puma L, Ruffo G (2020) Structural inequalities emerging from a large wire transfers network. Appl Netw Sci. https://doi.org/10.1007/s41109-020-00314-x
Serena L, Ferretti S, D'Angelo G (2022) Cryptocurrencies activity as a complex network: Analysis of transactions graphs. Peer-to-Peer Netw Appl. https://doi.org/10.1007/s12083-021-01220-4
Shun J, Dhulipala L, Blelloch GE (2015) Smaller and faster: parallel processing of compressed graphs with ligra+. In: 2015 data compression conference, pp 403–412. https://doi.org/10.1109/DCC.2015.8
Tarjan R (1971) Depth-first search and linear graph algorithms. In: 12th annual symposium on switching and automata Theory (swat 1971), pp 114–121. https://doi.org/10.1109/SWAT.1971.10
Times F (2023) Armenia: on the new silk road for goods to sanctions-hit Russia. Accessed 18 July 2023. https://www.ft.com/content/0fc846f7-aac8-4a34-a7dd-3b0615bce983
Vandewalle N, Brisbois F, Tordoir X (2001) Non-random topology of stock markets. Quant Finance 1(3):372–374. https://doi.org/10.1088/1469-7688/1/3/308
Vilella S, Lupi A, Fornasiero M, Moncalvo D, Ricci V, Ronchiadin S, Ruffo G (2025) Weirdnodes: centrality based anomaly detection on temporal networks for the anti-financial crime domain. Appl Netw Sci 10:14. https://doi.org/10.1007/s41109-025-00702-1
Wagner R (2023) Higher German exports to Russia's neighbours fuel sanctions evasion fears. Accessed 17 May 2023. https://www.reuters.com/world/german-exports-russias-neighbours-fuel-sanctions-evasion-fears-2023-05-16/
Warrick J (2023) In Central Asia, a hidden pipeline supplies Russia with banned tech. Accessed 18 July 2023. https://www.washingtonpost.com/national-security/2023/07/18/russia-sanctions-weapons-china-drones/
Weber M, Chen J, Suzumura T, Pareja A, Ma T, Kanezashi H, Kaler T, Leiserson CE, Schardl TB (2018) Scalable graph learning for anti-money laundering: a first look. CoRR arxiv:1812.00076
Yan X-G, Xie C, Wang G-J (2014) The stability of financial market networks. Europhys Lett 107(4):48002. https://doi.org/10.1209/0295-5075/107/48002
šubelj L, Furlan-Bajec M (2011) An expert system for detecting automobile insurance fraud using social network analysis. Expert Syst Appl 38(1):1039–1052. https://doi.org/10.1016/j.eswa.2010.07.143

Publisher's Note