SentiLink

The SentiLink Fraud Report

Identity Fraud Rates & Trends 1H 2026

Table of Contents

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 of the **15** largest U.S. banks

8 of the **10** largest U.S. credit unions

2 of the **3** largest U.S. telecoms

500+ other companies and government organizations

Executive Summary

6.12%
Mean identity theft rate

0.62%
Mean synthetic fraud rate

2.00%
Mean first-party fraud rate

- **Identity theft remained the top identity fraud threat**, driven primarily by high rates in Demand Deposit Accounts (DDA) and Telecom. And although the identity theft rate declined over the course of the half, the mean rate we observed in 1H 2026 (6.12%) is still higher than what we observed in the first half of 2025 (4.81%).
- **Synthetic fraud stayed relatively stable** over the half, averaging 62 bps overall.
- **First-party fraud** averaged 2%, **but was significantly higher in Auto Lending and Telecom**.
- **Identity theft attacks continued to exhibit characteristics of organized fraud campaigns**, including geographically concentrated victim clusters and automated application activity.
- **Fraudsters increasingly used sophisticated residential proxies** that routed traffic through devices near victims' homes, making their activity harder to detect. (They're also increasingly using legacy PII, a trend we covered in the previous edition of this report.)
- **Organized fraudsters and scammers are moving up the value chain**, targeting victims' home equity and retirement accounts rather than trying to scam them for gift cards.
- **Top signals of identity theft risk** across all industries involved the manipulation of the phone number, email address, or IP address linked to an application.
- **Top signals of synthetic fraud risk** included SSN manipulation and patterns in SentiLink's cross-partner activity data that were not visible in the application itself.
- **Rates of all fraud types**, including first-party fraud, **were highest in the early hours of the morning**.

Introduction

You probably know what fraud looks like at your institution. But how does it compare with what others are seeing, both inside your industry and beyond it? What is the nationwide rate of fraud, and how is it changing?

This report represents SentiLink's attempt to answer a very simple question — "what is the fraud rate?" — using data from hundreds of different institutions, analyzed with our labeling methodology to ensure consistency. Specifically, we looked at more than 170 million applications from the first half of 2026 to answer questions such as:

- What are the rates of identity theft, synthetic fraud, and first-party fraud?
- How are those rates fluctuating and changing, within and across industries?
- What are the most common signs of application fraud?
- What do these types of fraud cost institutions when they go unnoticed?
- How are fraudsters' tactics evolving and changing?

As always, our intent is to share what we're seeing transparently so that every fraud, identity, and risk team across the industries we serve will have the information they need to make informed decisions. To that end, we have included full methodology details in the appendix of this report.

What these rates are, and what they are not

This report highlights **the rate of identity fraud attempts, not the rate of successful fraud**. Because these high-risk applications are flagged by SentiLink in real time, the vast majority of them were unsuccessful.

We call these "fraud rates" rather than "attempted fraud rates" because stealing a real identity or creating a fake identity and using it on an application *is* fraud, regardless of whether the application is actually approved.

It is also important to note that identity fraud includes a variety of modi operandi (MOs) beyond the forms of identity theft, synthetic fraud, and first-party fraud (FPF) flagged by our scores. There are thus types of identity fraud that are not included in this report's analysis. Please refer to the appendix for a more detailed discussion of how SentiLink defines fraud subtypes and the methodologies employed for the various analyses in this report.

What's new in this report?

- Fraud rates for Alternative Finance (AltFi) and Buy Now, Pay Later (BNPL) within the Other Consumer Lending industry category.
- Analysis of costs of various fraud MOs, based on real-world charge-off data from 2022 — 2026.
- Additional first-party fraud analysis, including a slightly updated methodology.

Comparisons to previous reports

SentiLink has previously released *Fraud Reports* for 2H 2024, 1H 2025, and 2H 2025. While this 1H 2026 report uses largely the same methodology as previous reports, when making comparisons between reports, please bear in mind that each report's sample is slightly different, as SentiLink adds partners over time and as partner usage changes.

In this report in particular, readers should be aware that during the first half of 2026, a number of partners included in this analysis upgraded from previous model versions to our most current Identity Theft Score and Synthetic Fraud Score models.

Additionally, when compared to our 2H 2025 report, this report uses a newer version of our First-Party Fraud Score and a slightly different methodology, so the first-party fraud rate results are not directly comparable.

Overall Identity Fraud Rates

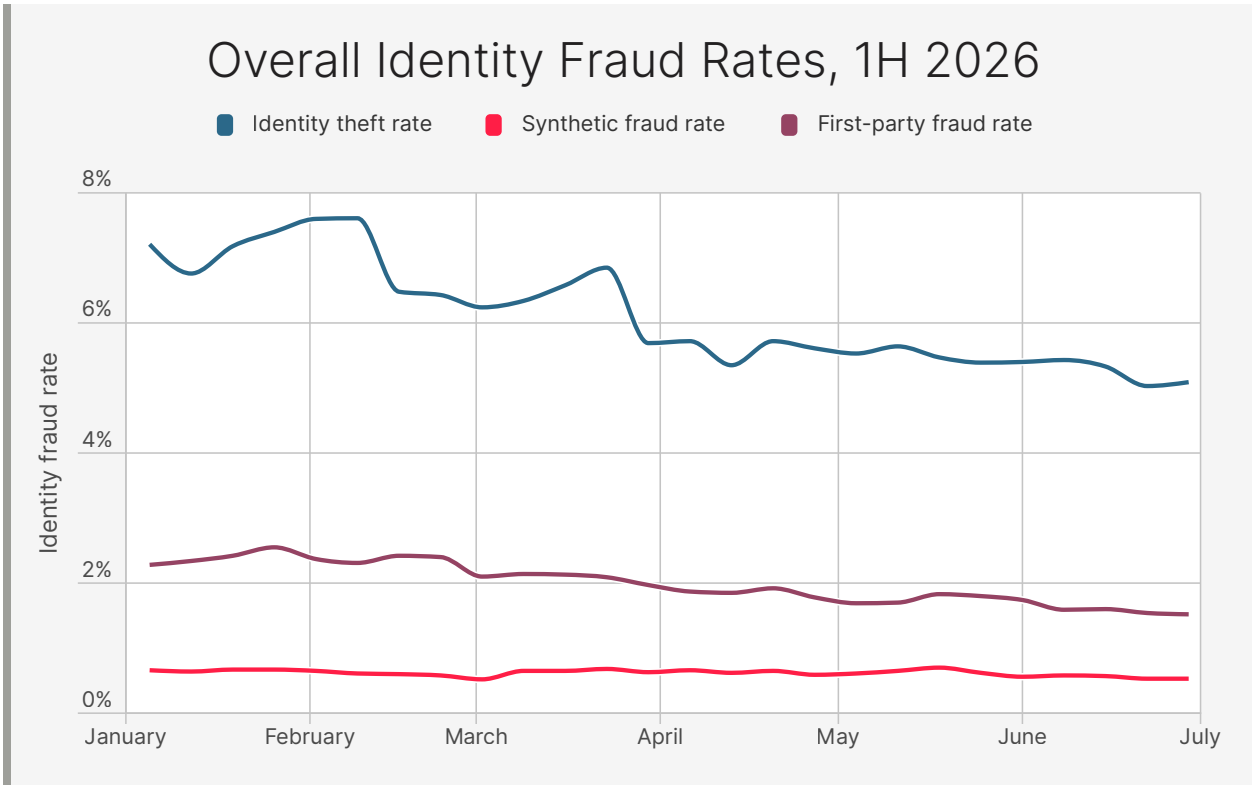

6.12%
Mean identity theft rate

0.62%
Mean synthetic fraud rate

2.00%
Mean first-party fraud rate

In the first half of 2026, we observed **the highest overall identity theft rate** we've reported to date: 6.12%.

At the same time, though, this rate dropped significantly over the course of the half. The primary driver of this drop is the banking industry, which has seen a decline from the extremely high rates observed during the winter of 2025-2026. Several other industries have also experienced less dramatic declines in identity theft rates. This drop may represent a reversion to the mean; the overall mean identity theft rate during the last two months of the half — May and June — was 5.37%, quite close to the 5.61% rate we observed during the second half of 2025.

We also observed a much smaller decline in the first-party fraud rate. The synthetic fraud rate remained relatively flat over the course of the half, which is consistent with what we've seen in previous reports.

Identity Theft

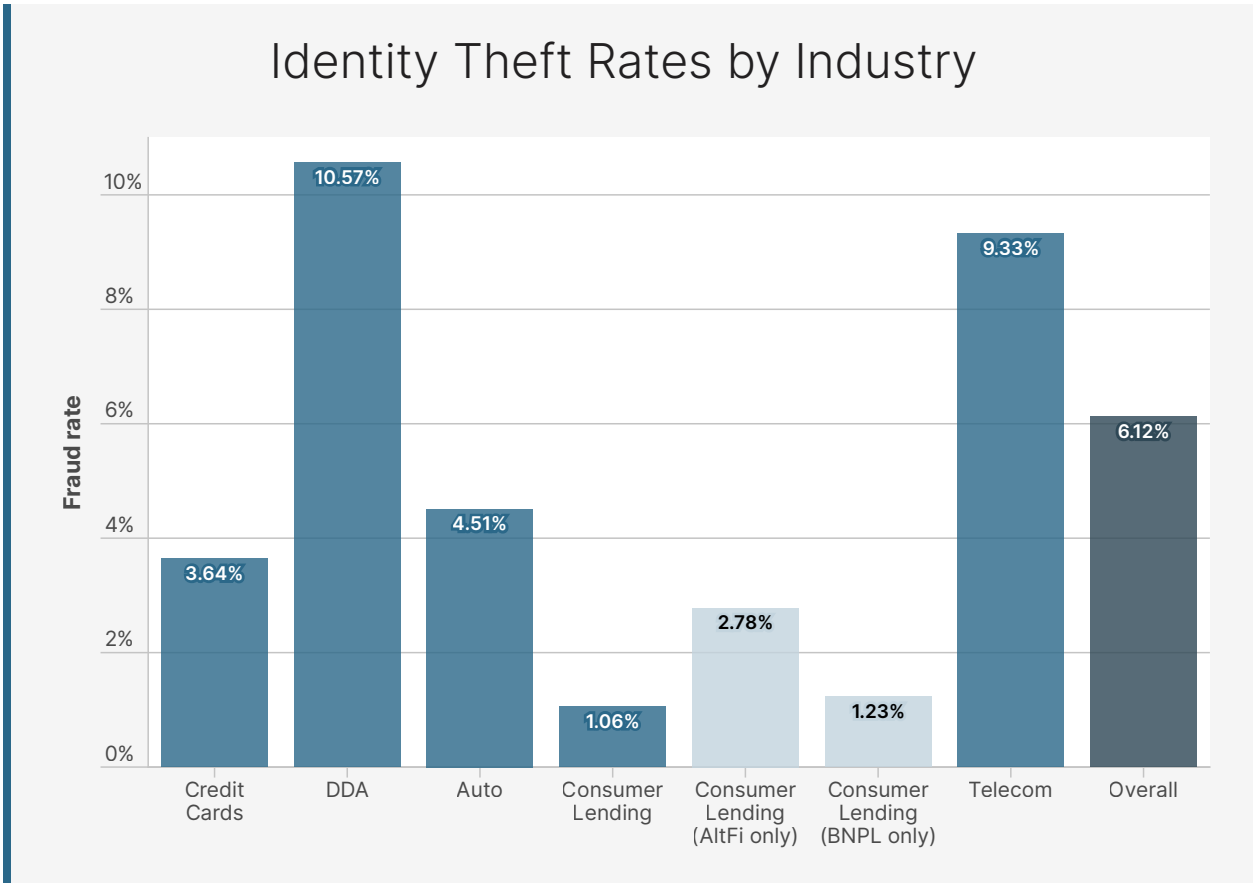

As we have seen in previous reports, the rates of identity theft in DDA and Telecom dwarfed the rates observed in other industries. Rates outside of DDA and Telecom were highest in Auto, followed by Credit Cards. Other Consumer Lending posted the lowest identity theft rates. This distribution of industries is consistent with what we've seen in previous reports, but the rates in most industries were higher in the first half of 2026 than what we observed in our 1H 2025 report.

Where are identities being stolen from?

Identity theft is everywhere. While it is difficult to pinpoint the location of the fraudsters, we can get a window into where victims are located by looking at the rate of high-risk applications from a given county:

Browse an interactive version of this map:

Identity Theft Rates by County

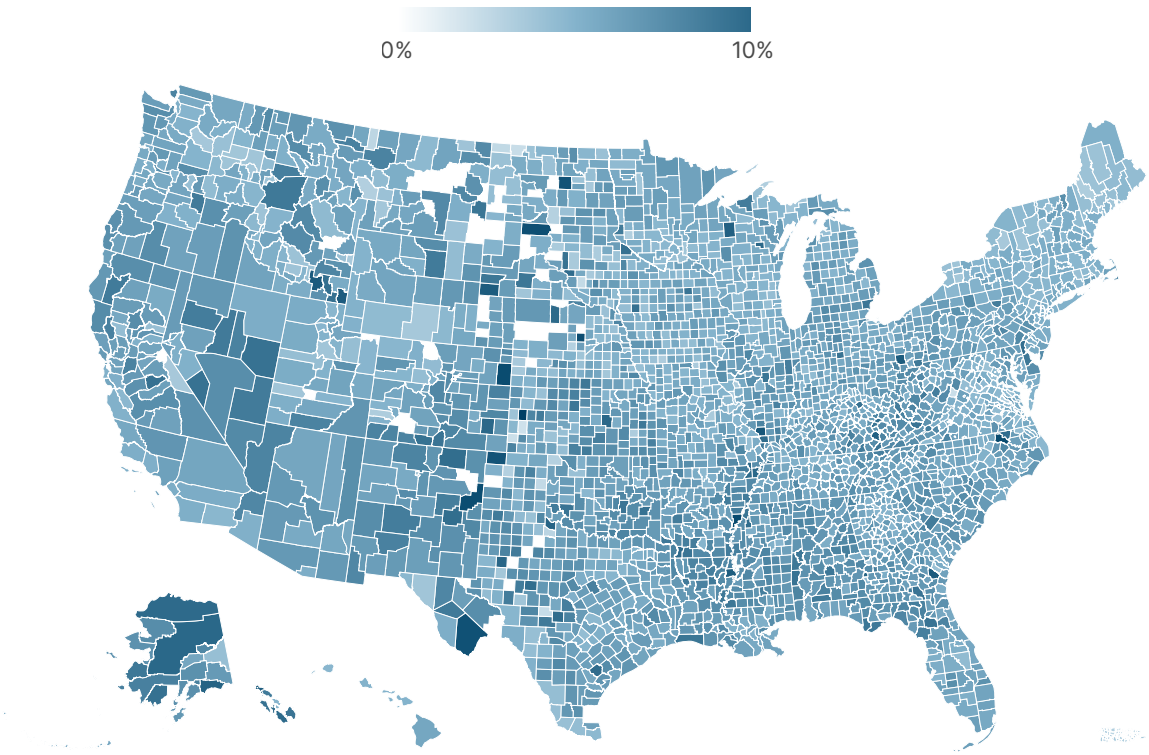

Note: counties with no color shade had insufficient application volume.

Hotspots on this map of identity theft victims likely indicate areas where fraudsters may have acquired a set of local identities that were then leveraged for targeted attacks. For example, let's zoom in on Power County, ID, which has a county-level identity theft rate of nearly 11% (compared to mean and median county identity theft rates of roughly 6%).

