Tackling Trade Credit Fraud: Challenges and Solutions

Modern supply chain critically relies on trade credit to function properly. With a trade credit, a supplier is not paid immediately upon shipping the goods to the buyer but at some later point, often when the buyer has sold the goods and realized the revenue. This type of credit plays a central role in managing cash flow within the supply chain and is a key factor in the relationship between suppliers and their customers.

However, as the digital economy grows, so does the prevalence of trade credit fraud—a serious issue that creates significant financial strain for businesses. According to a recent study, 98% of B2B businesses reported fraud attacks in 2021, losing on average 3.5% of their annual sales revenues while 47% of B2B businesses had chosen not to onboard new clients due to fraud concerns (Payments). Financial institutions use sophisticated AI-base tools to detect fraud. However these tools cannot be directly applied to supply chain trade credit due to several factors inherent to how trade credit operates.

Why Trade Credit Is Different

Supply chain trade credit has unique characteristics that lead to unique challenges in fraud detection. Many of these unique aspects stem from the fact that trade credit decisions are made, and credit is issued, by suppliers to buyers directly, while other types of credit are commonly extended by banks and other financial institutions. The following are some of the noteworthy differences between trade credit and financial loans.

Business Imperative

A financial institution would like to make a loan if it is likely to be profitable but it does not have to. The available funds can always be invested in other ways, and there is no urgency as money does not have a shelf life. In contrast, suppliers must move their products. Finding an alternative buyer may be difficult and not quick enough – the product may spoil. In many cases, a bulk of supplier’s business occurs on trade credit.

Time Sensitivity

In industries like construction, manufacturing, and logistics, materials and resources are often needed urgently. Unlike banks, which have the flexibility to conduct detailed credit assessments, review trade references, and verify financial documents over days or even weeks, suppliers often need to approve credit within hours or lose sales. This accelerated time frame limits the thoroughness of risk assessments and can increase susceptibility to fraud or defaults (Incomlend, BrightQuery).

Incomplete Documentation

Unlike financial credit applications, which are typically accompanied by extensive documentation such as identification, financial statements, and utility bills, trade credit applications often come with incomplete or inconsistent information. Even when suppliers request essential documents like balance sheets, profit and loss statements, or trade references, these may be incomplete, forged, altered, or outdated. Despite this, suppliers must make rapid credit decisions due to the need of maintaining customer relationships and meeting timeliness requirements.

Unsecured Credit

Banks typically require collateral—such as property, equipment, or inventory—for business loans, making these loans "secured" because the asset can be seized and liquidated if the borrower defaults. In contrast, in trade credit, suppliers rarely take collateral from buyers, creating a greater opportunity for fraudsters, who don't have to front resources before a credit is issued, and a high risk of loss for the suppliers, especially in sectors with repeated, high-value transactions (Incomlend, BrightQuery, Deloitte United States). To mitigate the risk, some suppliers explore alternative strategies, such as credit insurance or factoring. However, these tools come with additional costs, and they do not fully replicate the security provided by traditional collateral, leaving the inherent risks of unsecured trade credit largely unaddressed in the current supply chain credit framework (Deloitte United States).

Ongoing Transactions

Unlike a traditional loan, which is typically a one-time transaction with a fixed repayment schedule, trade credit often involves repeated orders where goods are shipped before payment on prior credit has been made. This structure can allow buyers to accumulate substantial debt, effectively turning trade credit into a revolving line of credit. The risk is heightened in industries with high material turnover, like construction or manufacturing, where quick delivery of supplies is essential to keep projects on track (Incomlend, BrightQuery, Deloitte United States) and, thus, fraudsters may exploit the revolving nature of credit to rack up debt over multiple transactions, often without detection until it's too late.

Limited Resources and Know-How

Financial fraud detection solutions, especially in consumer fraud detection, typically rely on advanced algorithms, machine learning, and access to vast amounts of personal and transactional data to train machine learning models and flag potentially fraudulent activities in real-time. Financial institutions have teams of data scientists dedicated to fraud detection efforts. However, suppliers don't have specialized knowledge, resources, and access to necessary data, to replicate these activities.

