

Detecting Trade-Based Money Laundering: A Literature Review of Red-Flag Frameworks and Network-Analytic Detection Methods

Malki Senevirathne

Independent Researcher, United Kingdom

Abstract: *Trade-based money laundering (TBML) is widely regarded as one of the hardest forms of money laundering to detect, because it hides inside ordinary import-export documentation rather than a suspicious financial transfer. Regulators have responded largely with red-flag indicator lists, while separate and much larger literatures have applied network analysis, machine learning and customs-fraud analytics to adjacent detection problems. This paper reviews these strands to establish what is known about detecting TBML and where the evidence base is thin. Thirty-five academic papers were identified through seven structured searches of an academic database, supplemented by four regulatory and institutional publications, and synthesised thematically. The review finds that regulatory red-flag taxonomies are detailed and well established but designed for manual, document-level review; that TBML-specific peer-reviewed research remains genuinely scarce, a finding independently confirmed by three prior systematic and bibliometric reviews of the field; and that a large, mature and well-validated customs-fraud-detection literature exists on directly adjacent problems, such as under-invoicing and misclassification, yet is rarely cross-referenced by TBML-specific research. The general anti-money-laundering network-analysis and machine-learning literature reports strong detection performance but is drawn almost entirely from financial transaction or cryptocurrency data rather than trade documentation. The review closes by setting out this gap, a research agenda for closing it, and the methodological limitations of the search process used here.*

Keywords: Trade-based money laundering, Literature review, Red flags, Network analysis, Customs fraud detection

1. Introduction

International trade moves goods worth trillions of dollars every year, and that scale is precisely what makes it useful to criminals. Trade-based money laundering (TBML) uses ordinary import-export transactions, invoices, letters of credit and shipping documents to move value across borders while presenting the appearance of legitimate commerce. Because the underlying activity looks like trade rather than a suspicious wire transfer, TBML is widely regarded as one of the hardest forms of money laundering to detect [1].

The scale of the problem is not in dispute even where the precise mechanism is hard to pin down. Global Financial Integrity's analysis of a decade of UN Comtrade data found value gaps of US\$835 billion between 134 developing countries and their advanced-economy trading partners in 2018 alone, rising to US\$1.6 trillion once all trading partners are included [1]. The Financial Action Task Force (FATF) and the Egmont Group, in a joint 2020 study, were blunt about why so little of it gets caught: recognising TBML is difficult, particularly where financial institutions, customs authorities and law enforcement each hold only a fragment of the picture and lack a shared understanding of the technique [2]. FATF and the Egmont Group followed this with a detailed set of red-flag indicators in 2021 [4], and the Wolfsberg Group, together with the International Chamber of Commerce and BAFT, issued parallel Trade Finance Principles aimed specifically at the correspondent and trade finance side of banking [3].

At least four distinct literatures bear on how TBML might be detected, and this review finds they engage each other only partially. The first is regulatory and typology-focused, producing the red-flag indicator lists described above. The second consists of a small number of reviews and bibliometric

studies that have already attempted to map the TBML research field itself, offering a meta-level view this paper draws on directly. The third is a body of research on trade-instrument-specific fraud, principally letters of credit, and on customs fraud detection, a problem that overlaps substantially with TBML's under-invoicing and misclassification typologies without always being framed as TBML research. The fourth, by far the largest, is computer science research applying network analysis and machine learning to money laundering detection generally, most of it aimed at financial transaction or cryptocurrency data rather than trade documentation. This paper reviews all four with the aim of establishing what each has found, how far the methods developed for adjacent problems could plausibly transfer to TBML detection specifically, and where the combined evidence base is thin.

The review is organised as follows. The Literature Survey section synthesises the four strands identified above, drawing on the sources described in the Methodology section. The Problem Definition section sets out the specific research questions the review addresses. The Results and Discussion section works through the review's thematic findings and their implications, and the paper closes with a Conclusion and a Future Scope section setting out a research agenda.

2. Literature Survey

On the typology side, Girich and colleagues trace how the "red flags" approach to TBML developed internationally and note that countries have taken different regulatory routes, with Russia leaning on currency-control mechanisms while other jurisdictions favour the FATF-style risk-based indicator system now used by most Western regulators [5]. Marzouk's interview-based study of UK banks after Brexit found that

trade finance teams frequently lack the skills, resourcing and technology to act on TBML intelligence even where the regulatory expectation to do so is clear [6], and his later book-length treatment argues that much of what passes for TBML compliance in practice is "symbolic": audit trails, checklists and metrics that create the appearance of vigilance while the underlying laundering continues largely undisturbed, because static indicator lists are easy for sophisticated launderers to route around once they are known [7].