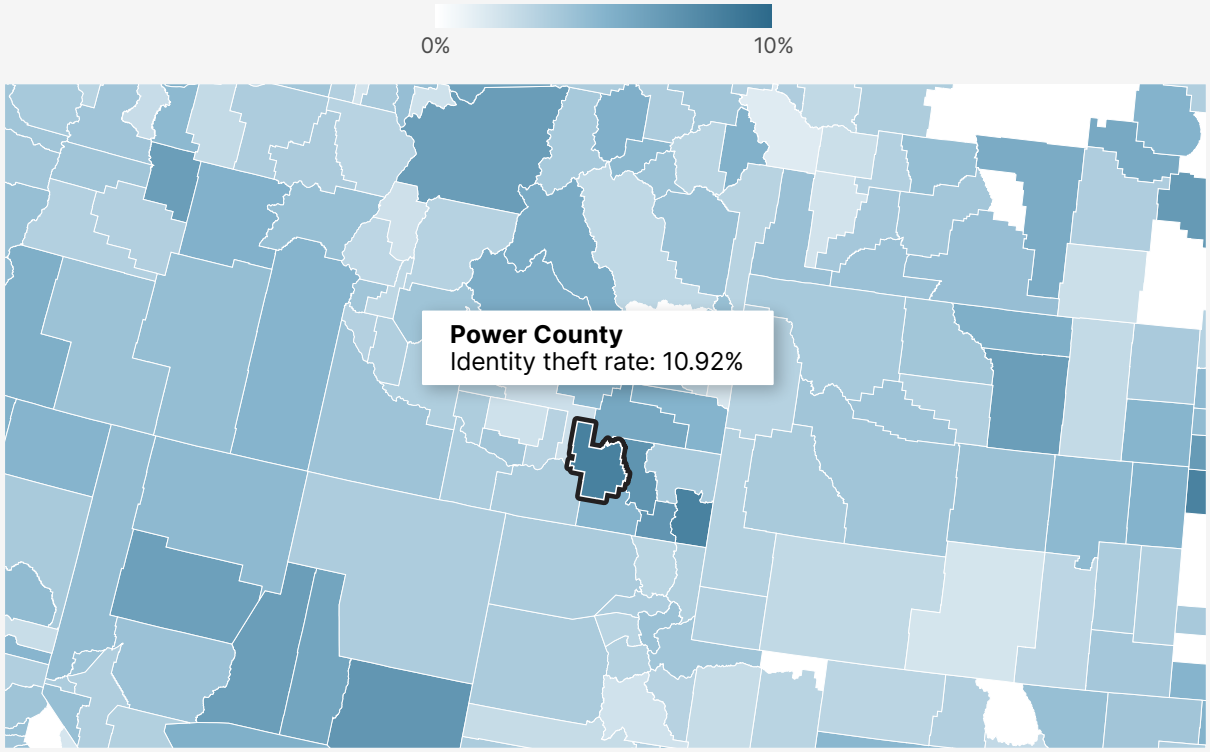

Note: counties with no color shade had insufficient application volume.

This rate is high primarily because identities from Power County and a number of neighboring counties in Idaho were used in a sustained attack that primarily targeted partners in consumer lending.

This attack leveraged a mix of stolen identities and identities that SentiLink calls “IDT synthetic” — stolen identities with one or more incorrect elements of personally identifiable information (PII), such as the wrong SSN, mixed in. It also clearly made use of automated tooling — applications were submitted from a wide variety of IP addresses, most of them hundreds or even thousands of miles from the identity’s home address, and listed obviously automated email handles like USER_5252366642_104352@YAHOO.COM.

The average age of the identities used for this attack was higher than what we observe in applications generally. This is a pattern we observe in many identity theft attacks. In Power County, for example, 58% of the identities associated with high-risk applications were age 55 or older, while just 16% of low-risk applications came from that age bracket.

This is likely a reflection of the fact that organized fraudsters often buy identities from online fraud markets in packages. The identity information in these packages is often sourced from data breaches; in this case, it is likely that the fraudsters acquired a package of identities sourced from the breach of some Idaho-based organization or institution.

Common signals of identity theft

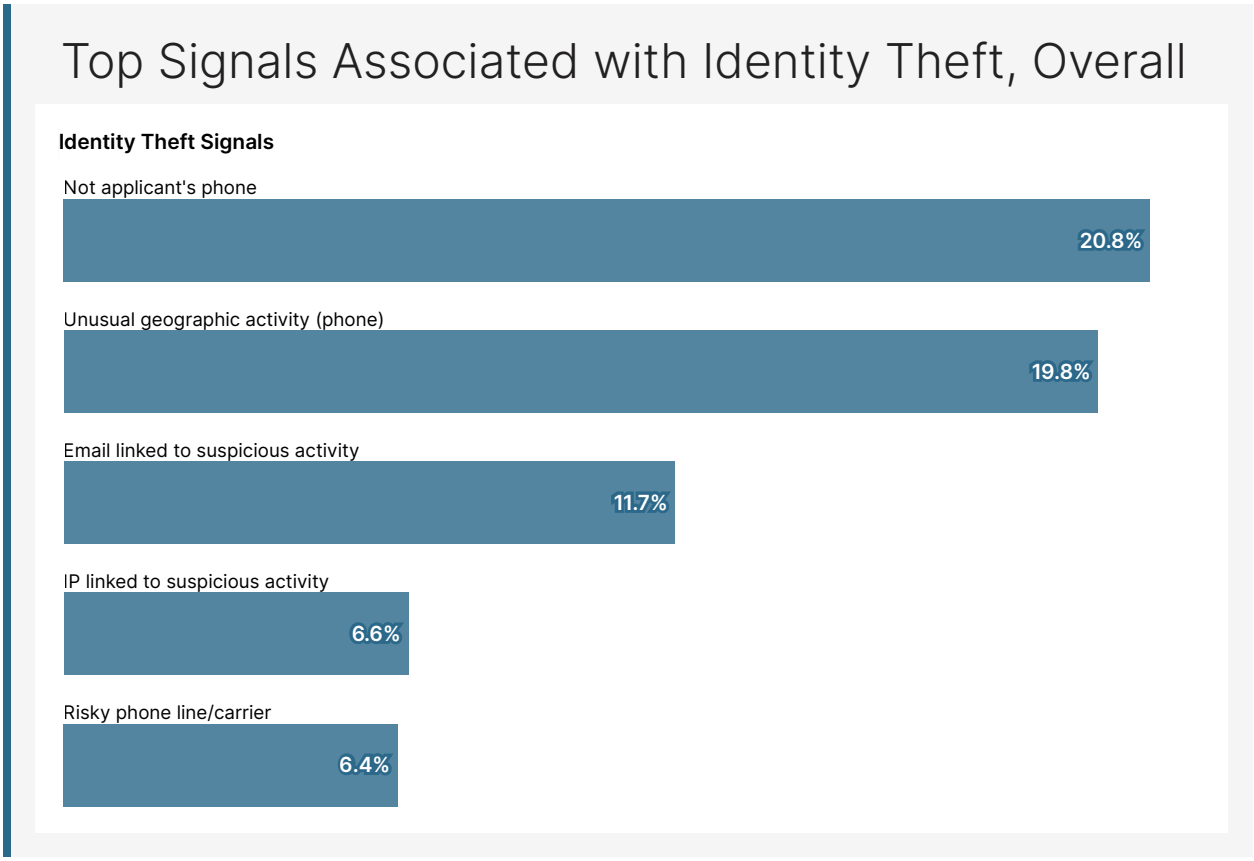

SentiLink's fraud models do not just score applications for risk; they also generate reason codes that indicate which signals in a given application are most indicative of fraud risk. Here, we've shared the five signals that appear most often as the top signal indicating high risk in identity theft applications.

The results here are consistent with what we've observed in previous reports: phone, email, and IP address signals remain the most common red flags when it comes to catching identity theft.

Identity theft risk and ITINs

Individual Taxpayer Identification Numbers (ITINs) are issued to immigrants for the purposes of tax filing, but they share the same nine-digit number format as SSNs. They are thus sometimes used in place of SSNs on credit and DDA applications, presumably by immigrants who may not yet have an SSN.

We are sometimes asked whether applications with ITINs represent an elevated identity theft risk. The answer in the first half of 2026 is the same as it was in 2025: no.

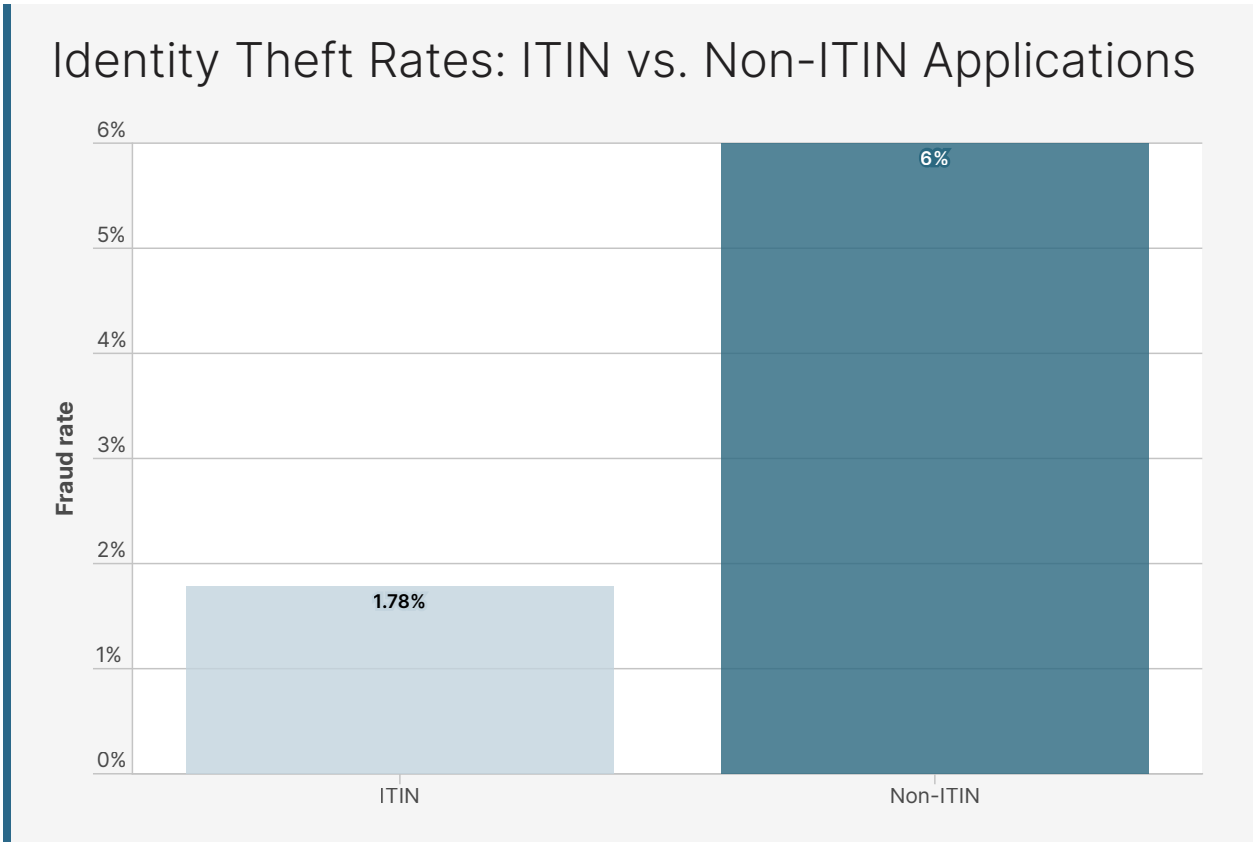

This finding also held across each of the individual industries that we examined for this report; in all cases, applications with ITINs represented a lower risk of identity theft than applications listing non-ITINs (primarily applications with SSNs).

Industry	ITIN Identity Theft Rate	Non-ITIN Identity Theft Rate
Credit Cards	2.02%	**3.52%**
DDA	2.16%	**10.10%**
Auto Lending	1.02%	**4.30%**
Other Consumer Lending	0.62%	**1.82%**
Telecom	1.65%	**2.25%**

Temporal patterns in identity theft

We also looked at identity theft rates by day of the week and by time of day across all industries. As in previous reports, we have observed no interesting trend in terms of day of the week; the higher Sunday rates are simply an artifact of the fact that many SentiLink partners are closed and do not process applications on Sundays.

When looking at fraud rates by hour of the day (normalized to U.S. Eastern time), we once again found that fraud rates were highest during the late night/early morning hours.

Identity Theft Rate by Time of Day (U.S. Eastern)

