Whitepaper

Real Identities, Real Losses

Reducing First-Party Fraud

SentiLink

Real Identities, Real Losses

Reducing First-Party Fraud

03 Introduction

03 The Cost

05 Defining First-Party Fraud

09 Detecting First-Party Fraud

11 Flagging First-Party Fraud Risk More Precisely

15 Operationalizing the First-Party Fraud Score

17 Conclusion

18 About SentiLink

Introduction

In recent retrostudies with more than a dozen companies across a variety of industries, SentiLink found that, on average:

- **2.89% of funded loans/booked accounts scored as high-risk for first-party fraud.**
- Institutions likely have **five times more first-party fraud on their books** than identity theft and synthetic fraud combined.
- High-risk first-party fraud applications accounted for **72% of preventable identity-fraud losses.**

First-party fraud is difficult to define, and even more difficult to identify at account opening. But significantly reducing first-party fraud losses is possible. In this whitepaper, we'll take a closer look at:

- How much first-party fraud is likely hiding on your books right now.
- Why first-party fraud is difficult to define and identify at account opening.
- Why traditional fraud solutions don't catch enough first-party fraud.
- How to more effectively flag first-party fraud risk from the start.

The Cost

In SentiLink's 1H 2026 Fraud Report, we found that the first-party fraud rate in account-opening applications was 2%. The identity theft rate, by comparison, was 6.12%. But unfortunately, while the incidence of first-party fraud may be lower than identity theft, the success rate for first-party fraudsters is much higher because first-party fraud is much harder to identify.

Based on recent SentiLink retrostudies with more than a dozen companies, across industries ranging from banking to auto lending to telecom, we found that **on average, 2.89% of funded loans/booked accounts scored as high risk for first-party fraud:**

2.89%

average percent of **funded loans/booked accounts** that score as high-risk for first-party fraud

5x

on average, institutions likely have **5x more first-party fraud on their books** than identity theft or synthetic fraud

72%

on average, **first-party fraud accounted for 72% of preventable fraud losses** across all three types of identity fraud

When we tested funded accounts for all three types of identity fraud (identity theft, synthetic fraud, and first-party fraud) we found that, on average, institutions likely had **5x more high-risk first-party fraud accounts on their books** than identity theft and synthetic fraud accounts combined.

The results look similar in real-dollar terms: on average, **72% of these companies' preventable identity fraud losses** came from applications that scored as **high-risk for first-party fraud**.

In short: while first-party fraud may be less prevalent on a per-application basis than identity theft, it is generally costing companies more because a much greater proportion of those applications make it onto their books.

What makes first-party fraud so difficult to detect?

Defining First-Party Fraud

In common parlance, the term first-party fraud refers to fraud committed using a person's real identity (name, DOB, SSN) to obtain credit they do not intend to pay back and/or funds they are not entitled to keep.

But intent is impossible to measure, particularly at the point of application. And what specific MOs does first-party fraud include?

When we begin to look at first-party fraud in the context of fraud prevention, that generic definition isn't precise enough to be useful. Getting more precise presents a challenge.

The problem

The first step in detecting any type of fraud is defining it. Without a clear definition, case labels will be inconsistent. Any models built on those labels will be inconsistent, too.

The line between first-party fraud risk and credit risk is blurry. In the generic definition above, first-party fraud is discussed as a question of intent: does the borrower intend to pay back the loan? But intent is impossible to measure. As a result, the same account that's marked as a first-party fraud loss at one institution may be considered a credit loss at another.

And with no victim to report a crime, even if those two institutions agree on the definition of first-party fraud, they may never agree on how to categorize a given account. Without a clear, agreed-upon definition and empirical signals to evaluate against, the line between first-party fraud loss and credit loss often comes down to a judgement call.