A small number of studies have already attempted to map the TBML research field as a whole, and their findings usefully frame this review's own. Tiwari and colleagues' 2024 systematic literature review, the first of its kind on TBML specifically, used a formulaic search string and both thematic and meta-analysis to categorise the existing literature into four areas: TBML risk assessment, TBML detection, the role of professionals, and general understanding of the phenomenon, concluding that the limited number of available studies constrains what can be said about the most effective way to combat TBML [8]. Akartuna and colleagues' PRISMA-compliant scoping review went further, coding 105 money-laundering typologies and trends reports and identifying over 200 value instruments, over 200 actors and entities, and 2,565 distinct red-flag indicators across the money-laundering literature broadly, arguing for a more standardised, cross-typological approach to reporting so that data sharing and cooperation across jurisdictions can improve [9]. Gaviyau's bibliometric review of 164 Scopus-indexed TBML documents found only a 1.4 percent annual growth rate in scholarly output and highlighted the growing influence of free trade zones, sanctioned jurisdictions and digital transactions as emerging themes [10]. Herbert's rapid review for the UK context reached a blunter conclusion: despite drawing on international-organisation publications, government reports and expert interviews, the availability of relevant, recent, peer-reviewed academic papers on TBML was limited [11]. Tiwari and colleagues' earlier, broader review of the money-laundering literature found that most detection-focused research concentrates on banking transactions or real estate and trade, while research on shell-company-based laundering remains comparatively scarce [12]. Read together, these five reviews independently corroborate the same observation this paper's own search process arrives at below: TBML-specific peer-reviewed research is a genuinely thin literature relative to the scale of the problem it addresses.

Sector- and jurisdiction-specific studies fill in what the generic FATF indicators look like when applied to a particular corner of trade finance. Makkink and colleagues interviewed compliance and shipping specialists in South Africa and found that freight forwarding companies sit unusually close to shipping-related red flags such as inconsistent routing or mismatched cargo descriptions, yet often lack awareness of the formal TBML guidance that would let them connect those anomalies to a laundering typology [13]. Singh and colleagues took a forensic risk-scoring approach to Reserve Bank of India data on external commercial borrowings, showing that loans from related-party lenders or routed through opaque jurisdictions cluster in a handful of sectors [14]. Chao and colleagues built a knowledge-driven classification model for TBML behaviour

monitoring using Chinese regulatory and financial market data [15]. Paul's conceptual paper on Bangladesh names the structural problem directly: TBML evidence is fragmented across customs declarations, financial intelligence units, banks and company registries, and no single actor sees the whole picture [16]. Singh and colleagues' cross-country panel study of forty economies found that trade openness raises TBML vulnerability, an effect roughly doubled in countries with the weakest regulatory quality [17]. Sivaguru and colleagues' review of the Sri Lankan context, drawing on international case studies given the scarcity of local TBML data, found that geographic position and the prevailing sociopolitical climate leave the country exposed to regional laundering flows [18].

A closely related but distinct sub-literature addresses fraud in specific trade instruments rather than TBML as a laundering category per se. Rafique and colleagues' study of letter-of-credit fraud in Bangladesh, based on thematic analysis of fraud cases and interviews with senior trade finance practitioners, found LC fraud escalating to a level threatening economic stability, and recommended that the Uniform Customs and Practice for Documentary Credits be updated with more specific anti-fraud guidance [19]. This literature is overwhelmingly legal and procedural rather than analytics-driven; Jain and colleagues' recent proposal to apply supervised and unsupervised machine learning to anomaly detection in letters of credit and bank guarantees specifically is a rare and notably recent exception, aimed at automatically sifting trade documents, transaction patterns and counterparty relationships for irregularities [20].