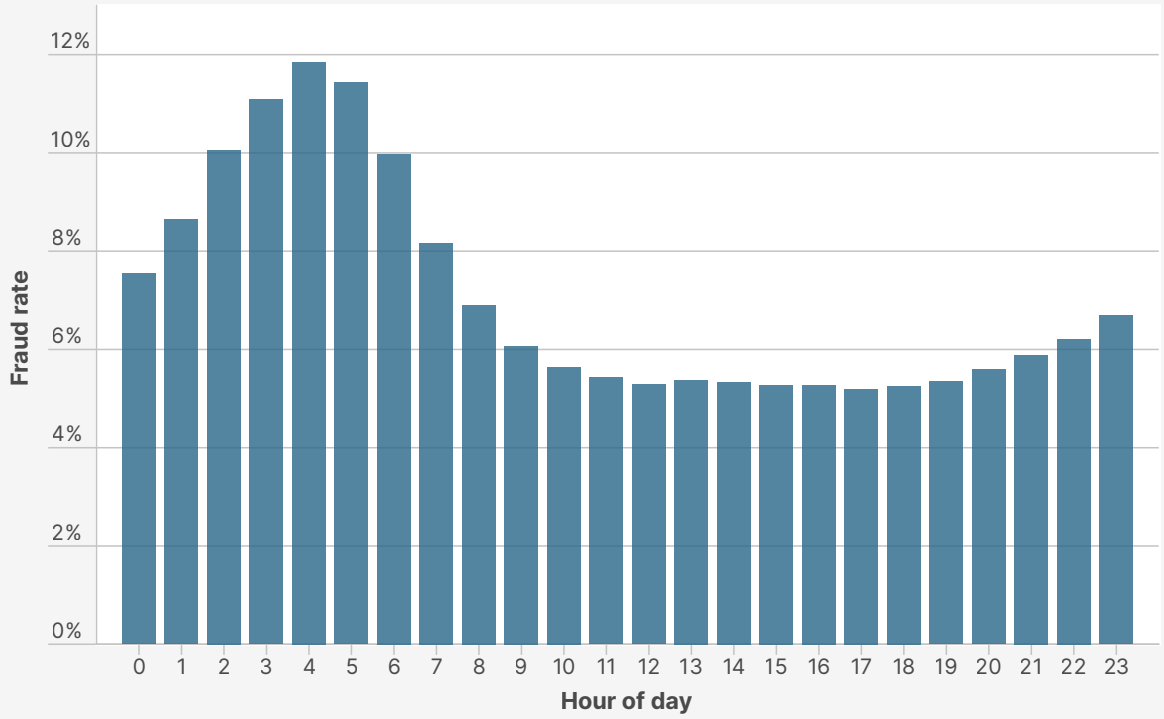

Identity Theft Rate by Day of Week

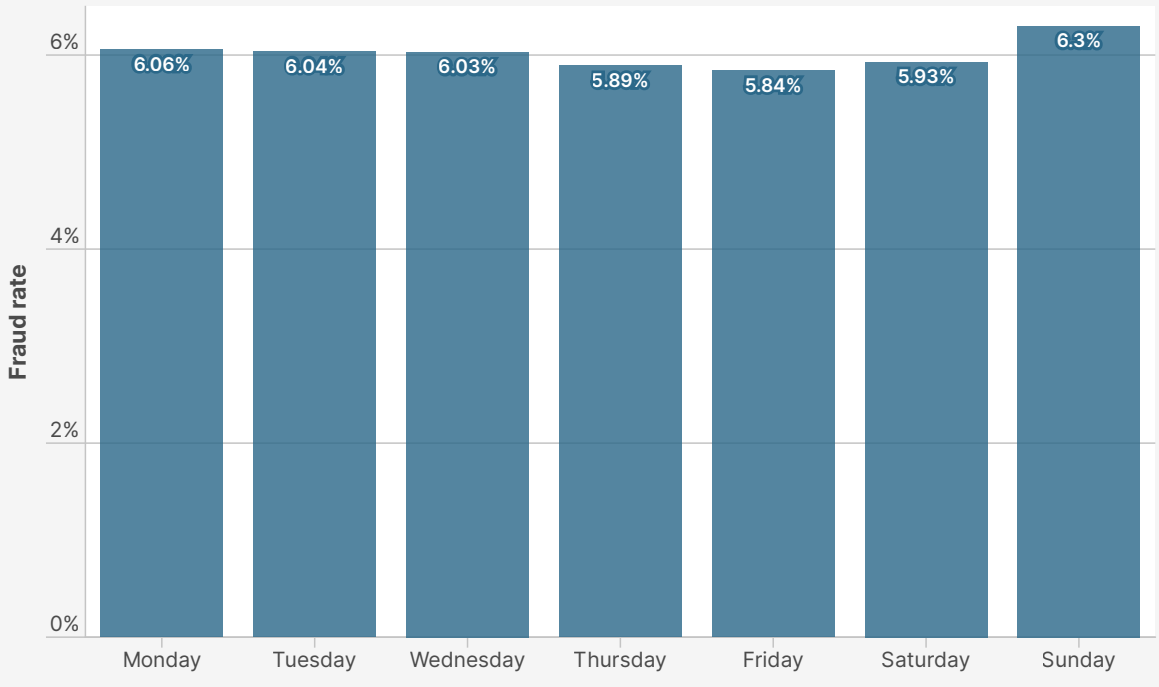

Victim age patterns in identity theft

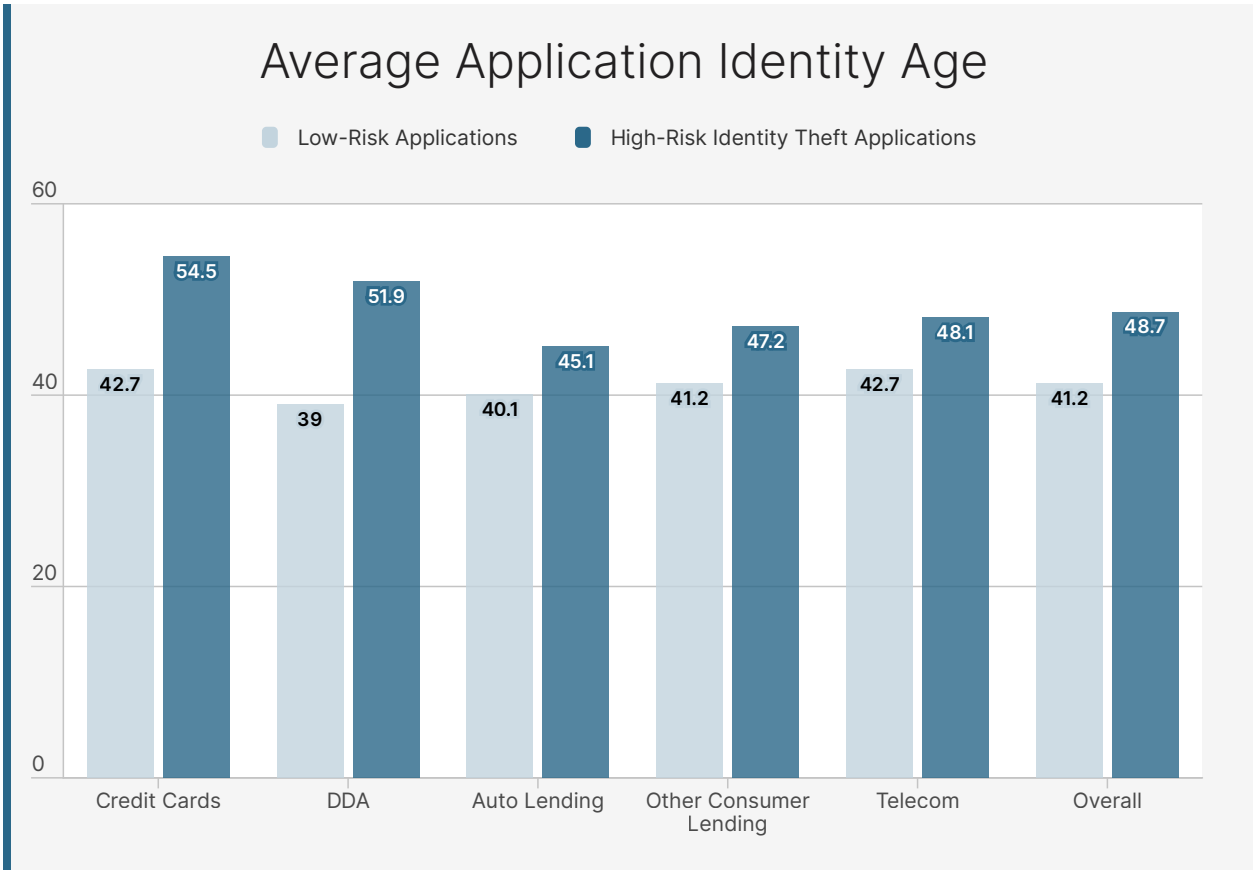

While people of *all* ages can be victims of identity theft, fraudsters leverage stolen identities that tend to skew older. While in some cases this might be an intentional choice (older identities may have better credit), it is probably mostly a reflection of the fraud supply chain. The stolen identities used for most identity theft are sourced in large part from hacks and data breaches and then sold in fraud markets on the dark web, on fraud-friendly Telegram channels, and even on the open web. And because older people have been credit active for longer, their identities are more likely to have been compromised in more data breaches, increasing the chances that they have unwittingly entered the identity theft supply chain.

Note: this chart is for context only. SentiLink models do not use demographic information to assess any kind of fraud risk, including identity theft risk.

Synthetic Fraud

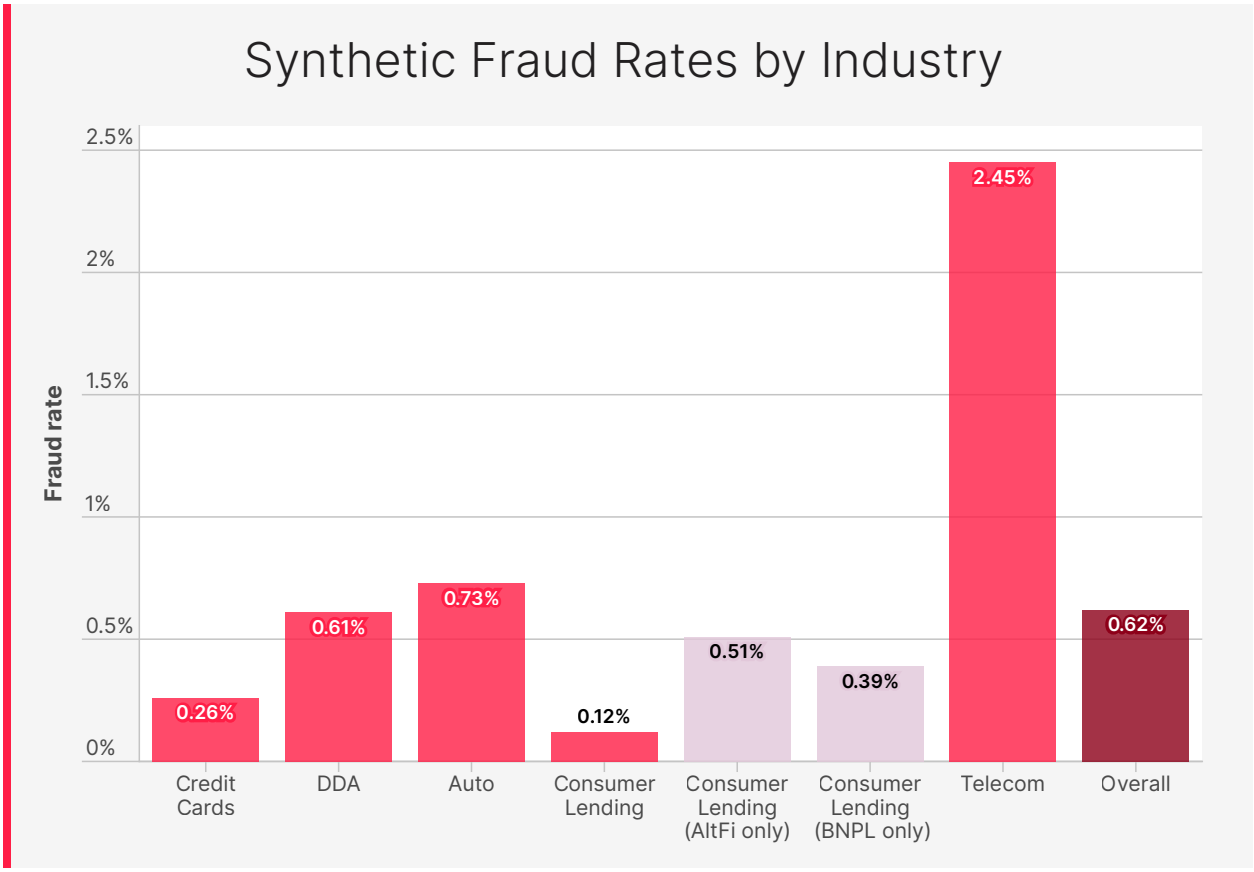

Synthetic fraud rates were low across the board with the notable exception of the Telecom industry. The potential causes of this elevated rate are discussed in more detail in the Telecom section, but this finding is consistent with what we have observed in previous reports.

Synthetic fraud subtypes

Synthetic fraud comes in many forms, and for our analysis, we also looked into the breakdown of applications between the two primary subtypes of synthetic fraud: first-party synthetic fraud and third-party synthetic fraud.

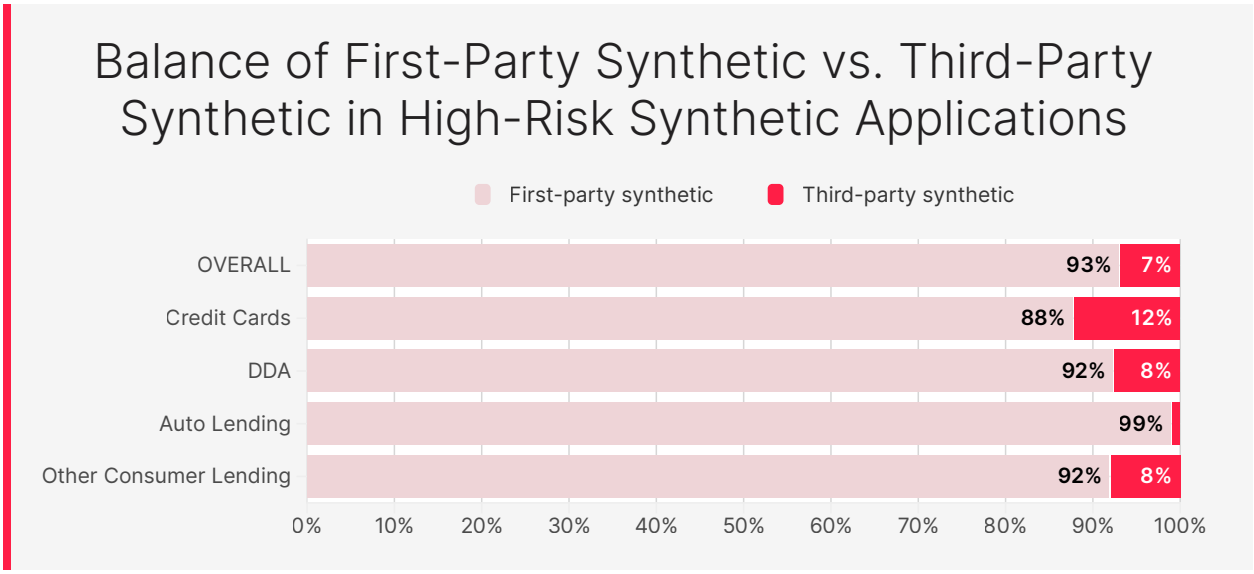

First-party synthetic fraud — in which the fraudster is manipulating their own identity by using a different SSN — makes up the majority of synthetic fraud overall. This is unsurprising, in part because the fraud ecosystem actively promotes and markets Credit Privacy Numbers, or "CPNs" (fake SSNs), to consumers as a form of "credit repair."

Third-party synthetic fraud, in which the fraudster creates an entirely fake identity, is rarer. Creating a wholly new fake identity and building a credit profile for it is more involved than simply replacing your SSN with a CPN. In some contexts, busting out with a third-party synthetic identity would also require forged identity documents.

For example, the rate of third-party synthetic applications in Auto Lending is very low because dealerships process auto lending applications in person. Applying with a wholly synthetic identity would require, at a minimum, a forged driver's license. In contrast, an applicant committing first-party synthetic fraud could simply use their real driver's license to confirm their identity and then provide a fake SSN.

Conversely, third-party synthetic fraud rates are comparatively high in Credit Cards. Applications are often completed online, and a low-limit or secured credit card is a common first step in establishing a new synthetic identity. Once that identity has built a credit history, credit cards offer some of the most profitable opportunities for a final bust-out.

Where is synthetic fraud happening?