Fraud in the Supply Chain

Fraud in the context of trade credit can take many forms: a more common first- and third-party fraud as well as more recently emerging synthetic and AI-assisted fraud. These types of fraud often involve deceptive tactics during the credit application process, where a credit applicant misrepresents their identity to gain access to credit they are not entitled to.

First-Party Fraud: Falsifying Company Existence

First-party fraud occurs when an applicant either pretends to be a non-existent company or by creating a fake business entity for the whole purpose of obtaining credit. This type of fraud often involves submitting fraudulent documents, such as fabricated financial statements, tax returns, and registration records, to convince the supplier that the company is legitimate. Once the goods are delivered, the fraudulent business ceases all communications, and the supplier is left with unpaid debt (HighRadius, Export Finance Blog).

Third-Party Fraud: Impersonating Another Entity

Third-party fraud, on the other hand, involves an applicant pretending to be someone else to access credit under false pretenses. In this scenario, the fraudster may impersonate an established, creditworthy company to obtain trade credit. This often involves the use of stolen identities or misrepresentation of ownership. Once the fraudster receives the goods, the legitimate company remains unaware of the fraudulent activity until it is too late (Export Finance Blog, HighRadius) .

Synthetic Fraud: A Distinct Threat

Synthetic fraud relies on crafting a fabricated entity from scratch, blending fake and potentially stolen information to create a new, seemingly real business. For example, a fraudster might register a fake company using a valid but stolen EIN from a dissolved or inactive business. They could create a website, set up a professional email domain, and even generate fake trade references and financial statements.

Additionally, AI-assisted fraud could involve using AI tools to generate hundreds of fake business profiles in bulk, complete with AI-generated websites, contact details, and even simulated business transaction histories, scaling up synthetic fraud to a new level of sophistication.

A Practical Approach to Fraud Detection in Supply Chain Trade Credit

I spent a number of years with a trade credit application processing company, which would receive credit applications for subscribing suppliers and facilitate their application processing workflow. Our customers – the suppliers – dealt with credit application fraud in various ways: from just absorbing fraud costs (increasing the cost of business) to sending personnel to physically check the applicants' identity and company affiliation (increasing time-to-decision and, again, cost of business due to labor involved).

As mentioned before, a typical approach to address fraud detection would employ AI-based technology to flag applications that are more likely to be fraudulent. Unfortunately, this approach wasn't feasible due to lack of labeled applications (i.e., a sample of credit application examples that have been identified as legitimate or fraudulent) to train the ML models. In other words, at least this company (although I suspect it to be true of other vendors in this space) does not receive (and hence obviously does not collect) the information from its customers if they end up finding some credit applications to be fraudulent.

Potentially, one could also attempt to infer the "fraud"/"legitimate" labels based on the data from the downstream systems in the Order-to-Cash (OTC) flow. For example, a collection system could provide collector's notes indicating that a buyer was unreachable or a payment system could show that the balance was never paid and eventually written off. If such data is available, these implied labels could prove useful in bootstrapping the training process.

In the meantime I was curious to see if other, more ad-hoc and immediately applicable, approaches might be beneficial to help with fraud prevention. Read on to see what I found. Spoiler alert – these approaches help a great deal as they can substantially (by almost 60% among the thousands of applications I analyzed) narrow the set of potentially fraudulent applications that the supplier might need to take a closer look.

My goal is to stratify credit applications with respect to their likelihood of being fraudulent. To this end I developed a number of metrics – obtained from various sources that are part of the Internet machinery rather than from application data – to assess the validity of application.

Verifying Email and Company Association

A typical trade credit application includes the name of the company that applies for credit, the shipping address for receiving the goods, and the name and email of the person submitting the application.

I start with the assumption that the applicant's email address is (or at least can be) verified to be valid and indeed belongs to the person who submitted the application. Email verification is ubiquitous in digital communication to screen out bots and mistyped addresses. In the context of trade credit, a supplier can send an email to the applicant's claimed email address that the applicant must respond to before the application proceeds.

However, verifying that the email exists only confirms the applicant's ability to communicate but not their legitimacy. To assess the applicant's legitimacy, I attempt to "bind" the verified email to the company listed in the application by searching independent information sources on the Internet for the appearance of this email. A successful binding means a third-party has associated the company with the applicant's email and hence that the company name is not misrepresented in the application. By excluding successfully bound applications, this process can narrow the scope of potentially fraudulent applications and reduce the manual review workload for suppliers.