Customs fraud detection is a larger, more mature and more heavily validated adjacent literature that overlaps substantially with TBML's under-invoicing, over-invoicing and misclassification typologies without always being framed in money-laundering terms. Vanhoeyveld and colleagues analysed 9.6 million customs records from the Belgian customs administration, showing that behavioural and high-cardinality features, such as the specific consignee, consignor and declarant involved, are highly predictive of fraud once a purpose-built ensemble method is used to handle the severe class imbalance customs data presents [21]. Singh and colleagues' GraphFC model addresses the related problem of label scarcity directly, using a domain-specific graph neural network validated on real customs data from two countries to generalise well even where manually inspected ground truth is limited [22]. Alwanin and colleagues' dual-learning gradient-boosting approach jointly classifies fraud and estimates revenue risk, reporting that auditing only 10 percent of declarations, prioritised by the model's ranking, recovers close to 88 percent of at-risk revenue [23]. Chen and colleagues addressed a narrower but consequential problem, verifying whether a declared Harmonized System code matches its commodity description using an unsupervised embedding approach, since incorrect codes cause both duty losses and downstream classification errors in any system relying on customs data [24].

The general anti-money-laundering network-analysis and machine-learning literature is by far the largest of the four strands, though most of it targets financial transaction networks rather than trade documentation specifically.

Colladon and Remondi's widely cited study of an Italian factoring company's transaction data showed that social network metrics, including how peripheral or central a client sits in the transaction network, predict risk more effectively than transaction-level attributes alone [25]. Vilella and colleagues developed a centrality-based anomaly detection method for over eighty million cross-border wire transfers [26]. Mohtasam and colleagues built a three-stage unsupervised ensemble on the Elliptic transaction-graph dataset, explicitly to catch laundering patterns that have never been labelled before [27]. Ejime and colleagues, working with the FinCEN Files dataset, combined network centrality measures with machine-learning classifiers and reported ensemble models reaching an area under the ROC curve of around 0.97 to 0.98, with degree and betweenness centrality among the strongest predictors [28]. Earlier graph-based work by Rocha-Salazar and colleagues used neural networks and a bespoke transaction-abnormality indicator to cut false positives relative to a rule-based system at a Mexican financial institution [29], while Dumitrescu and colleagues built features from a reduced egonet structure specifically designed to surface circular payment flows, a pattern that maps closely onto the FATF red flag for funds returning to their country of origin [30]. More recent graph neural network architectures, including Wang and colleagues' reinforcement-learning-enhanced attention network [31] and Schmidt and colleagues' GAMLNet [32], report strong classification performance on synthetic and semi-synthetic money laundering datasets, and Huang and colleagues' earlier CoDetect framework demonstrated the value of combining network structure with node-level features rather than treating them separately [33]. Hilal and colleagues' broad survey of anomaly detection techniques for financial fraud, and Karim and colleagues' evaluation of scalable semi-supervised graph learning across four large transaction datasets, both confirm that this literature has matured well beyond proof-of-concept, with the latter explicitly noting the United Nations' estimate that 2 to 5 percent of global GDP is laundered annually as the scale-of-problem motivation for the field [34, 35].

Explainability recurs as a concern across several of the strands above independently. Mazumder proposes a predictive analytics framework fusing customs, banking and trade finance data using supervised and unsupervised models alongside SHAP and LIME explainability techniques, arguing interpretability is what will let such a system be adopted by regulators who need to justify their decisions [36], and extends this in later work to an explainable-AI framework aimed jointly at sanctions evasion and TBML [37]. Kute and colleagues' critical review of deep learning combined with explainable AI for money laundering detection found that explainability techniques remain largely absent from the AML domain despite growing academic interest in deep learning methods, with over half of reviewed machine-learning approaches non-interpretable and most studies relying on ageing datasets [38]. Konstantinidis and colleagues' more recent deep neural network combined with SHAP achieved very high reported precision and recall on a synthetic dataset while explicitly framing the SHAP layer as necessary for transparency rather than optional [39].

3. Problem Definition

The four strands surveyed above rarely cite one another fully. The typology literature is detailed about what a red flag looks like on a document but says little about how an institution facing a large alert volume should decide which red flags to act on first. The meta-level reviews confirm that TBML-specific peer-reviewed research remains thin, but do not themselves propose how the much larger adjacent literatures, customs fraud detection and general AML network analysis, might close that gap. The customs fraud detection literature is mature and heavily validated on real, large-scale data, yet the papers reviewed here rarely frame their contribution in TBML or money-laundering terms at all, despite addressing typologies, under-invoicing and misclassification, that sit at the centre of FATF's own TBML red-flag categories. And the general AML network-analysis and machine-learning literature reports strong detection performance but is drawn almost entirely from financial transaction or cryptocurrency data rather than trade documentation.