The addresses provided in synthetic fraud cases often correspond to the fraudsters' own locations. So while our identity theft map provides an approximation of victim locations, the Synthetic Fraud Rates by County map below likely indicates where *fraudsters themselves* are located:

Browse an interactive version of this map:

Synthetic Fraud Rates by County

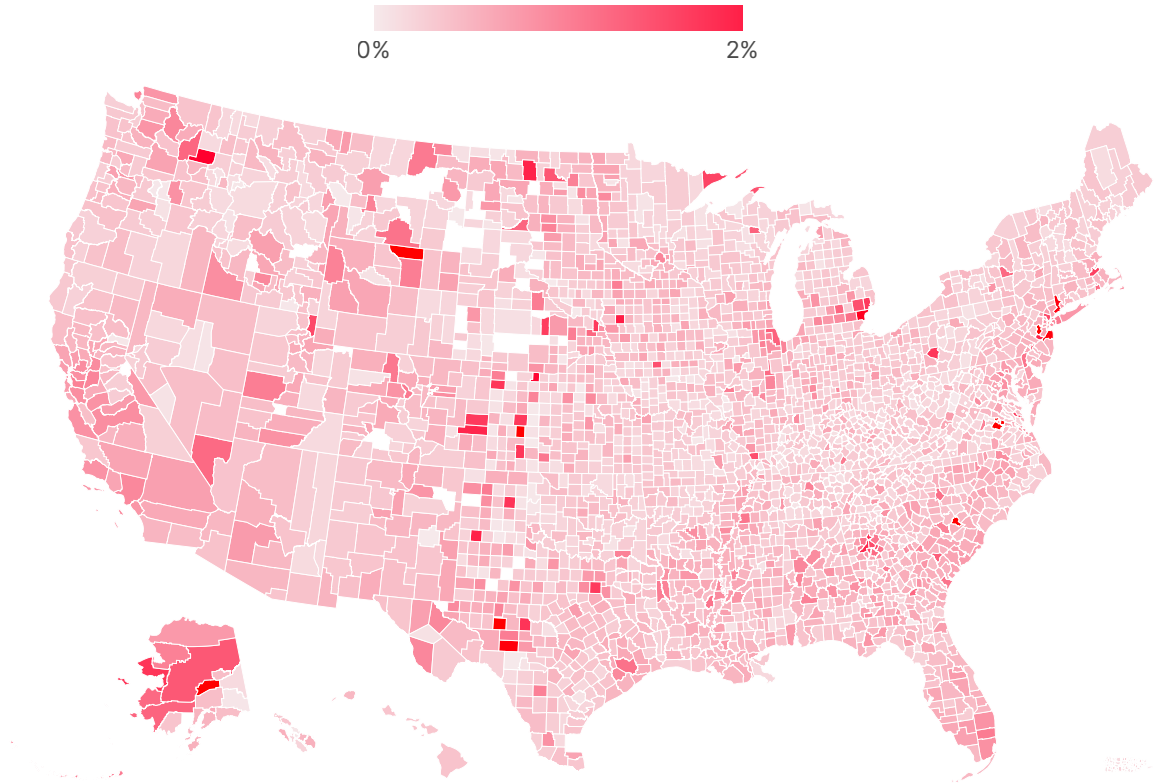

Note: counties with no color shade had insufficient application volume.

This map contains some familiar hotspots, such as Sheridan County, WY, which we wrote about in detail in our 2H 2024 Fraud Report. But let's zoom in on a different hotspot that we've observed before but haven't highlighted: southeastern Michigan.

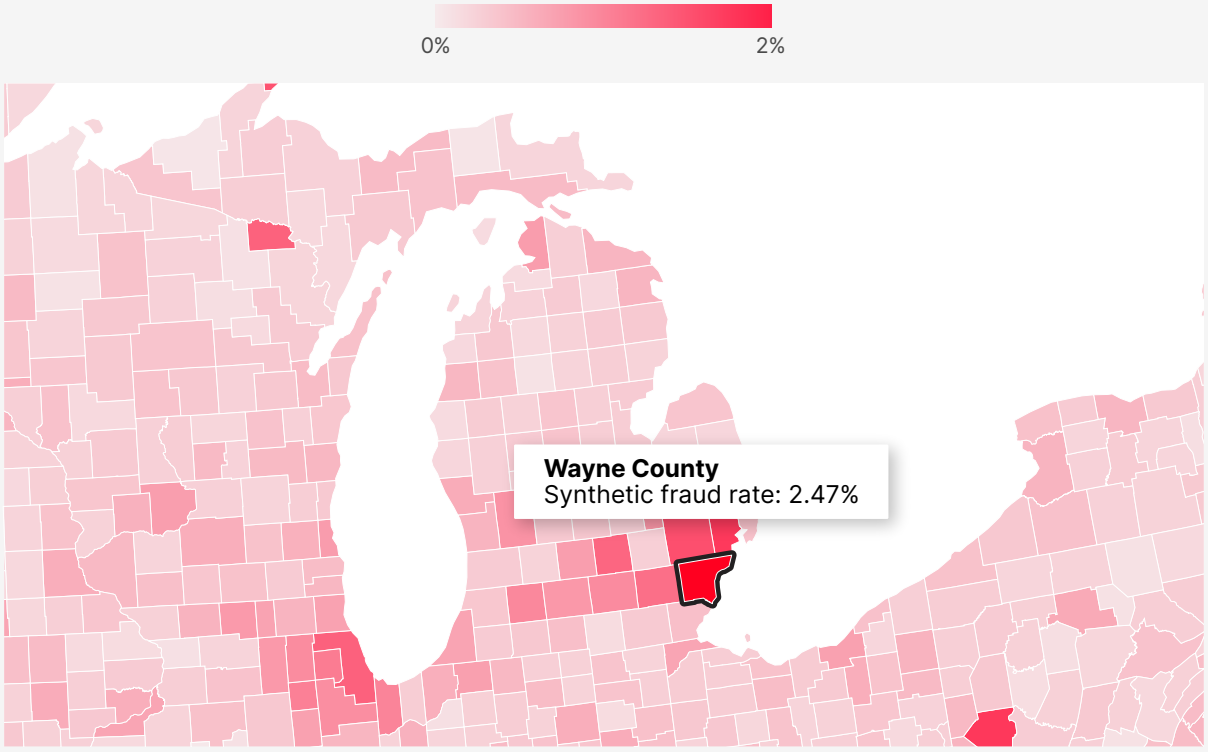

Note: counties with no color shade had insufficient application volume.

Wayne County, the location of Detroit, is another area where we have consistently seen elevated synthetic fraud rates. During the first half of 2026, its mean synthetic fraud rate was 2.47%, nearly 6X the mean per-county synthetic fraud rate.

High-risk synthetic applications from Wayne County have primarily targeted mobile carriers, consistent with the elevated synthetic fraud rates seen across Telecom. They are mostly first-party synthetic applications, meaning the applicant is using their real base identity in combination with a fake or stolen SSN. Notably, these high-risk applications tend to be associated with younger-age identities; 48.8% of high-risk applications came from identities ages 25-34, whereas just 25.4% of low-risk identities came from that age bracket.

It is also worth noting that the county-level aggregation in this synthetic fraud map hides some nuance. For example, synthetic fraud rates in specific zip codes of Burbank and Glendale, CA, remain high — one Burbank zip code has a synthetic fraud rate of well over 3% — but lower rates elsewhere in Los Angeles County bring down the area's overall rate. The reason for elevated synthetic fraud rates in some Los Angeles County zip codes is discussed in 2H 2025 report.

Common signals of synthetic fraud

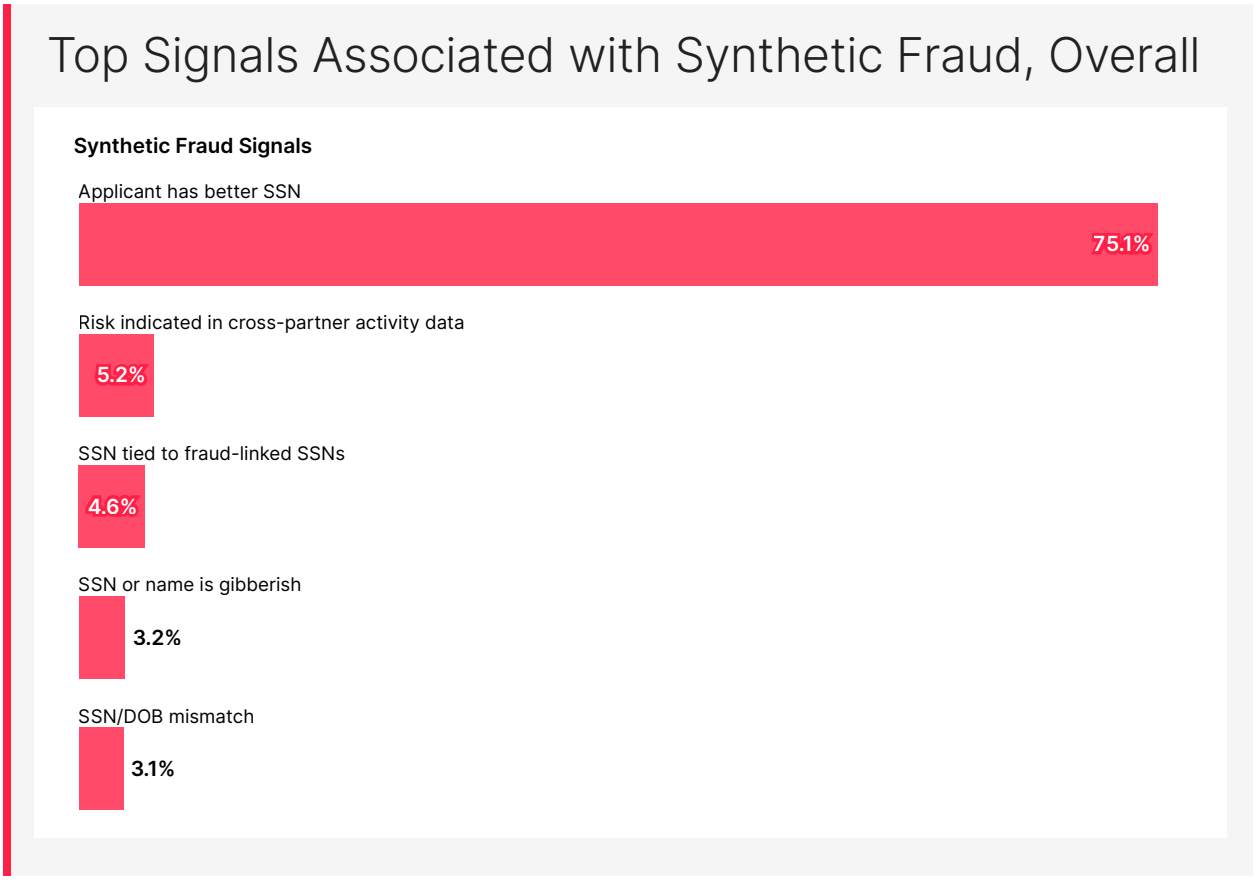

The top signals of synthetic fraud mostly relate to the application SSN, which is to be expected given that SSN manipulation defines synthetic fraud. However, it's worth noting that risk indications in cross-partner activity data — e.g., signs in *other* applications that indicate risk connected to a given application — are also a top signal of synthetic fraud risk.

Temporal patterns in synthetic fraud

We also looked at synthetic fraud rates by day of week and by time of day across all industries. As in previous reports, we have observed no interesting trend in terms of day of the week; the spike in Sunday rates is simply an artifact of the fact that many SentiLink partners are closed and do not process applications on Sundays.

When looking at fraud rates by hour of the day (normalized to U.S. Eastern time), we once again found that fraud rates were highest during the late night/early morning hours.

Synthetic Fraud Rate by Time of Day (U.S. Eastern)

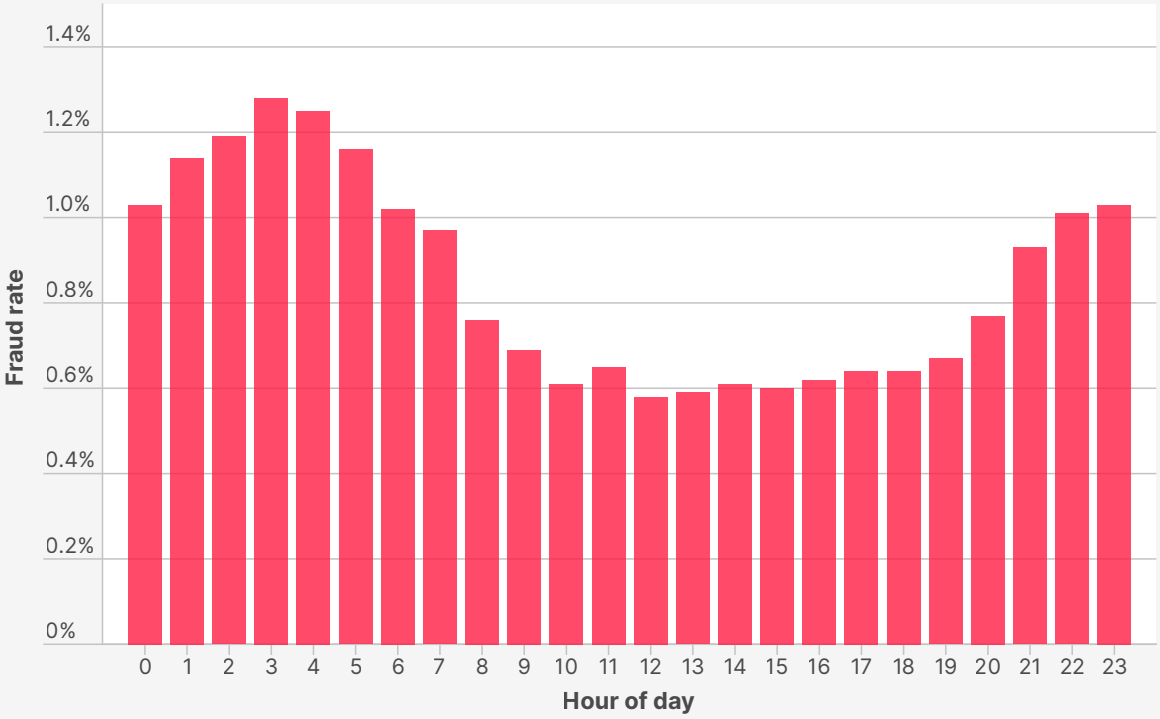

Synthetic Fraud Rate by Day of Week

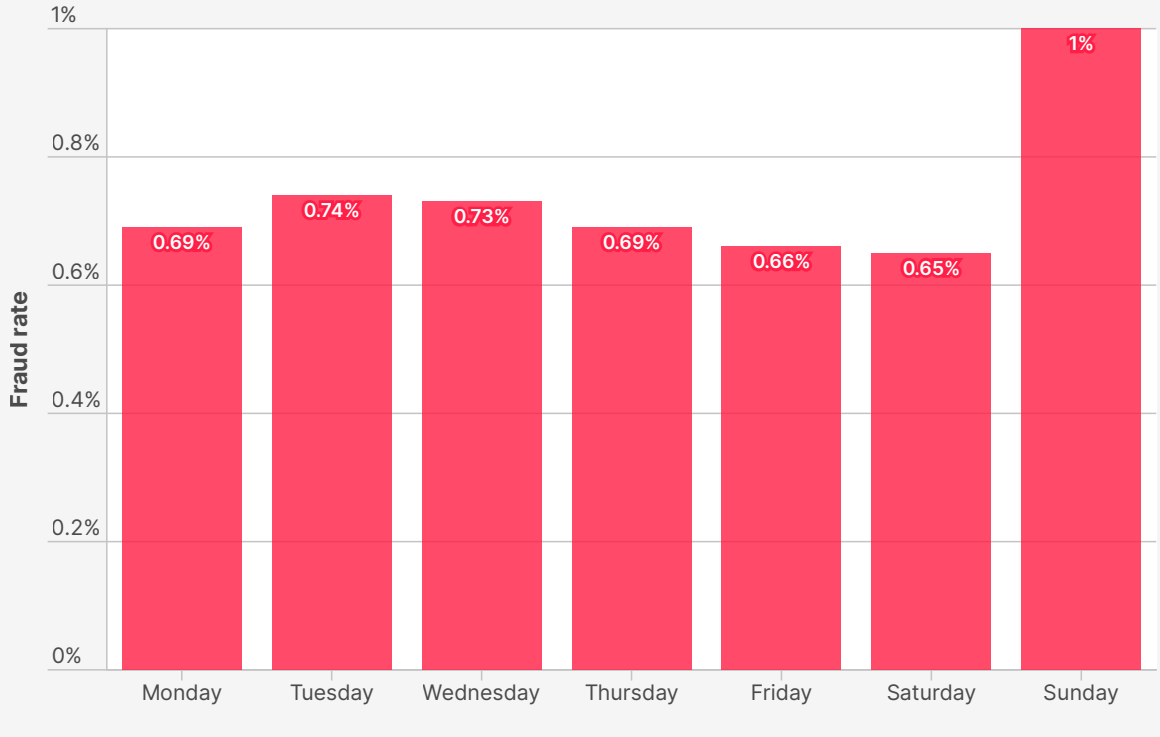

First-Party Fraud

A note about first-party fraud rates: These rates are derived using a different methodology and a smaller sample size than our other fraud rates (please refer to the methodology section for details). Additionally, these numbers are not directly comparable to the first-party fraud rates in our 2H 2025 report because of minor methodological changes and the adoption of the latest version of SentiLink's First-Party Fraud Score model. We feel these results present a more accurate picture of the first-party fraud landscape, but they should still be considered indicative rather than definitive.