Further complicating the definition problem, first-party fraud includes a wide variety of fraud MOs, and these impact institutions differently. The following is a partial list of common first-party fraud tactics, but this is really just the tip of the iceberg:

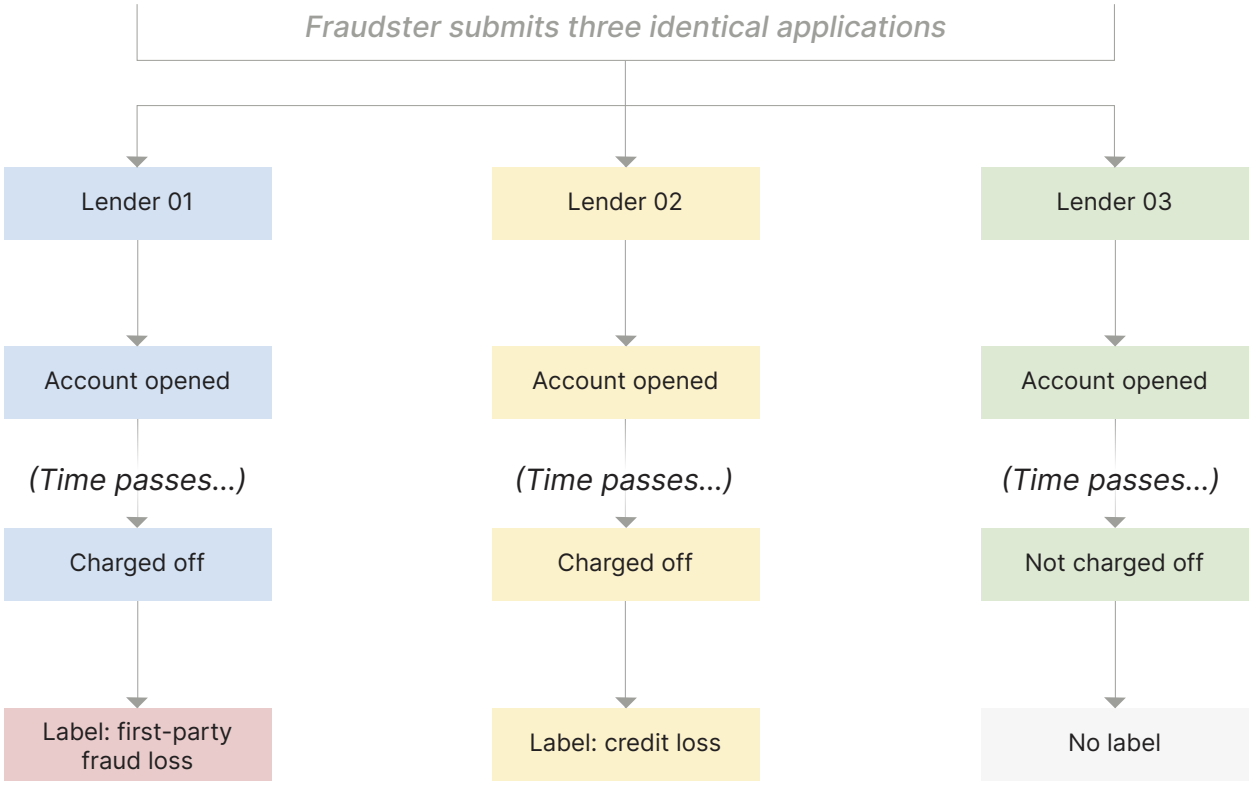
- **Credit washing** — The fraudster maxes out credit lines and then falsely claims to be victim of identity theft, using this "victim" status to have the bad tradelines removed from their credit report.

- **Piggybacking** — The fraudster illegitimately boosts their credit score by purchasing aged tradelines.
- **Synthetic business** — The fraudster has opened an entirely fabricated business or is misrepresenting the legitimacy of their business to obtain business loans.
- **Check fraud** — The fraudster deposits a fake or fraudulently modified check and moves the money to a different account or withdraws it as cash before the bank discovers the check is illegitimate. (Banks are required to comply with funds availability rules and thus sometimes must make deposited funds available before they're able to fully assess a check's legitimacy).
- **ACH fraud** — The fraudster creates multiple checking accounts, initiates an ACH transfer from one, and reports it as unauthorized from another account.
- **Bust-out fraud** — The fraudster maxes out lines of credit with no intent to pay them back, often after having first slowly built up and paid off credit to increase the amount they're able to borrow for the bust-out.
- **Employment or income fraud** — The fraudster falsifies documentation or otherwise misrepresents their employment situation or income to increase their chances of approval for credit products.
- **Auto loan stacking** — The fraudster applies for auto loans at several lenders at once, exploiting the reporting lag so that no individual lender can see the other loans in progress at the point of application.