I attempt to bind the applicant to their claimed company by looking at:

- **Direct Associations:** Linkage between the applicant's email and the company's website or other reliable internet resources.
- **Third-Party References:** Linkage between the applicant's email and the company on reputable platforms like government databases or Better Business Bureau listings.
- **Transitive Bindings:** Indirect linkage between the applicant's email and the company, where the email is bound to the applicant's name in the application and the latter is bound to the company.
- **Proximity Associations:** A binding has been established between applicant's email and a name listed on the application *and* someone with that name resides in a proximity to the (separately verified) company address.

These different types of associations carry a different weight depending on the reliability of the resource. For example, mentions by city service departments, local government databases, or chamber of commerce websites indicate higher credibility, as these institutions generally vet businesses before endorsement. Similarly, recommendations from verified community platforms like Nextdoor, where real neighborhood users discuss and endorse local businesses, can be more reliable. These endorsements are typically harder to falsify than a self-created social media profile because they come from multiple, unrelated users who have interacted with the business directly.

Domain Analysis for Business Authenticity

In addition to linking the applicant's email address to a company on the credit application, email domain analysis also helps assess the authenticity of the business. Corporate domains (e.g., @companyname.com) are generally more credible than a free domain (e.g., @gmail.com). In the case of a corporate email domain, I further check the age (and potentially reputation) of the applicant's email domain. For example, fraudsters often use newly created domains, whereas older, established domains indicate more credibility (Incomlend, Deloitte United States). For newly created domains, the domains that have been purchased for a longer time period are more trustworthy. Thus I introduce a ***Business Durability Score*** – a metric that takes into account both the age of the domain as well as the time to expiration but could also add age and expiration of other Internet resources (e.g., SSL certificates) and reputation of Internet resource providers (e.g., certificate authorities, domain registrars, or Web hosters) the company uses. By analyzing this score, suppliers can better judge if they are dealing with a legitimate business.

Reference Quality

A trade credit application also typically includes trade references, each including both the name of the company providing the reference and contact information (name and email) of a person representing that company. References offer important fraud detection opportunities, even when the supplier hasn't received the actual reference responses. The contact information for references, such as email addresses and phone numbers, can reveal red flags early in the process:

- **Duplicate or Self References**: Credit applications usually require multiple trade references. Repeated reference names or contact details, or listing the applicant themselves as a reference raise concerns about the legitimacy of the applicant.
- **Reference Email validation**: Malformed emails or emails at non-existent domains could indicate fraudulent intent.
- **Email-to-Company Association for References**: Just as applicants are scored, references themselves can be evaluated based on their online digital footprint, helping to identify whether the contact person actually represents the reference company listed.

These various attributes of reference contact information contribute to the ***Reference Quality Score.*** Even if a supplier doesn't receive the actual response to a reference request in time to make a credit decision, the reference contact information itself, through the Reference Quality Score, can still indicate potential fraud.

Results

This approach turned out to be surprisingly effective at narrowing the set of applications for further review. I further applied the described strategies to associate applicant email addresses with companies to 19,644 credit applications across 66 suppliers submitted in 2023. Even when excluding low-confidence proximity matches, high-confidence bindings would lead to a 58% reduction in applications requiring manual reviews for third-party fraud prevention (see Fig. 1).