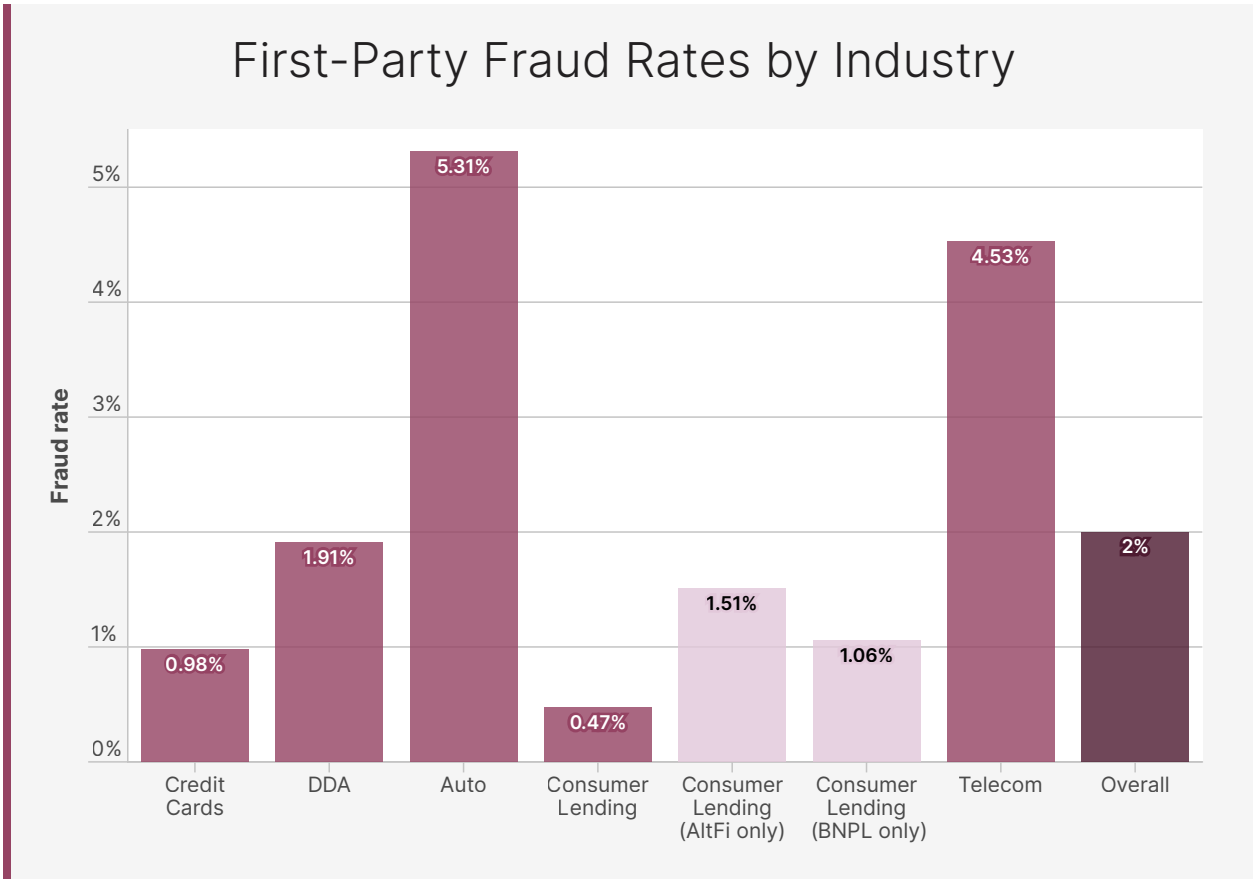

SentiLink's First-Party Fraud Score estimates the likelihood that an application involves first-party fraud across a range of fraud types, including ACH fraud, check return fraud, and credit card fraud. Because those risks vary by industry, elevated scores may indicate different behavior in different contexts. In DDA, for example, they are more likely to reflect ACH or check return fraud, while in Credit Cards, they may be more indicative of bust-out fraud.

The most notable trend here is the strikingly high rate of first-party fraud in Auto Lending. This is not a universal trend; some individual auto lending partners in our cohort had first-party fraud rates as low as 1%. But several partners saw first-party fraud rates north of 10%. This finding is discussed in more detail in the Auto Lending section of this report.

First-party fraud rates were also high in Telecom, although this is less surprising — historically, the telecom industry has seen high rates across all fraud MOs.

Where is first-party fraud happening?

The addresses provided in cases of first-party fraud often correspond to the locations of the fraudsters themselves:

Browse an interactive version of this map:

First-Party Fraud Rates by County

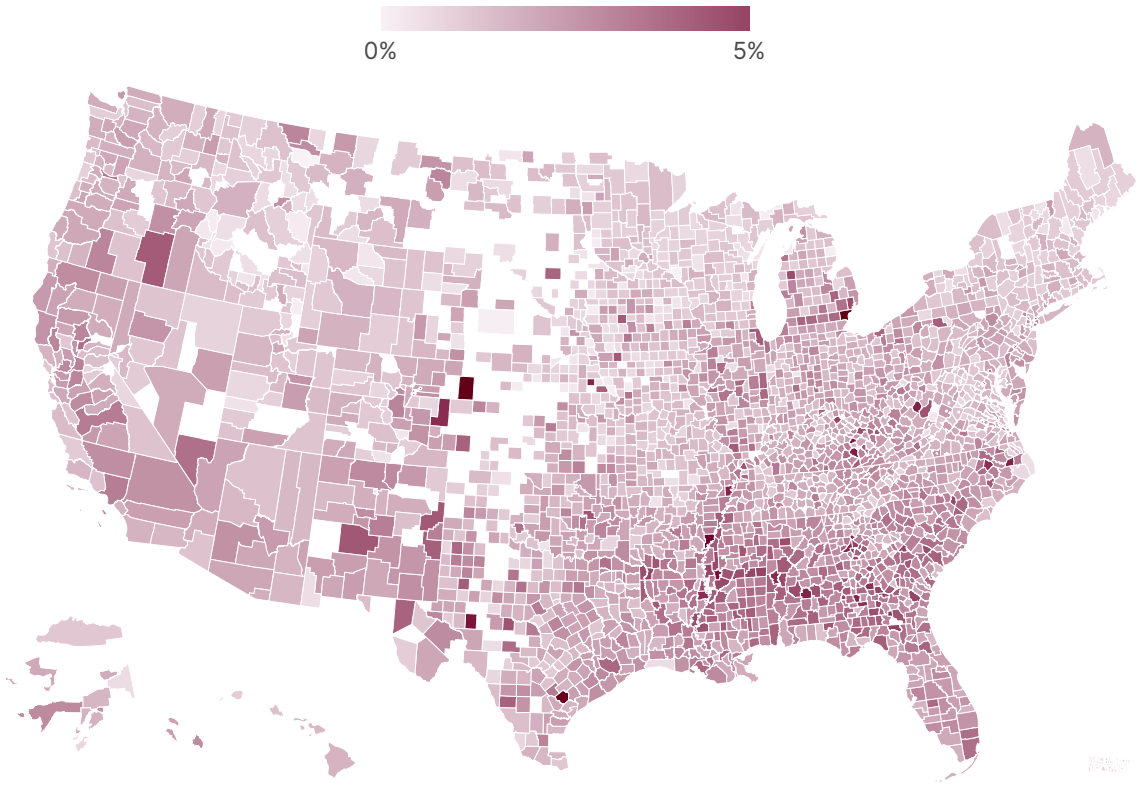

Note: counties with no color shade had insufficient application volume.

First-party fraud is by its nature an individual act, so we don't see organized rings or coordinated attacks in the same way that we might with identity theft (or, more rarely, synthetic fraud).

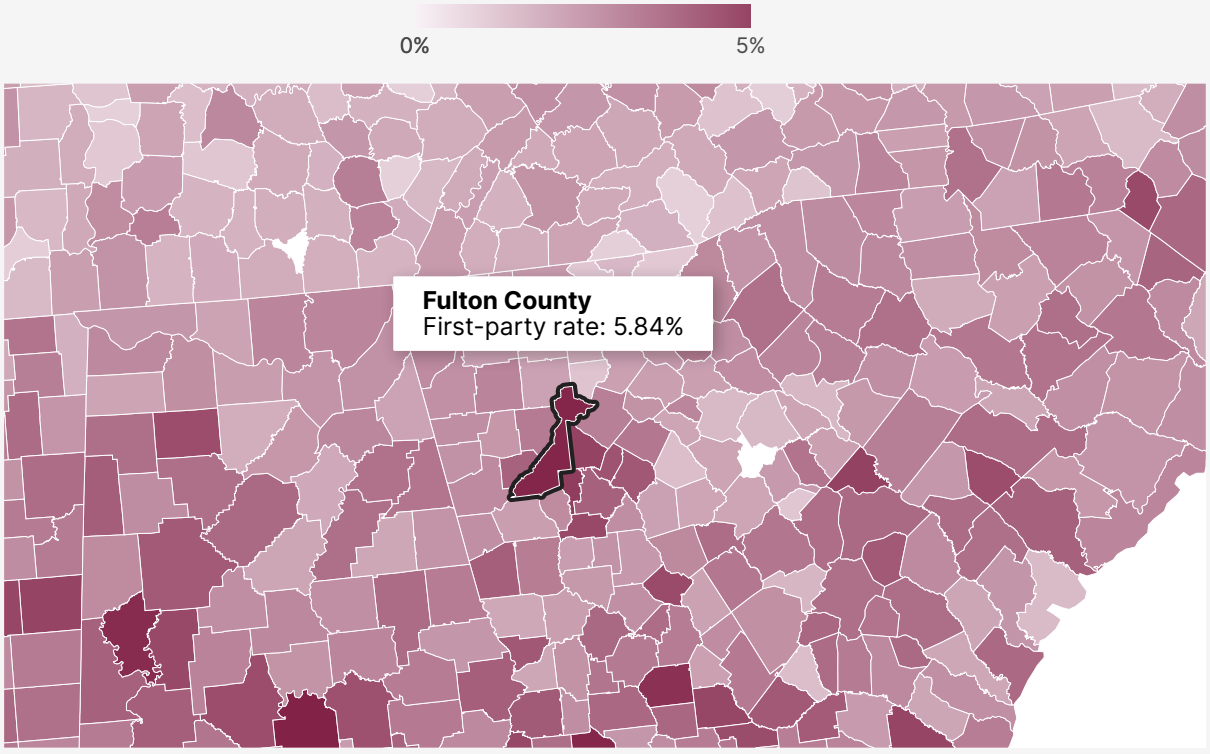

Note: counties with no color shade had insufficient application volume.

If we drill down to a specific location and examine the applications there, we don't find much in the way of cohesive patterns. But one thing we do see consistently is an identity age on high-risk applications that is significantly lower than what we observe for low-risk applications. In Fulton County, GA, for example, 58% of high-risk FPF applications were submitted listing identities between the ages of 18 and 34, compared with just 32% in that bracket for low-risk applications. Overall, the average age on a high-first-party-fraud-risk application from Fulton County was just under 34 years old, compared to an average age of over 43 for low-risk applications. (As noted elsewhere, SentiLink does not consider demographic characteristics like age in its fraud risk assessments.)

First-party fraud risk and ITINs

As with identity theft, we are sometimes asked whether applications listing ITINs represent a higher risk of first-party fraud. The answer, as with identity theft, appears to be: no. Applications that listed ITINs were associated with a slightly lower first-party fraud risk, although the difference is less pronounced than what we saw in identity theft.

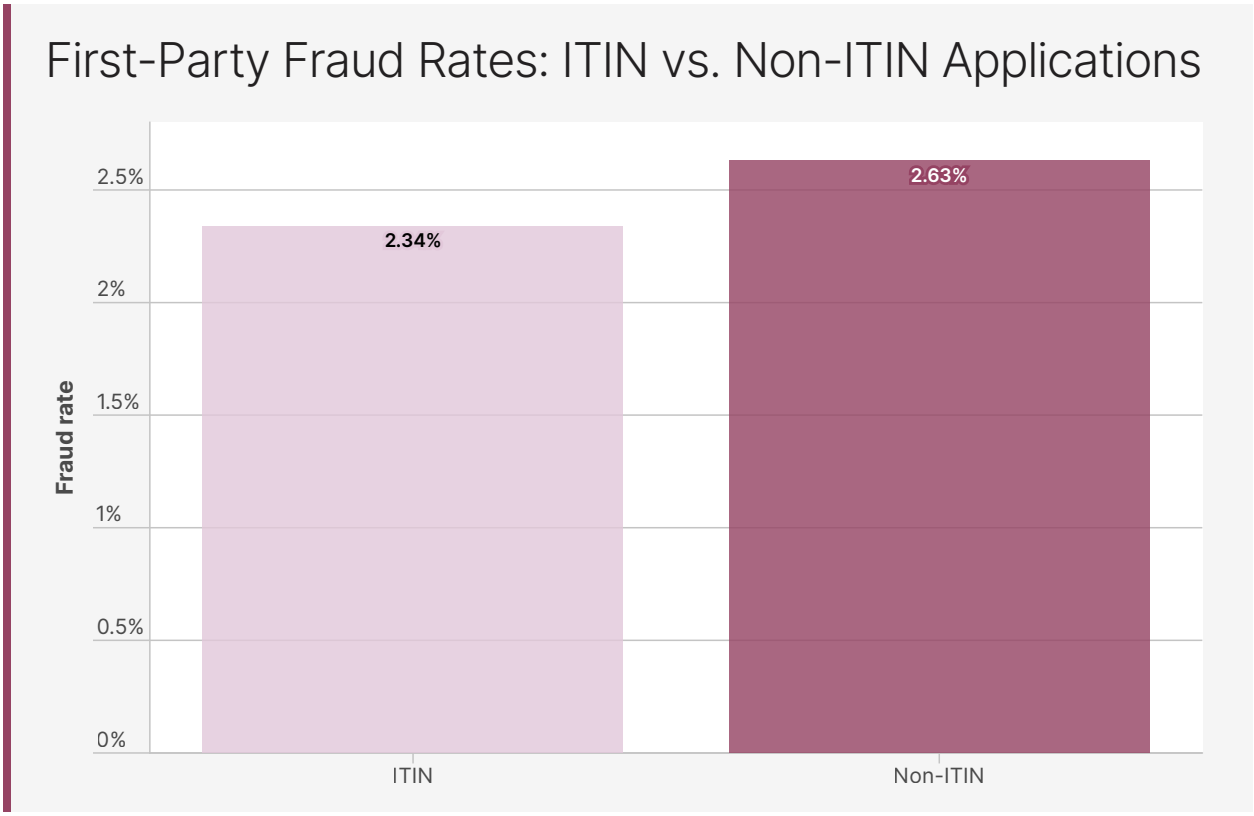

Broken down by industry, there were two industries where the first-party fraud rate was slightly higher for ITIN applications: DDA and Other Consumer Lending.