This review sets out to answer three questions from the literature surveyed. First, what do published sources, including the meta-level reviews identified here, establish about the effectiveness and maturity of TBML detection research as it currently stands. Second, how does the customs-fraud-detection literature's evident scale and methodological maturity compare with the much smaller literature explicitly addressing TBML, and what does that gap suggest about where research effort has and has not concentrated. Third, what would need to happen for the general AML network-analysis and machine-learning literature's strong results, and the customs fraud literature's validated methods, to transfer usefully to trade finance data specifically, while preserving the interpretability that red-flag-based regulatory frameworks require.

4. Methodology

This review is a structured narrative synthesis rather than a formal systematic review with a pre-registered protocol; it does not claim exhaustive coverage of the literature, and the search process below is reported in full so the review's scope can be assessed on its own terms. Readers seeking a formal PRISMA-compliant systematic review of the TBML literature specifically are directed to Tiwari and colleagues [8] and Akartuna and colleagues [9], both reviewed as primary sources above.

Academic sources were identified through seven structured searches of Consensus, an academic search platform covering Semantic Scholar, PubMed, Scopus and ArXiv, run with exclude_preprints set to true so that only peer-reviewed or otherwise formally published work was returned. The search queries, run across two research sessions, were: "trade-based money laundering detection red flags financial institutions"; "network analysis graph-based anomaly detection trade-based money laundering customs declarations"; "explainable AI trade based money laundering detection"; "trade based money laundering typology systematic literature review"; "letters of credit bank guarantee fraud detection trade finance anomaly"; "customs fraud detection machine learning international trade declarations misclassification"; and "graph

neural network money laundering benchmark dataset cross-border transaction network". Each search returned up to ten results; across the seven searches this yielded approximately sixty results before removing duplicates. From these, thirty-five papers were retained on the basis of direct topical relevance to TBML detection or to one of the three closely adjacent problems identified in the Introduction, trade-instrument fraud, customs fraud detection, and general AML network analysis; papers addressing money-laundering typologies without a detection, review or methodological contribution, and papers on cryptocurrency-specific graph neural network benchmarks with no stated relevance to trade finance, were generally excluded once the thematic saturation of that sub-area became apparent.

Four regulatory and institutional sources, FATF and the Egmont Group's Trends and Developments and Risk Indicators reports, Global Financial Integrity's trade misinvoicing analysis, and the Wolfsberg Group's Trade Finance Principles, were retrieved directly from the issuing body's website to ensure the review's typology strand was not limited to how academic papers describe regulatory frameworks second-hand [1, 2, 3, 4]. No formal quality-appraisal tool was used to score the included studies; sources are instead discussed with their methodology and scope stated directly in the Literature Survey and Results and Discussion sections, so the reader can weigh the strength of each claim individually.

Two limitations follow from this process. First, retention decisions after the seven searches were made by the reviewer's judgement of topical relevance rather than a pre-specified, replicable inclusion protocol, which is a materially weaker standard than the PRISMA-ScR approach Akartuna and colleagues used [9] or the formulaic search-string approach Tiwari and colleagues used [8]. Second, the search was run entirely against one platform; relevant work indexed primarily in trade-finance, customs-administration or supply-chain-management outlets that a money-laundering-centred query does not surface well may be under-represented, a risk the customs-fraud-detection search queries were specifically added to mitigate but cannot be assumed to have fully resolved.

5. Results and Discussion

Four findings emerge from the literature surveyed. First, and most directly answering the review's first research question, TBML-specific peer-reviewed research is genuinely thin, and this is not merely this review's own search limitation speaking. Tiwari and colleagues' formal systematic review explicitly states that the limited number of available studies constrains what conclusions can be drawn about the most effective way to combat TBML [8]; Herbert's rapid review, despite drawing on international organisations, government reports and expert interviews to compensate, found the availability of relevant, recent, peer-reviewed academic papers limited [11]; and Gaviyau's bibliometric analysis found only a 1.4 percent annual growth rate in scholarly TBML output despite the field's evident real-world importance [10]. Three independent reviews using three different methods, systematic review, rapid review, and

bibliometric analysis, reach the same conclusion, which is a stronger signal than any one of them stating it alone.