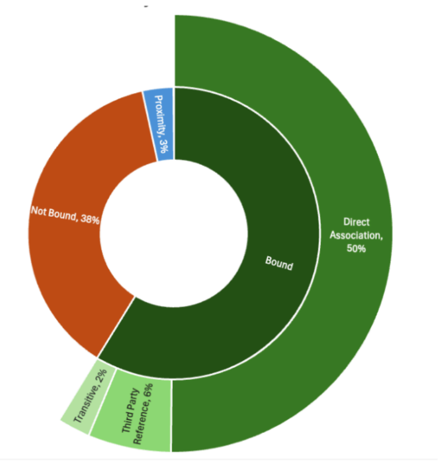

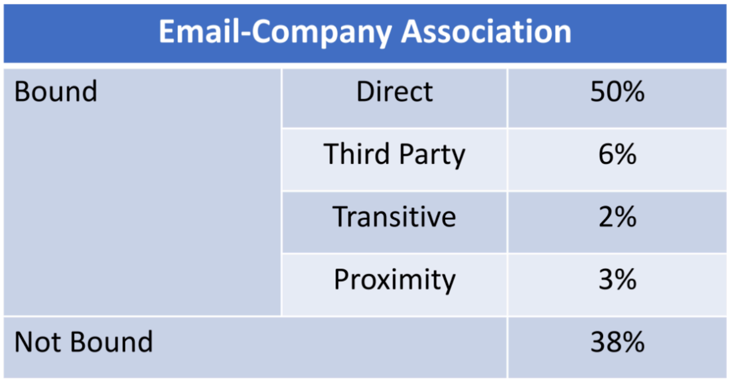

Email-Company Association		
Bound	Direct	50%
	Third Party	6%
	Transitive	2%
	Proximity	3%
Not Bound		38%

Fig 1: Bound vs Not Bound Applications

To confirm that bound applications do reflect lower likelihood of fraud, I used the same methodology on 13,748 applications submitted between 2020 and 2022, and then checked whether the applicant's domain remained active in 2023. While the distribution of "bound" versus “not bound” applications stayed consistent with the 2023 data, the prevalence of inactive applicant domains from not bound applications was over five times greater than those from the bound category.

I also analyzed (see Fig.2), the mean values for the References Quality Score and Business Durability Score, finding significant differences ($p < 0.0001$) between bound and unbound credit applications. Similarly, the bound group received significantly ($p < 0.0001$) more reference responses compared to the unbound group. These correlations reinforce the intuition that these individual metrics all relate to the likelihood of fraud.

Interestingly, both bound and not bound applications were approved at nearly identical rates. Thus, the suppliers did not appear to effectively factor fraud detection in their credit decisions.

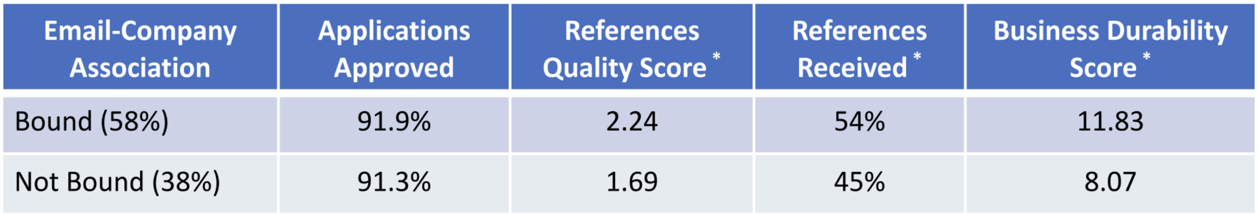

Email-Company Association	Applications Approved	References Quality Score *	References Received *	Business Durability Score *
Bound (58%)	91.9%	2.24	54%	11.83
Not Bound (38%)	91.3%	1.69	45%	8.07

Fig 2: Characteristic of Bound vs Not Bound Applications

Conclusion

The above analysis shows that even without machine learning, the described approach can substantially narrow the scope of potentially fraudulent applications. It is very important to realize that excluding applications from further review does not have to be precise: as long as applications excluded from further scrutiny are less likely to be fraudulent than the remaining applications, this solution already reduces the business costs relative to the current state of affairs. Over time, as suppliers adopt these techniques and the number of manual reviews is reduced, credit managers can spend more time on each application, detecting more fraudulent submissions and accumulating valuable labeled data. These labels, in turn, will lay the foundation for more advanced AI-driven solutions, which can scale fraud detection efforts and develop more robust classifiers to combat trade credit fraud.

Future enhancements could include analyzing discrepancies between shipping and business addresses, though one must recognize that such mismatches often have legitimate explanations, such as construction companies shipping directly to work sites. Similarly, comparing the location of IP address from where the application was submitted to the declared business address can offer insights but must account for remote work and travel. Additionally, behavioral patterns commonly used in consumer fraud detection could be adapted to trade credit contexts—after all, while the focus is on business credit applications, it is ultimately a person behind the submission.