Industry	ITIN First-Party Fraud Rate	Non-ITIN First-Party Fraud Rate
Credit Cards	1.13%	**1.15%**
DDA	**2.86%**	2.21%
Auto Lending	6.41%	**9.90%**
Other Consumer Lending	**1.09%**	0.82%
Telecom	0.93%	**2.35%**

Temporal patterns in first-party fraud

We analyzed first-party fraud rates by day of week and by time of day across all industries, and found that they do not differ significantly from what we observe with identity theft and synthetic fraud rates: flat rates over the course of the week (excepting Sunday, which is an outlier only because many partners do not accept applications on Sundays, making the sample size for that specific day much smaller), and higher fraud rates in the early-morning hours U.S. Eastern time.

First-Party Fraud Rate by Time of Day (U.S. Eastern)

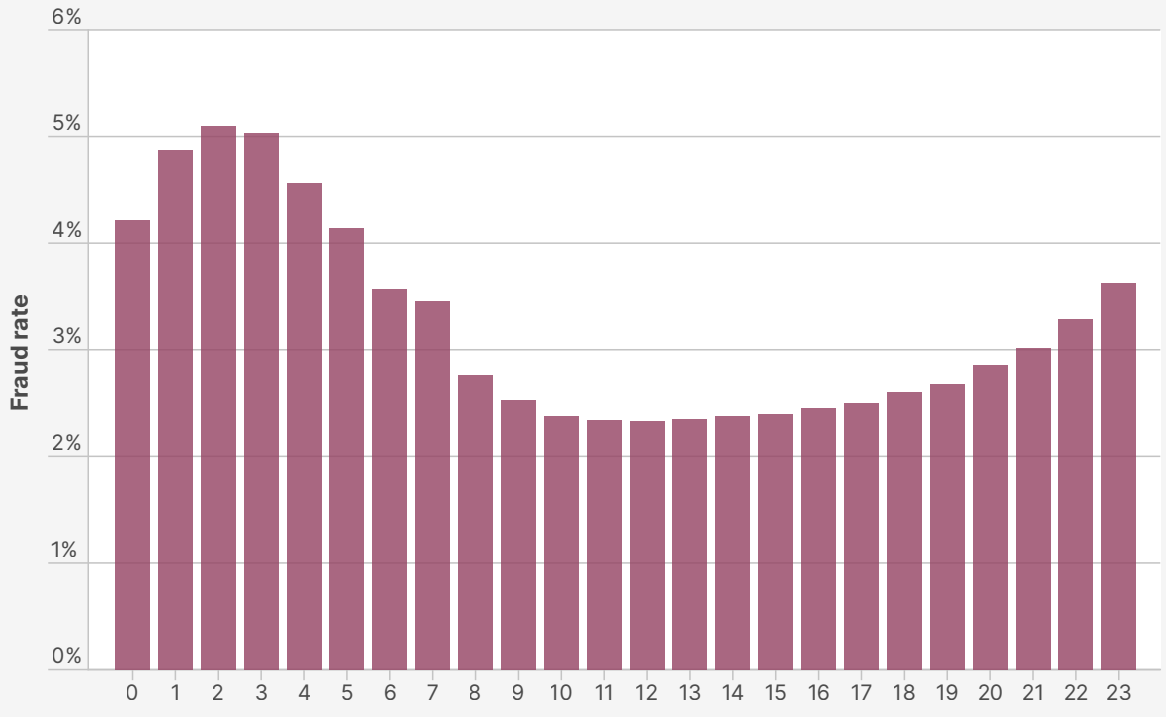

First-Party Fraud Rate by Day of Week

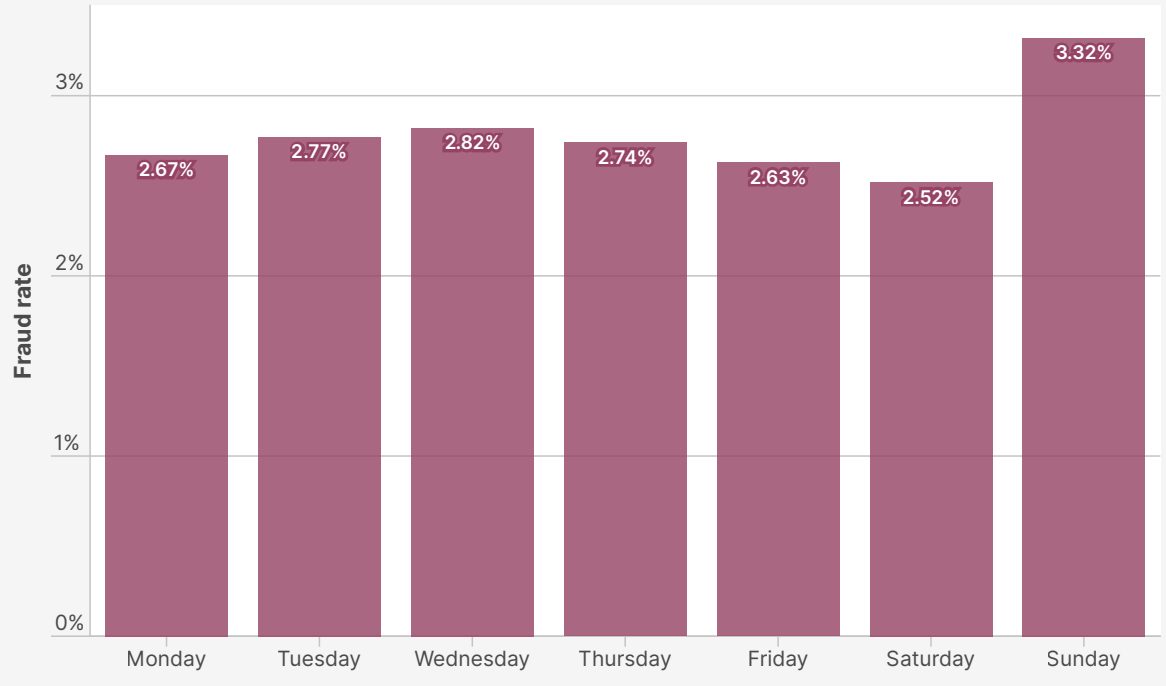

Fraudster age patterns in first-party fraud

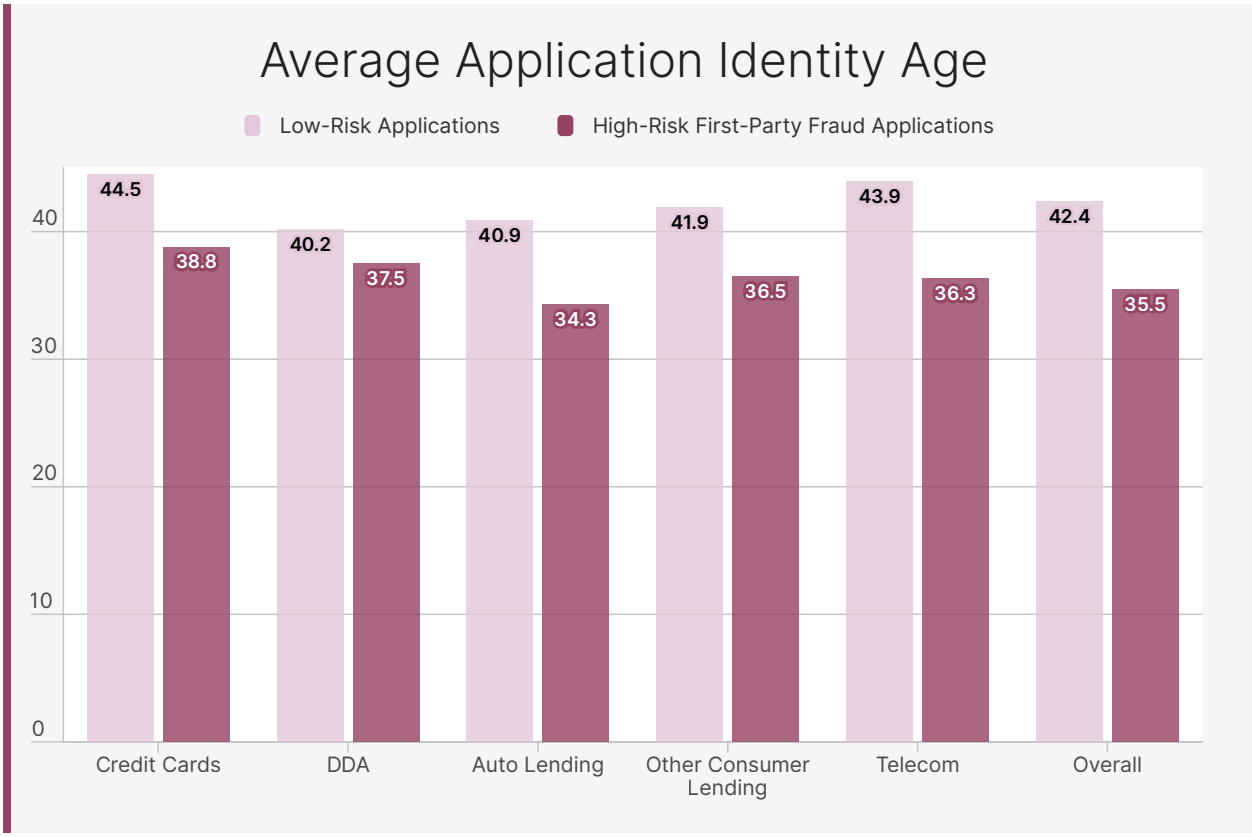

While people of all ages can and do commit first-party fraud, the identities on high-risk FPF applications are, on average, younger than the identities we typically see on applications that score as low-risk for first-party fraud.

Note: this chart is for context only. SentiLink models do not use demographic information to assess any kind of fraud risk, including first-party fraud risk.

Industry Breakdown

Credit Card Identity Fraud Rates

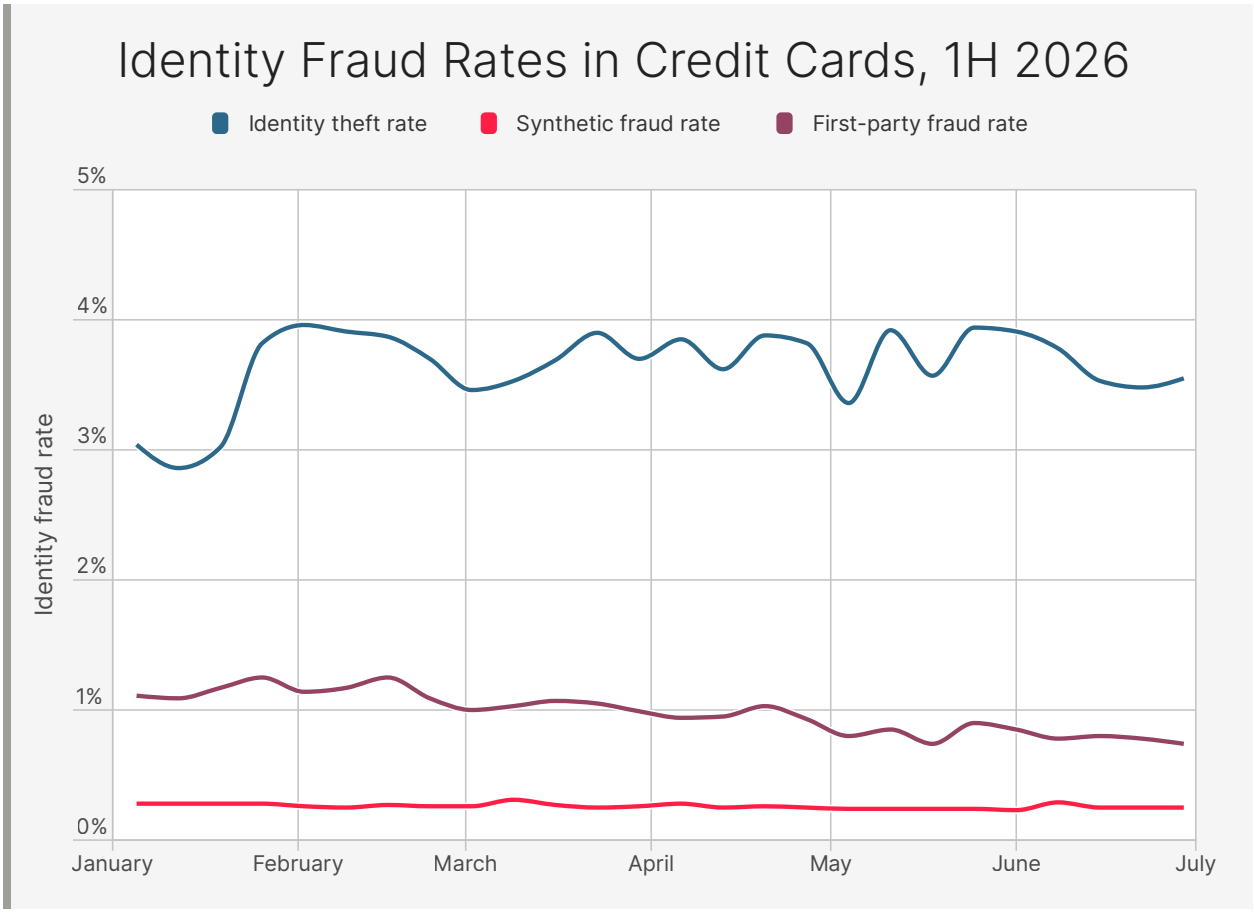

3.64%
Mean identity theft rate

0.26%
Mean synthetic fraud rate

0.98%
Mean first-party fraud rate

In Credit Cards, rates remained relatively steady throughout the half after rising from roughly 3% to nearly 4% in January. Part of the change was driven by targeted attacks. For example, one credit card partner suffered a late-January attack that briefly drove their identity theft rate above 25%:

SentiLink does not share partner data publicly. The above chart represents the real attack pattern experienced by the partner, but is not the partner's actual data.

This attack leveraged older-than-average identities: high-risk applications listed identities that were on average about ten years older than what we observe for low-risk application identities with this partner, and 2.5% of high-risk identities were over age 80, compared to just 0.3% of low-risk identities. Disproportionate percentages of these identities live in Florida (20.1% of high-risk applications) and Colorado (11.8% of high-risk applications), suggesting that the fraudsters behind the attack leveraged identities that may have been acquired via data leaks connected with institutions in these states.

High-risk applications were also much more likely to be associated with legacy email providers such as Hotmail and Yahoo, and more than 10% of high-risk applications came from an obscure email provider whose website appears risky.