While SentiLink categorizes all of these (and more) as first-party fraud MOs, some organizations may define first-party fraud differently — often more narrowly — based on the specifics of their business and the MOs they see most often.

This poses an obvious challenge when it comes to building a solution to capture first-party fraud. Sourcing labels from a formal consortium won't produce consistent results because different institutions within the consortium likely define first-party fraud differently. And even when first-party fraud definitions align, labeling accuracy across institutions varies quite a bit. Without a clear taxonomy and empirical signals to evaluate against, the end result will produce far too many inconsistencies.

The scenario below illustrates how the exact same incoming application information at three different lenders (for example) can lead to three different outcomes:



This is the problem with an inconsistent definition and by extension inconsistent labeling standards: it produces inconsistent labels, and no model trained on inconsistent labels is likely to perform well.

An effective model for catching first-party fraud before it's on your books, therefore, has to start with a single definition of first-party fraud, based on measurable data points, from which consistent labels can then be derived.

Our definition

While it doesn't perfectly capture every type of fraud MO, to build a model that detects first-party fraud at application, we first had to define first-party fraud and build a taxonomy to ensure consistent targets. And that definition could not be based on intent, since there is no way to measure that. (Even if a definition could be based on intent, it would miss the crux of what SentiLink is equipped to detect: identity-associated fraud risk signals.)

Instead, SentiLink defines first-party fraud like this:

First-party fraud is the act of presenting a financial institution or other lender with accurate core personal information but inaccurate, misleading, or incomplete financial identity information.

For example, an application with elevated risk of first-party fraud might include the applicant's real name, DOB, and SSN, but incorrect income information. Or, the application might contain accurate income information, but leave out that the applicant's credit history has been distorted by claims that debts they incurred legitimately were the result of identity theft.

In other words, the fraudster is effectively creating a fake financial identity for the FI to lend to — one that shares the fraudster's real name and SSN, but not their true financial and credit history.

Detecting First-Party Fraud

The problem

So why is first-party fraud so tough to detect? The definition problem discussed above is a core reason, but there are quite a few others, too:

The core identity information provided on the application is accurate. Identity theft and synthetic fraud can be difficult to detect, but they virtually always contain some core PII element (SSN, DOB, phone, email, IP address, etc.) that doesn't belong. These inconsistencies are more rare with first-party fraud.

There are no authoritative signals. With identity theft, there is a victim who can report the crime. With synthetic fraud, there are authoritative services such as eCBSV that can confirm the veracity of an identity. First-party fraud almost never comes with an answer this clear.

Outcome only manifests downstream. Identity theft and synthetic fraud can be very obvious — even provable — at the point of application. But first-party fraud is more nebulous. Until an account has actually charged off, for example, institutions can only assess relative fraud risk signals. Even after an account has charged off, it can be difficult to determine with certainty whether it was first-party fraud.

First-party fraud is easily re-branded. The past several years have seen a spate of first-party fraud spikes spread through social media as "glitches" or "hacks." Some first-party fraud methodologies are actively marketed to consumers as forms of legitimate credit repair. A viral video can leave an institution inundated with first-party fraud overnight, with hundreds or thousands of applications from consumers whose history may not indicate risk — in some cases, these consumers may not even be aware they're committing fraud.

First-party fraud signals cross industries, institutions, and years. In identity theft and synthetic fraud applications, the evidence of fraud is often in the application itself. First-party fraud signals, on the other hand, tend to become more obvious over time and across industries. A single company assessing an application generally doesn't have enough of a window into the application identity's full history to see those signals, so they're blind to the threat.