Second, and this is the finding this review's own search strategy was best positioned to surface, a large, mature and heavily validated customs-fraud-detection literature exists on problems that sit at the centre of FATF's own TBML typology, over- and under-invoicing, misclassification, and undervaluation, yet this literature and the TBML-specific literature identified above show remarkably little cross-citation. Vanhoeyveld and colleagues' 9.6-million-record study with the Belgian customs administration [21] and Singh and colleagues' cross-country-validated GraphFC model [22] represent exactly the kind of real, large-scale, validated detection research the TBML-specific literature lacks, yet neither paper frames its contribution in money-laundering terms, and none of the TBML-specific papers reviewed above cite this customs-fraud-detection literature directly. This is a concrete, addressable gap rather than a general observation about data scarcity: the detection methods already exist, validated on real government data at a scale no TBML-specific study reviewed here approaches, and the barrier appears to be disciplinary framing rather than technical feasibility.

Third, the trade-instrument-specific fraud literature, concentrated on letters of credit, remains almost entirely legal and procedural. Rafique and colleagues' Bangladesh study and the broader body of doctrinal LC-fraud scholarship it sits within address how courts and the Uniform Customs and Practice for Documentary Credits should treat fraud once it is discovered, not how to detect it in the first place [19]. Jain and colleagues' proposal to apply machine learning to LC and bank guarantee anomaly detection is recent and, on the evidence of this search, close to alone in its space [20], suggesting trade-instrument-level analytics is an emerging rather than established sub-field, several years behind where customs-fraud analytics already stands.

Fourth, the general AML network-analysis and machine-learning literature remains the deepest, most methodologically advanced, and most consistently high-performing of the four strands, but its headline results are drawn almost entirely from financial transaction or cryptocurrency data rather than trade finance documentation. Ejime and colleagues' 0.97-to-0.98 AUC result [28] and Karim and colleagues' evaluation of scalable graph learning across four large transaction benchmarks [35] both demonstrate what is achievable on clean, structured transaction data; Dumitrescu and colleagues' circular-flow features [30] and Colladon and Remondi's centrality metrics [25] are conceptually transferable to a trade finance counterparty graph, since both operationalise patterns FATF's own red flags already describe, but no paper identified in this review tests that transfer directly on trade documentation. Explainability recurs as a requirement across this strand and the explainability-focused sources independently, from Mazumder's SHAP- and LIME-based framework [36, 37] to Kute and colleagues' finding that over half of reviewed AML machine-learning methods remain non-interpretable [38] to Konstantinidis and colleagues' explicit framing of explainability as a necessity rather than an enhancement [39],

which is a stronger convergent signal than any single source raising the point alone.

It is worth restating the limitations of this review's own process plainly. Thirty-five academic papers curated from sixty results across seven search queries against one platform, using relevance judgement rather than a pre-specified protocol, is a narrower and less rigorous base than Tiwari and colleagues' or Akartuna and colleagues' formal systematic reviews [8, 9], both cited throughout this section precisely because they offer a more rigorous version of the same underlying claim. The finding that TBML-specific research is thin should be read as independently corroborated by those two more rigorous reviews rather than resting on this review's own search alone; the finding that a disconnected customs-fraud-detection literature exists is this review's own contribution and has not, to this review's knowledge, been previously stated in these terms, so it should be read with the corresponding caution that follows from a single, narrower search process.

6. Conclusion

This review set out to establish what the published literature says about detecting trade-based money laundering, drawing on thirty-five academic papers identified through seven structured searches of an academic database and four regulatory publications identified through targeted web search. TBML-specific peer-reviewed research is genuinely thin, a finding this review's own search corroborates and that three prior systematic and bibliometric reviews of the field have already established independently. A large, mature and well-validated customs-fraud-detection literature exists on directly adjacent problems, over-invoicing, under-invoicing, misclassification, but shows little cross-citation with TBML-specific research despite addressing the same underlying typologies. Trade-instrument-specific analytics, particularly for letters of credit, remains an emerging rather than established sub-field. And the general AML network-analysis and machine-learning literature reports strong detection performance but almost entirely on financial transaction or cryptocurrency data rather than trade documentation, leaving the transferability of its methods to TBML largely untested in the sources reviewed here. Explainability emerges as a convergent, independently stated requirement across every strand that engages with analytics-driven detection at all.