Top Fraud Signals in Credit Cards, 1H 2026

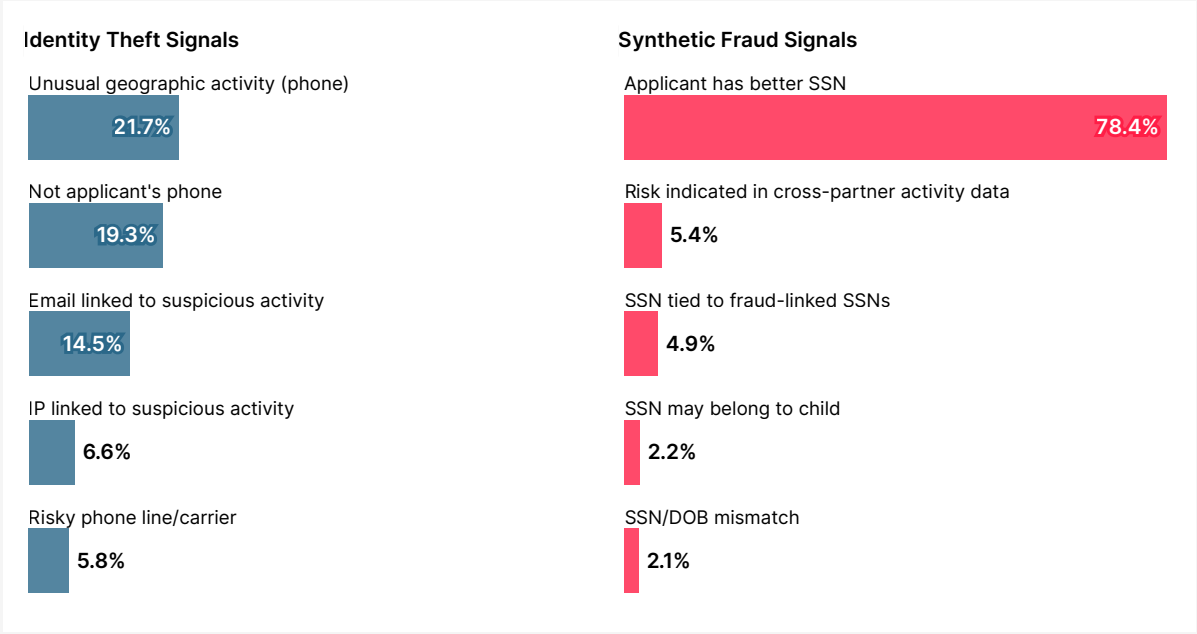

The top signals of identity theft in Credit Cards are similar to what we observe across all industries, with most signals relating to two elements of a stolen identity that fraudsters frequently replace with their own: the phone number and the email.

Synthetic fraud signals, by their nature, are centered around issues related to SSN, although risk indications in our cross-partner activity data were also a significant signal of synthetic risk. Of interest in Credit Cards specifically is the presence of “SSN may belong to child,” which did not make the list of top synthetic fraud signals in any other industry. This may indicate that one or more fraud groups are intentionally leveraging children’s SSNs to submit applications, as these applications are clustered at just a few partner institutions — three partners alone account for more than half of all credit card applications with this as their top synthetic fraud signal.

Demand Deposit Account (DDA) Identity Fraud Rates

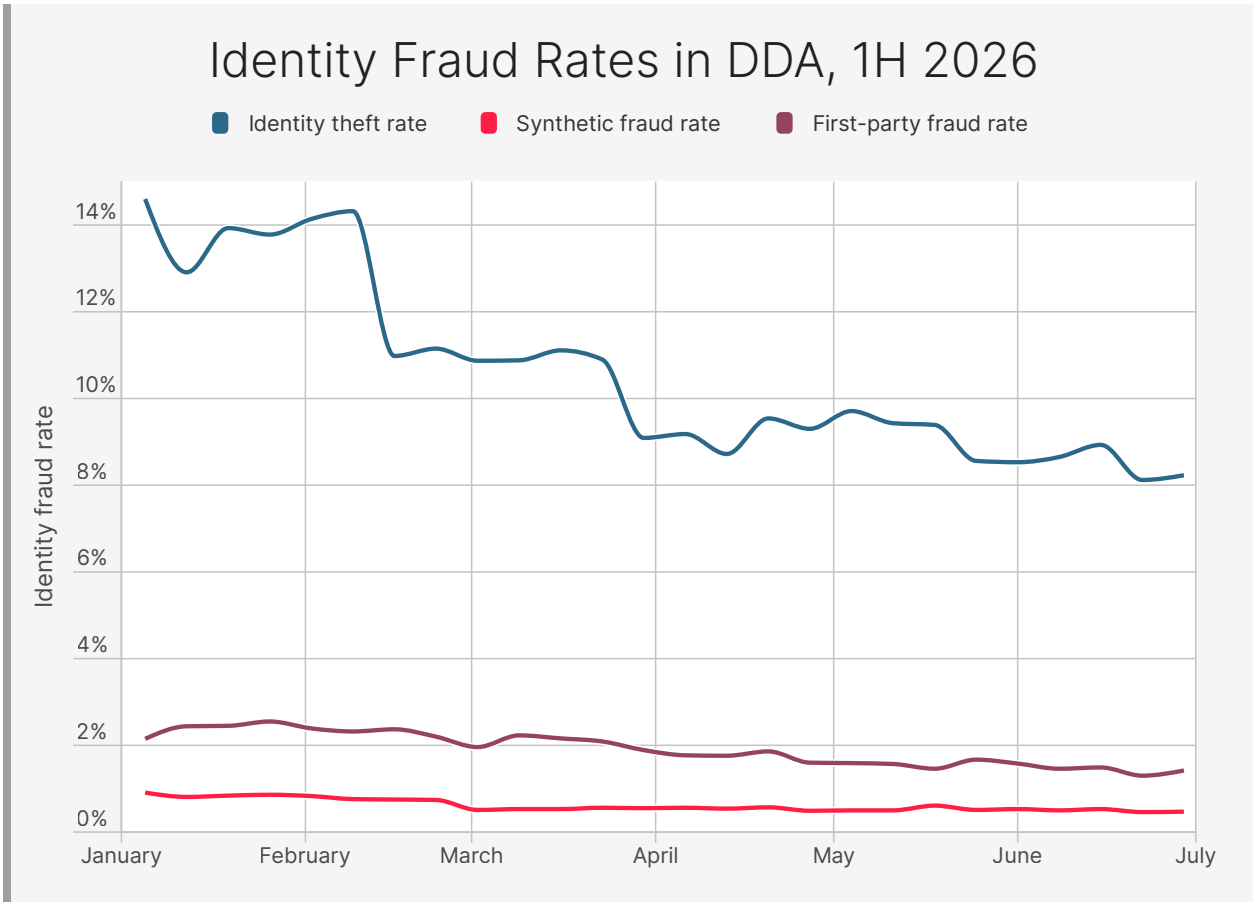

10.57%
Mean identity theft rate

0.61%
Mean synthetic fraud rate

1.91%
Mean first-party fraud rate

While identity theft rates remained high in DDA, they fell significantly from the holiday highs we observed at the end of 2025 and the beginning of 2026.

In our previous fraud report, we noted that the high identity theft rates during the holiday period appeared to reflect a seasonal pattern, but the rate was propped up through January and into February thanks primarily to several targeted attacks. One FI, for example, saw a late January attack that, on its worst day, spiked their identity theft rate to more than 60%. This attack leveraged identities that were, on average, 15 years older than low-risk identities, and

typically paired with emails from legacy providers Hotmail and Outlook. Many email handles were gibberish, a sign of automated account creation. The fraudsters also leveraged the same set of IP addresses for application submission repeatedly, with some IPs being reused for 20+ applications. By comparison, no IP address appeared more than three times on low-risk applications during the same period.

Another partner experienced a similar attack in early February: older-than-average victim identities, email handles that appeared unrelated to the victim's name, and significant signs of risk in the application cluster, including phones and IP addresses linked to other high-risk activity.

The synthetic fraud rate and first-party fraud rate were less volatile than the identity theft rate, but both also dropped. The synthetic fraud rate began the half around 85 bps and had dropped to around 50 bps by late June; the first-party fraud rate was over 2% early in the half but was closer to 1.5% as July approached.

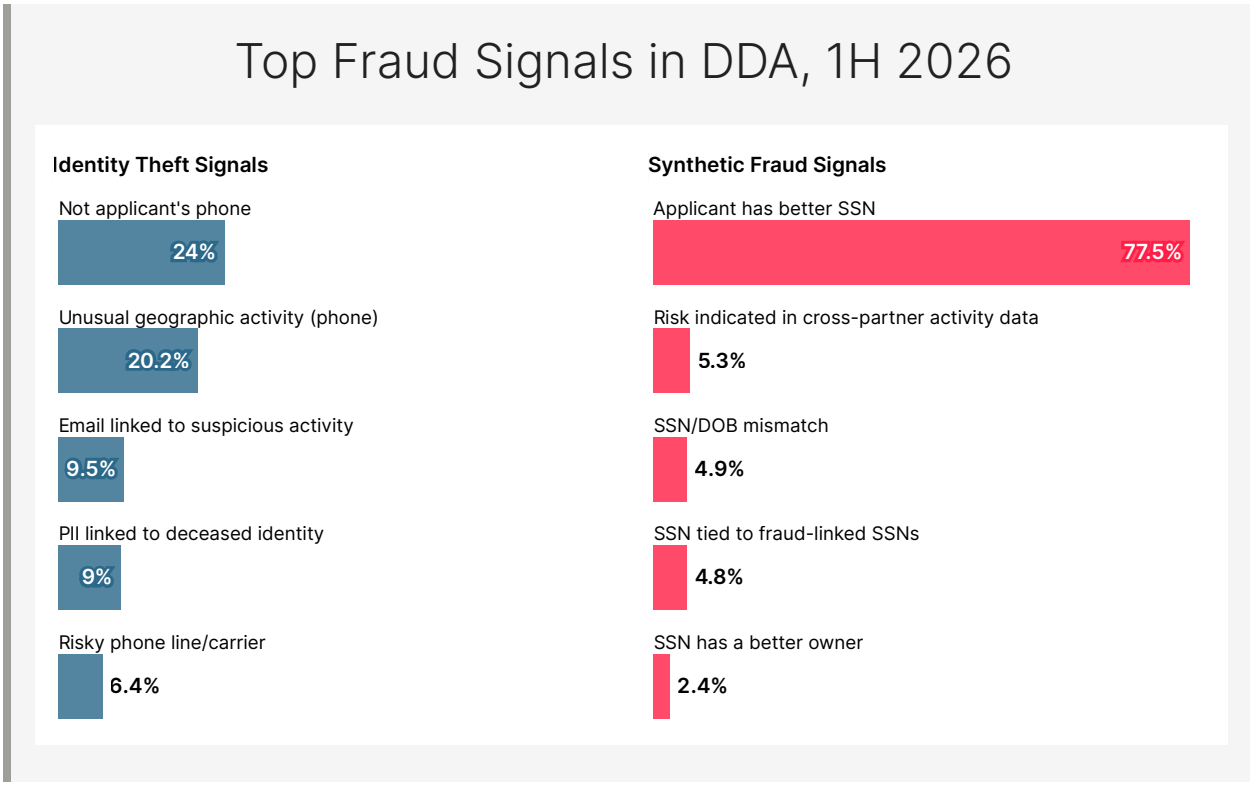

The top identity theft signals in DDA were similar to those observed across industries, with most involving two elements of a stolen identity that fraudsters commonly replace with their own: the phone number and email address. PII associated with deceased individuals was also a significant signal in DDA this half.

Synthetic fraud signals naturally tend to involve inconsistencies related to an applicant's SSN; four of the top five signals shown above fall into this category. Risky activity at other institutions in our cross-partner data was also a meaningful indicator of synthetic fraud.

Auto Lending Identity Fraud Rates

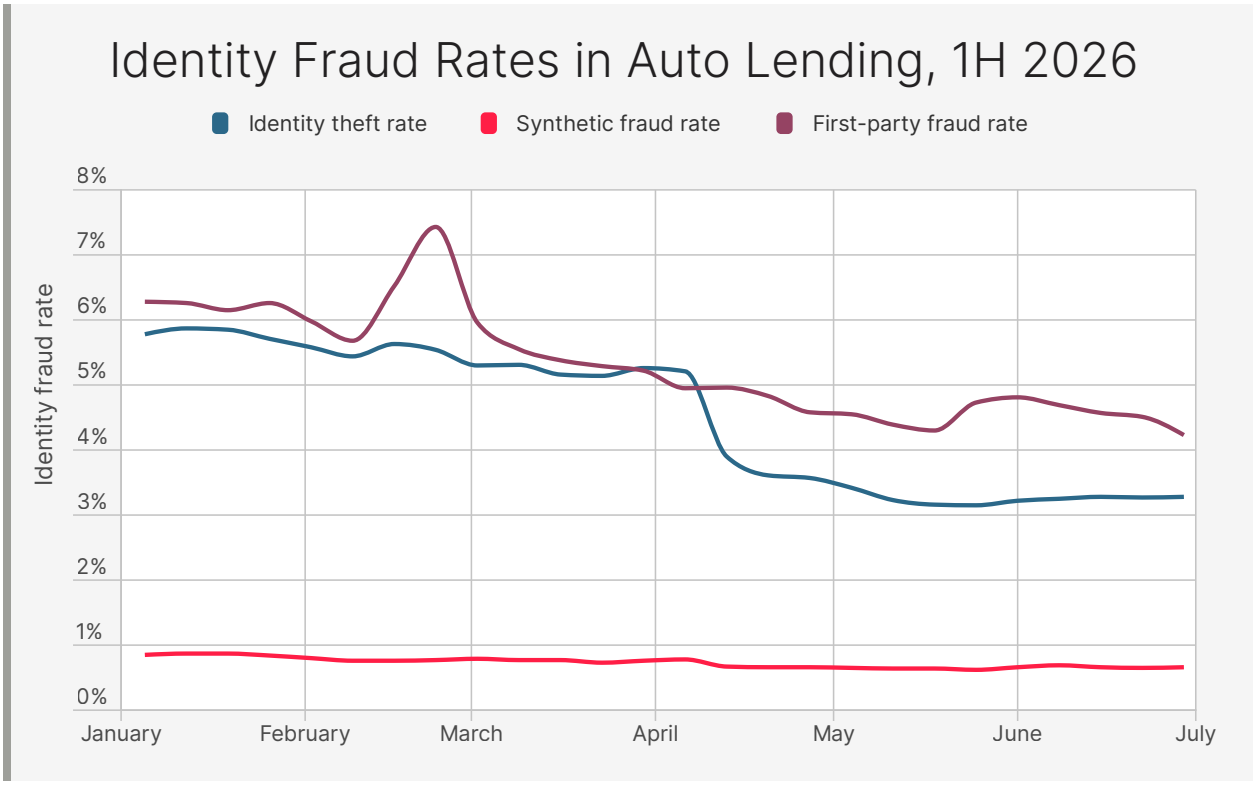

4.51%
Mean identity theft rate

0.73%
Mean synthetic fraud rate

5.31%
Mean first-party fraud rate

In Auto Lending, the biggest story was the first-party fraud rate, which — outside of a single week in early April — remained higher than the identity theft rate for the entire half.

Auto Lending is the only industry in which we see first-party fraud eclipsing identity theft. This result is consistent with what we've observed in recent retrostudies, and while we can't say with certainty *why* first-party fraud is higher in Auto Lending, there are a number of structural reasons that we suspect contribute, including:

- In-person interactions at dealerships and increased documentation requirements can make identity theft and third-party synthetic fraud more difficult in auto lending than in other industries.

- Dealers are incentivized to sell. Less scrupulous dealers may overlook red flags or encourage applicants to misrepresent figures such as income in order to help close the sale.