Formal consortiums don't go deep enough

Ultimately, first-party fraud isn't something that institutions can identify and stop effectively using purely their own data. The window of a single institution is simply too narrow to get a clear signal. But sharing data with other institutions doesn't work well either — as discussed earlier in this paper, inconsistent definitions and labeling practices make it extremely difficult to leverage other companies' outcome labels in a way that meaningfully impacts first-party fraud losses.

Even if an application looks like it presents a high risk of first-party fraud, it won't be labeled as such until the account busts out, for example, which might be years later. Before a consortium member will be aware of the risk associated with a fraudster, all of the following things need to have happened:

- Fraudster applies to a consortium member institution (and consortiums are often isolated to one industry)
- Fraudster busts out or otherwise commits first-party fraud as determined by that institution's internal definition (typically months or years after application)
- Institution labels the account a fraud loss

Even in a best case scenario, the "bad actor" flag simply takes too long to trickle down to member institutions for it to be particularly useful, even if they could have complete confidence in the accuracy of other members' labels.

Finally, formal consortiums require sharing application and performance data with a wide variety of institutions — often including competitors. This introduces a variety of legal, bureaucratic, and strategic obstacles that make the idea unworkable for many companies.

Flagging First-Party Fraud Risk More Precisely

To surface and reduce first-party fraud risk at the point of application, then, it's more effective to turn to identity attributes and patterns that are associated with common first-party fraud tactics.

Information captured in an application can provide indications that it warrants further review. For example, an application identity might be associated with opening four deposit accounts in a very short window, which signals some level of fraud risk. This could be one factor that invites further scrutiny.

This is the foundation upon which SentiLink's First Party Fraud Score model was built.

Measurable signals, consistent labels

Like all of SentiLink's fraud risk scoring models, effective risk signaling begins with a consistent taxonomy and consistent labeling.

Consistent taxonomy: Years of assessing application fraud risk across 500+ partners in industries ranging from traditional finance to telecom to education to utilities has led to a simple taxonomy that broadly classifies identity fraud across three axes of measurement:

- Cohesion — how real and consistent is the application identity?
- Control — who is behind the application?
- Impact — is the behavior malicious (like identity theft) or benign (like a typo in the SSN or DOB)?

Every application is assessed across these three axes and further broken down into distinct labels which broadly define the fraud category (like "first-party fraud") and sub-labels, which describe the MO at a more granular level (like "referral abuse").

Some examples of fraud risk signals assessed by our model because they are empirically associated with higher levels of first-party fraud risk include:

- **High cross-network application velocity:** Application identity is tied to an unusually high number of applications across distinct industries and/or institutions in a short period of time.
- **High velocity of new DDA account applications:** Application identity is tied to an unusually high number of DDA account opening applications across SentiLink's network.
- **History of synthetic fraud:** Application identity is tied to past use of SSNs that are not their own.
- **Short-history phones:** Application identity is tied to an unusually high number of distinct phone numbers, with each used only for a short period of time.
- **Temporary phones:** Application identity has a history of using temporary or disposable phone numbers, or a phone number that was recently ported to a new carrier.
- **Short email history:** Application lists an email with little or no history of use in previous applications of that application identity.

(This is far from a complete list.)

Consistent labeling: SentiLink's First Party Fraud model, like all of our scoring models, is trained in part on labels from our Fraud Intelligence Team, which manually labels nearly 100,000 applications every year, and evaluated against real-world scenarios.

This approach directly addresses the inconsistent definitions and labeling discussed earlier in this whitepaper, and the results speak for themselves. Averaged across recent retrostudies, SentiLink's First Party Fraud Score prevented more than $7.5M in annualized fraud losses while flagging just 2.6% of booked loans per company.

And that's just an average; at larger companies the prevented losses were much higher. Across retros assessing all identity fraud types, SentiLink found that high-risk first-party fraud applications accounted for an average of 72% of the preventable identity fraud losses that were on organizations' books.

The model, built on real-world data and tens of thousands of labels from our FIT team, has already proven highly performant against millions of real-world applications across a wide variety of industries and first-party fraud MOs.