The central conclusion is that the field's next contribution is more likely to come from connecting existing, validated detection methods across disciplinary boundaries than from developing new ones from scratch. The customs-fraud-detection literature has already solved, at real government scale, problems that sit at the heart of TBML's own typology; the general AML network-analysis literature has already validated centrality- and graph-based features that map directly onto FATF's red-flag categories. What is missing is not the underlying technique but the bridge between fields that currently do not cite each other, and the explainability requirement that recurs independently across every strand suggests that bridge needs auditability built in as a design constraint from the outset, not retrofitted afterward.

This conclusion rests on a narrower and less formally protocolled search than Tiwari and colleagues' or Akartuna and colleagues' systematic reviews, for the reasons set out in the Methodology section, and should be read with that limitation in mind, particularly for the cross-literature gap finding, which is this review's own contribution rather than one independently corroborated elsewhere.

7. Future Scope

Three directions follow from the gaps this review identifies. The first, and most directly actionable given the maturity of the underlying methods, is a validation study applying customs-fraud-detection techniques already proven at scale, Vanhoeyveld and colleagues' behavioural feature engineering [21], Singh and colleagues' label-scarce graph neural network approach [22], or Alwanin and colleagues' joint classification-and-revenue-risk model [23], to trade finance data specifically framed as a TBML detection problem, to establish whether methods validated on customs declarations transfer to the letters of credit and trade documentation TBML red flags are built around. The second is a more rigorous, formally protocolled systematic review than this paper's own narrative synthesis, ideally extending Tiwari and colleagues' [8] and Akartuna and colleagues' [9] existing systematic and scoping reviews with an explicit search strategy spanning customs-administration and trade-finance outlets alongside money-laundering-specific ones, to test whether the cross-literature disconnect this review identifies holds under fuller, more rigorous search coverage. The third is a closer empirical look at the explainability requirement that recurs across every strand of this review: what specific evidentiary standard compliance and law-enforcement audiences require of an analytics-driven TBML alert before they will act on it, since Mazumder [36, 37], Kute and colleagues [38], and Konstantinidis and colleagues [39] all raise the requirement independently but none of the sources reviewed here measures it directly against practitioner judgement.

Beyond these three, the wider question this review leaves open is whether the disciplinary separation between TBML research, customs fraud detection, and general AML network analysis reflects a genuine difference in the underlying problems, or simply the accident of which academic and professional communities happen to study each one. Testing that question would require researchers willing to work explicitly across those boundaries, which very few of the sources identified in this review currently do.

References

- [1] Global Financial Integrity. Trade-Related Illicit Financial Flows in 134 Developing Countries: 2009-2018. Washington DC: GFI; 2021.
- [2] Financial Action Task Force, Egmont Group. Trade-Based Money Laundering: Trends and Developments. Paris: FATF; 2020.
- [3] Wolfsberg Group, International Chamber of Commerce, BAFT. Trade Finance Principles. 2019.
- [4] Financial Action Task Force, Egmont Group. Trade-Based Money Laundering: Risk Indicators. Paris: FATF; 2021.

- [5] Girich M, et al. Trade-based money laundering: development of the «red flags» system. *Int Trade Trade Policy*. 2022.
- [6] Marzouk M. Trade-based money laundering (TBML) empowers criminals to run free post-Brexit. *J Money Laund Control*. 2021.
- [7] Marzouk M, et al. Trade-Based Money Laundering. 2026.
- [8] Tiwari M, et al. Trade-based money laundering: a systematic literature review. *J Account Lit*. 2024.
- [9] Akartuna EA, et al. Motivating a standardised approach to financial intelligence: a typological scoping review of money laundering methods and trends. *J Exp Criminol*. 2024.
- [10] Gaviyau W. Exploring trade-based money laundering research: a bibliometric review. *Int J Appl Res Bus Manag*. 2026.
- [11] Herbert G. Global developments in trade-based money laundering. *K4D Helpdesk Report*. 2022.
- [12] Tiwari M, et al. A review of money laundering literature: the state of research in key areas. *Pac Account Rev*. 2020.
- [13] Makkink IM, et al. The role of freight forwarding companies in detecting and investigating trade-based money laundering. *J Money Laund Control*. 2024.
- [14] Singh R, et al. External commercial borrowings and trade-based money laundering: a forensic risk analysis with evidence from India. *J Money Laund Control*. 2026.
- [15] Chao XR, et al. Behavior monitoring methods for trade-based money laundering integrating macro and micro prudential regulation: a case from China. *Technol Econ Dev Econ*. 2019.
- [16] Paul S. Macro-meso-micro analytics for trade-based money laundering risk screening in Bangladesh: a conceptual entity-graph extension. *J Econ Criminol*. 2026.
- [17] Singh R, et al. When trade openness meets weak institutions: evidence on trade-based money laundering vulnerability. *EPH Int J Bus Manag Sci*. 2026.
- [18] Sivaguru D, et al. The phenomenon of trade-based money laundering (TBML) – a critical review in Sri Lankan context. *J Money Laund Control*. 2023.
- [19] Rafique RB, et al. The risk of fraudulent letter of credit transactions in Bangladesh: a growing threat to Bangladesh's economy. *J Financ Crime*. 2023.
- [20] Jain PKT, et al. Leveraging artificial intelligence for trade-based money laundering detection: a machine learning approach for anomaly detection in letters of credit and bank guarantees. *Int J Innov Sci Res Technol*. 2025.
- [21] Vanhoeyveld J, et al. Customs fraud detection. *Pattern Anal Appl*. 2019.
- [22] Singh K, et al. GraphFC: customs fraud detection with label scarcity. In: *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*; 2023.
- [23] Alwanin R, et al. Customs fraud detection using a gradient boosting approach for joint classification and risk estimation. *Sci Rep*. 2025.
- [24] Chen H, et al. The use of machine learning to identify the correctness of HS Code for the customs import declarations. In: *2021 IEEE 8th International Conference on Data Science and Advanced Analytics (DSAA)*; 2021.
- [25] Colladon AF, Remondi E. Using social network analysis to prevent money laundering. *Expert Syst Appl*. 2021.
- [26] Vilella S, et al. Weirnodes: centrality- based anomaly detection on temporal networks for the anti-financial crime domain. *Appl Netw Sci*. 2025.
- [27] Mohtasam M, et al. Machine learning-based unsupervised ensemble approach for detecting new money laundering typologies in transaction graphs. *Int J Appl Math*. 2025.
- [28] Ejime OM, et al. Machine learning and network analysis for cross-border suspicious transaction detection: evidence from the FinCEN Files. *IEEE Access*. 2026.
- [29] Rocha-Salazar JJ, et al. Money laundering and terrorism financing detection using neural networks and an abnormality indicator. *Expert Syst Appl*. 2021.
- [30] Dumitrescu B, et al. Anomaly detection in graphs of bank transactions for anti- money laundering applications. *IEEE Access*. 2022.
- [31] Wang Q, et al. RMGANets: reinforcement learning-enhanced multi-relational attention graph-aware network for anti-money laundering detection. *Complex Intell Syst*. 2024.
- [32] Schmidt J, et al. GAMLNet: a graph based framework for the detection of money laundering. In: *2024 11th IEEE Swiss Conference on Data Science (SDS)*; 2024.
- [33] Huang DX, et al. CoDetect: financial fraud detection with anomaly feature detection. *IEEE Access*. 2018.
- [34] Hilal W, et al. Financial fraud: a review of anomaly detection techniques and recent advances. *Expert Syst Appl*. 2021.
- [35] Karim MR, et al. Scalable semi-supervised graph learning techniques for anti money laundering. *IEEE Access*. 2024.
- [36] Mazumder PT. Data driven detection of trade based money laundering (TBML): a predictive analytics framework for securing US supply chains and financial integrity. *SAMRIDDHI J Phys Sci Eng Technol*. 2023.
- [37] Mazumder PT. Explainable artificial intelligence for protecting the U.S. financial system from sanctions evasion and trade-based money laundering. *Int J Adv Ind Eng*. 2024.
- [38] Kute DV, et al. Deep learning and explainable artificial intelligence techniques applied for detecting money laundering – a critical review. *IEEE Access*. 2021.
- [39] Konstantinidis G, et al. Deep neural networks for anti-money laundering using explainable artificial intelligence. In: *2024 IEEE 12th International Conference on Intelligent Systems (IS)*; 2024.

Author Profile

Malki Senevirathne is a risk and compliance professional based in the United Kingdom with 12 years of experience in the financial industry. She is an ICA-qualified professional and holds an MSc from the University of Westminster, an MBA from Cardiff Metropolitan University, and a BBMgt (Bachelor of Business Management) from the University of Kelaniya. Her professional interests include anti-money laundering transaction monitoring, trade-based money laundering detection and financial crime risk management.