What the First Party Fraud Suite gives you

895 First-Party Fraud Score	Easy-to-understand score from 1-999. Higher numbers indicate more elevated risk for first-party fraud.
▪ DDA application activity in SentiLink's network data ▪ High velocity of applications across industry categories in recent history ▪ Whether the supplied phone number corresponds to a risky carrier or line type	Reason codes provide explainability and highlight signals that make the application look more or less risky.

SentiLink's **First Party Fraud Score** provides a clear, explainable assessment of the level of first-party fraud risk associated with an application, based on the signals both within that application and within linked applications across SentiLink's network of 500+ partners.

The score is portfolio-agnostic, meaning that it takes risk signals across product types and industries into account, and it has proven performant in the real world across a spectrum of each.

SentiLink's First Party Fraud solutions also include **High Risk Fraud Flags**, boolean flags that indicate high risk for specific first-party fraud MOs including:

- Serial disputers: Applications that demonstrate patterns consistent with repeated tradeline disputes (sometimes referred to as "credit washing").
- Authorized user purchasers: Applications using addresses associated with authorized tradeline sales.

Combined with our Identity Theft Score and Synthetic Fraud Score, the First Party Fraud solutions are designed to flag risk, at the point of application, across key identity fraud categories. These scores, reason codes, and flags make it easy to identify risky signals and route applications to appropriate step-up treatments accordingly.

What first-party fraud risk means, and what it doesn't

SentiLink's First Party Fraud Score should be considered a risk signal. It is not a determination that an application is fraudulent or that an applicant has committed, or intends to commit, fraud.

Like all SentiLink scores, the First Party Fraud Score can be used for the prevention of fraud, unauthorized transactions, claims or other liability, as contemplated under the Gramm-Leach-Bliley Act (GLBA). It assesses identity signals associated with first-party fraud risk. It does not attempt to — and cannot be used to — assess creditworthiness or ability to repay. It is used to assess fraud risk and inform fraud treatment paths, not to make credit decisions or any other action within the scope of the Fair Credit Reporting Act.

Application treatment should be tiered based on risk signals. A high First Party Fraud Score indicates that the application at hand warrants further scrutiny, and should be routed to the most-appropriate step-up treatment, not declined outright.

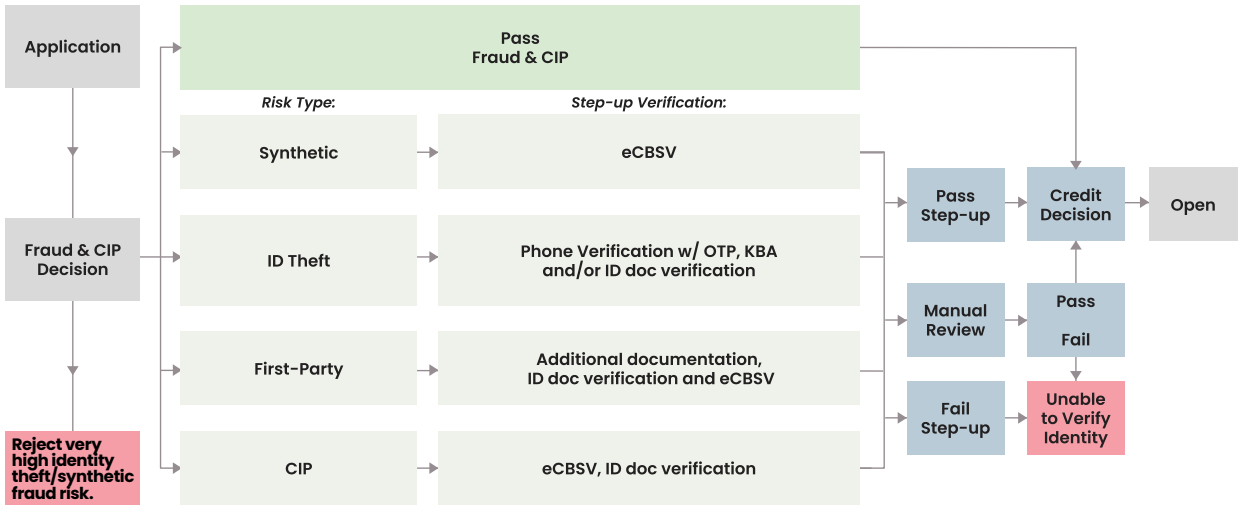
Operationalizing the First-Party Fraud Score

With the implementation of any new fraud control measure such as SentiLink's First Party Fraud Score, organizations must take care to ensure that it improves, rather than detracts from, the onboarding experience. For this reason, we suggest a few best practices:

Put fraud controls at or near the top of your funnel.	This enables you to avoid wasting time or resources on high-fraud risk applications before they've been routed for treatment. Typically, this is straightforward and SentiLink solutions can fit seamlessly into your existing fraud control stack.
Deploy fraud controls and step-up treatments by risk type.	Identity theft, synthetic fraud, and first-party fraud all look different, and a treatment strategy that's effective for one may be useless against another. Assessing for all three at the top of the funnel allows you to maximize efficiency by routing the application to the most appropriate treatment immediately, and also to pass low-risk applications on to the next step in the process with no disruption to the user experience.
Measure and monitor key risk indicators.	Fraud patterns change over time, and your business changes too. Track how all fraud rules and treatments are performing, and track each vintage against key risk indicators to flag when changing patterns require a change in your fraud logic or treatment strategies.

Sample and investigate high-fraud risk applications monthly.	Review a sample of high-fraud risk accounts monthly to assess performance and identify improvement opportunities.
Update your model when a new version is available.	Fraud evolves quickly, and fraudsters are always developing new tactics. SentiLink regularly updates our models to keep pace with these developments, and we suggest that partners implement these updated versions in production as soon as practicable.

Here is a sample diagram of what the fraud and identity verification process may look like in more detail. (Note that applications cannot be rejected on the grounds of being unable to verify identity because of a high First Party Fraud Score.)



The most appropriate treatment strategy for applications with high First Party Fraud Scores will depend to some extent on the specifics of your business, but treatment strategies we recommend considering for high-risk applications include:

- **(Non-notarized) Attestation:** can include legal language with digital attestation, which provides a deeper sense of gravity and acknowledgement.

- **Physical ID verification:** for suspected credit washers, mitigate the risk of a future dispute.
- **4506-C/T:** for suspected credit washers or tradeline purchasers, gather additional income data and introduce additional and more complete identity processes.
- **Manual investigation:** Conduct non-documentary verifications, and/or consult other fraud investigation tools under the GLBA to verify identity.

On a technical level, implementation of SentiLink's First Party Fraud Score is easy — simply pass the appropriate information to SentiLink via a REST API call and receive the First Party Fraud Score, along with reason codes and any other SentiLink scores and products associated with that application, instantly.

Getting started with SentiLink's First Party Fraud Score begins with a free retrostudy to check the score's performance against your historical applications to quantify ROI.

Get in touch to get started, or learn more about our retrostudy process.

Conclusion

- First-party fraud is the most difficult type of identity fraud to define.
- First-party fraud is the most difficult type of identity fraud to detect at application.
- First-party fraud losses probably account for most of the fraud losses on your books (but you might not know that they're fraud).
- SentiLink's First Party Fraud Score is performant across a variety of industries and first-party fraud MOs, preventing millions in fraud losses you're missing without adding additional friction for good customers.

Reach out to get started today

About SentiLink

SentiLink, the leading provider of innovative identity verification and fraud prevention solutions, empowers organizations to accurately identify customers and detect synthetic fraud, identity theft, and first-party fraud. Its best-in-class solutions leverage a deep understanding of identity and risk informed by machine learning models and insights from a team of the industry's best risk analysts. SentiLink proudly serves:

13/15
largest U.S. banks

8/10
largest U.S. credit unions

2/3
largest U.S. telecoms

500+
other companies and government organizations


SentiLink