- The high value (and resale value) of a vehicle may entice more fraudsters to take the risk of attaching their own name to a loan they intend to default on. All fraud is risky, but the potential payoff from successful first-party auto fraud is likely many times greater than what a fraudster could expect to obtain through first-party credit card fraud, for example.

We also observed an interesting spike in the first-party fraud rate in late February. First-party fraud is typically committed by individuals, and this pattern does not appear to reflect a coordinated or organized attack. Instead, it likely results from a confluence of two factors: Presidents' Day sales and other discounts offered by automakers and dealers, combined with the arrival — or anticipated arrival — of tax refunds.

However, we did still see a pattern associated with this spike: the average age on high-risk first-party fraud applications was over ten years younger than what we saw on low-risk applications. High-risk applicants were, on average, less than 35 years old.

Outside of first-party fraud, we also observed a marked drop in the identity theft rate in April. This was primarily driven by several SentiLink partners who experienced rate drops in April. At one partner, for example, the rate fell from roughly 4% at the start of April to about 2% by the end of the month.

The synthetic fraud rate in Auto Lending remained comparatively stable, though, as in other industries, it declined slightly over the half, from 85 bps at the beginning to 66 bps at the end.

Top Fraud Signals in Auto Lending, 1H 2026

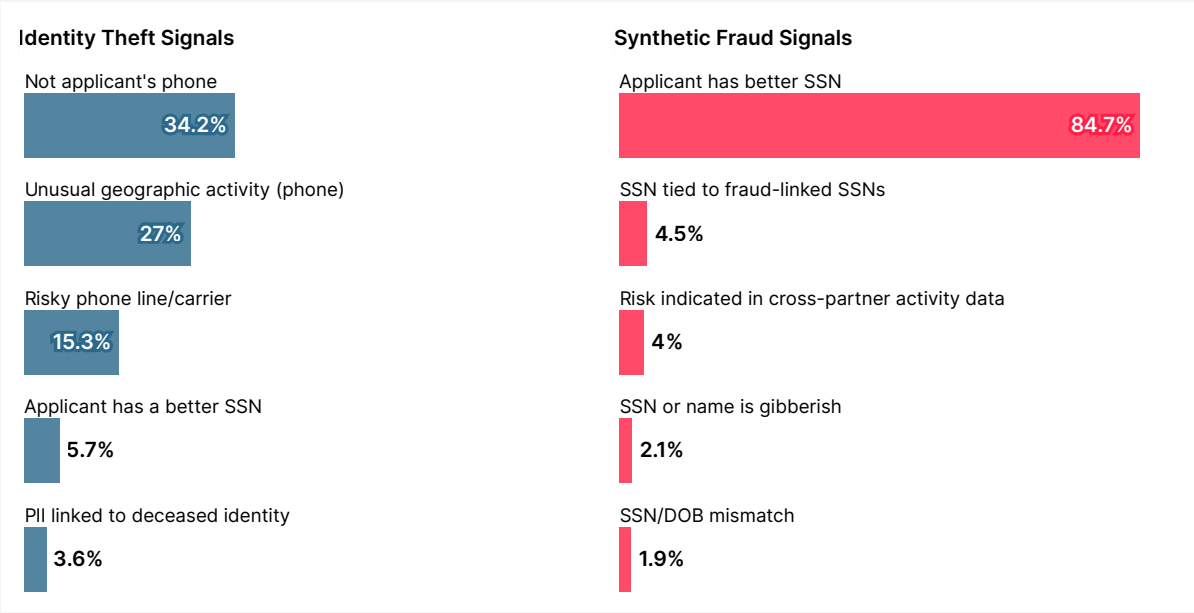

The top signals of identity theft in Auto Lending are similar to what we observe across all industries, with most signals relating to two elements of a stolen identity that fraudsters frequently replace with their own: the phone number and the email.

"Applicant has better SSN" appears as a risk signal for identity theft and synthetic fraud here because of attacks leveraging a fraud MO that blurs the lines. In this MO, fraudsters leverage the stolen identity of a real person, but combine it with an SSN that does not belong to that base identity. SentiLink calls this fraud subtype "IDT synthetic."

Other Consumer Lending Identity Fraud Rates

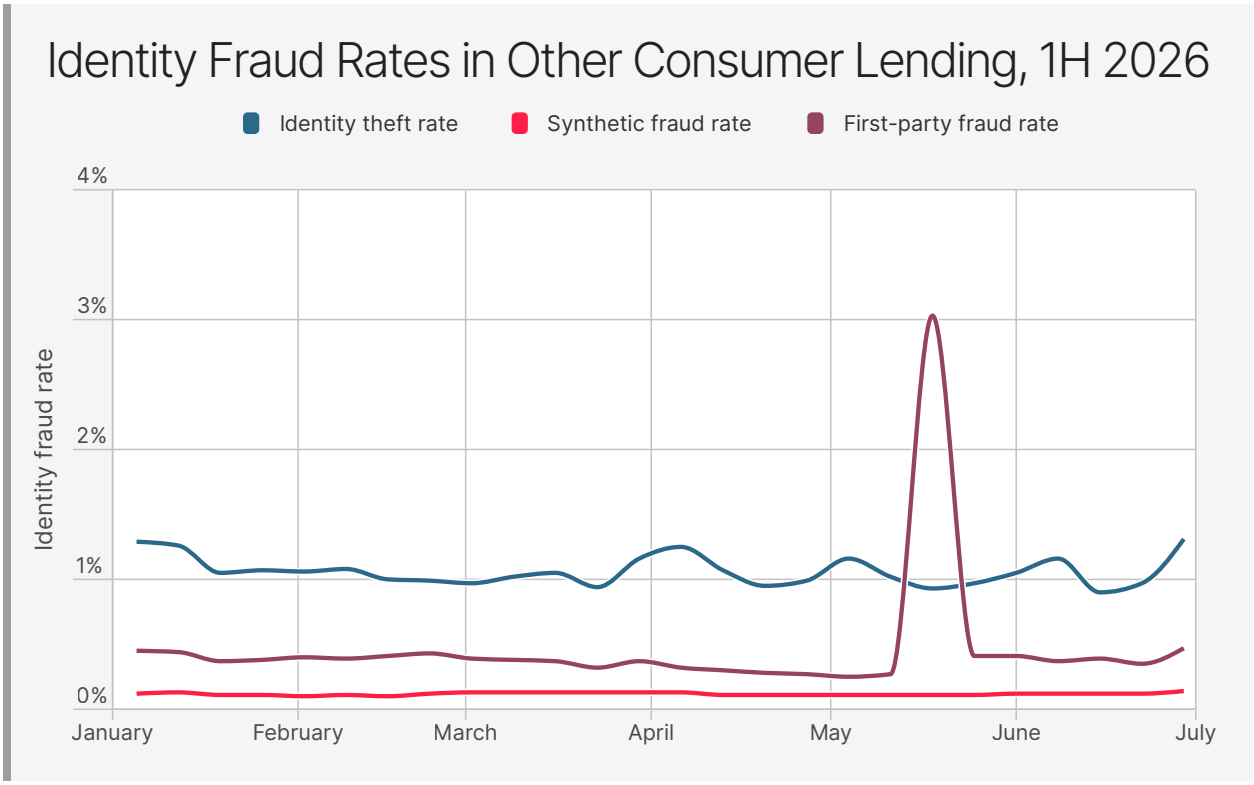

1.06%
Mean identity theft rate

0.12%
Mean synthetic fraud rate

0.47%
Mean first-party fraud rate

Fraud rates across the board were lower in Other Consumer Lending — our catch-all category for a variety of consumer lending models, including AltFi and Buy Now, Pay Later lenders — than in any other category, consistent with the pattern we have observed in previous years. One notable exception was a brief spike in first-party fraud in May. This spike was not universal, but it affected several high-volume partners. The timing of the spike coincided with Memorial Day, when many businesses offer promotions and discounts.

Similar to the first-party fraud spikes we've investigated in other industries and our observations about age on high-risk applications overall, the average identity associated with high-risk applications during this spike was meaningfully younger (37) than the average identity associated with low-risk applications (44).

For the first time in this report, we have also broken out rates for two subtypes of lenders in the Other Consumer Lending category: Alternative Finance (AltFi) lenders and Buy Now, Pay Later (BNPL) lenders.

Alternative Finance (AltFi) identity fraud rates

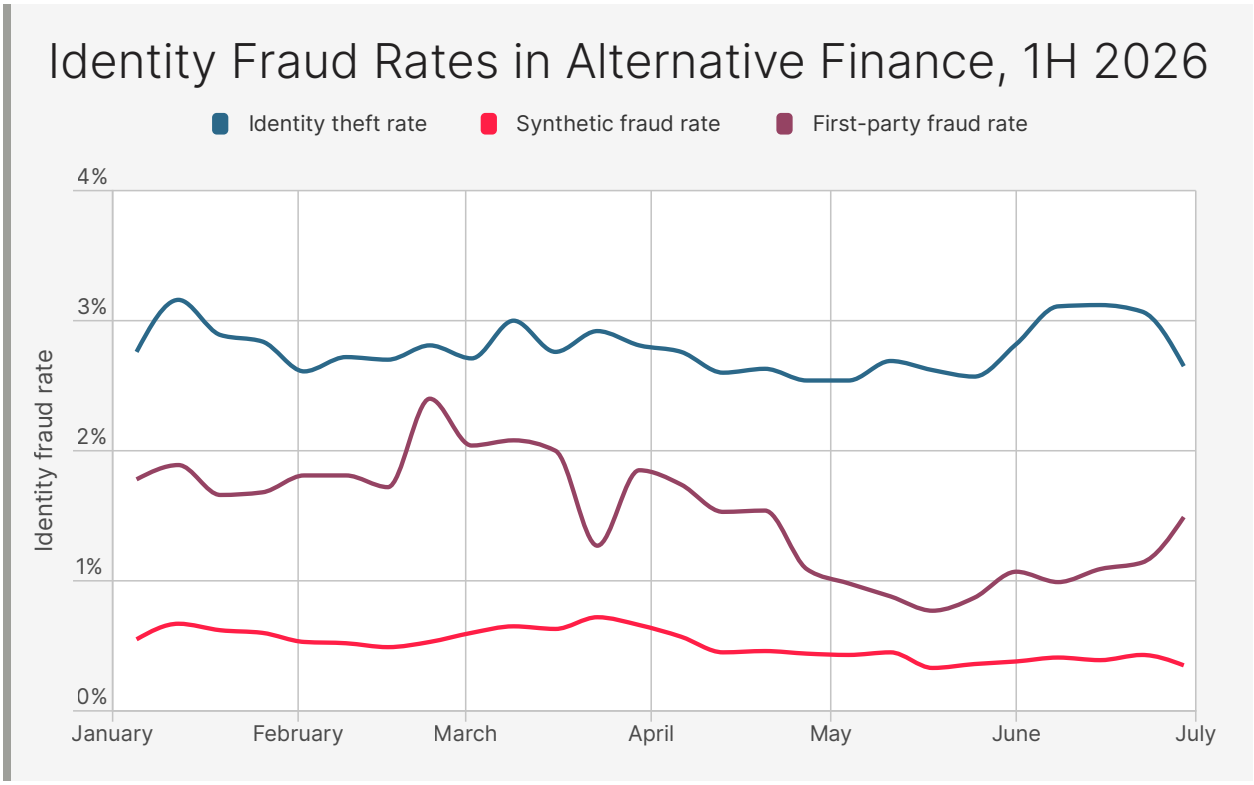

2.78%
Mean identity theft rate

0.51%
Mean synthetic fraud rate

1.51%
Mean first-party fraud rate

Unsurprisingly, fraud rates within AltFi were higher than rates in BNPL and the overall Other Consumer Lending category.

In identity theft, the sustained "bump" in June is thanks primarily to a targeted attack affecting several SentiLink partners intermittently over the course of two weeks. Interestingly, the attackers in this case submitted a high number of "Frankenstein" applications, in which PII from multiple real base identities is mixed.

In some cases, the identities used in this attack also belonged to deceased individuals.

In first-party fraud, the rate drop was driven by a combination of factors: a major shift in the application mix of partners in this category caused by a significant application volume increase from a partner with relatively low rates, in tandem with genuine rate declines observed at quite a few other partners.

Buy Now, Pay Later (BNPL) identity fraud rates

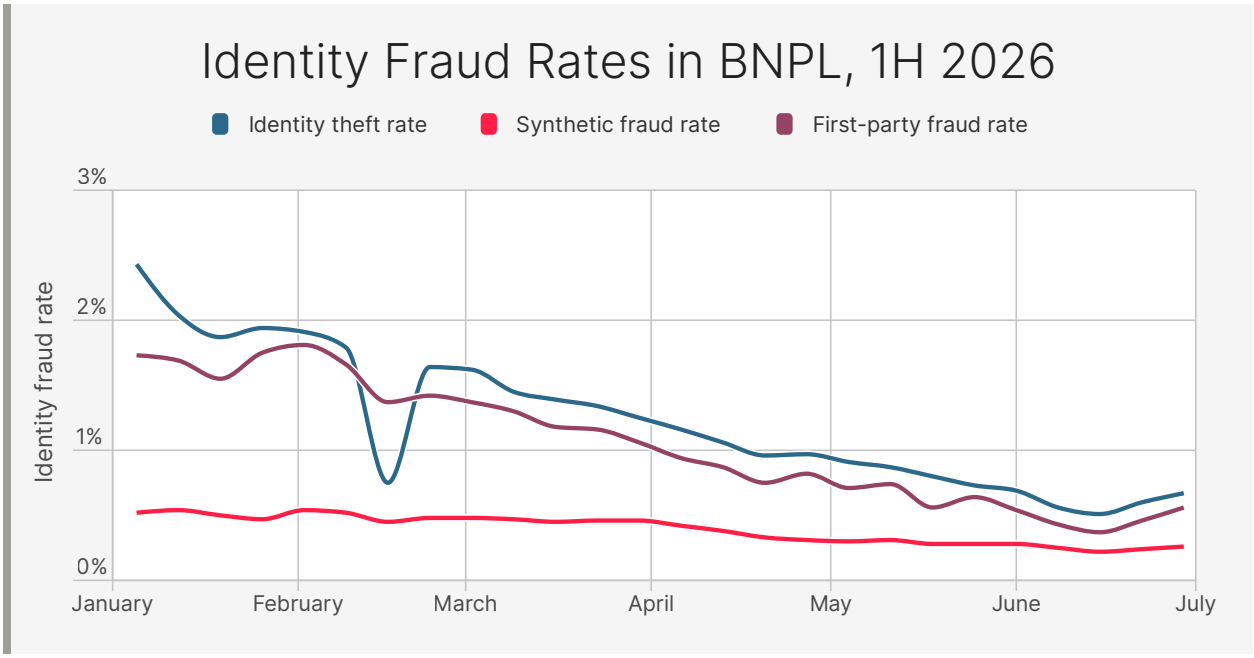

1.23%
Mean identity theft rate

0.39%
Mean synthetic fraud rate

1.06%
Mean first-party fraud rate

Rates in BNPL were low compared to other industries, and all three declined meaningfully over the course of the half. The dip in identity theft rates in February was driven primarily by a massive BNPL provider that saw a major influx of applications from low-risk customers with an average age slightly older than their normal clientele. This dip also correlated with Presidents' Day, and thus may have been related to sales or promotions at retailers that leverage this BNPL lender as a payment option.

Top Fraud Signals in Other Consumer Lending, 1H 2026

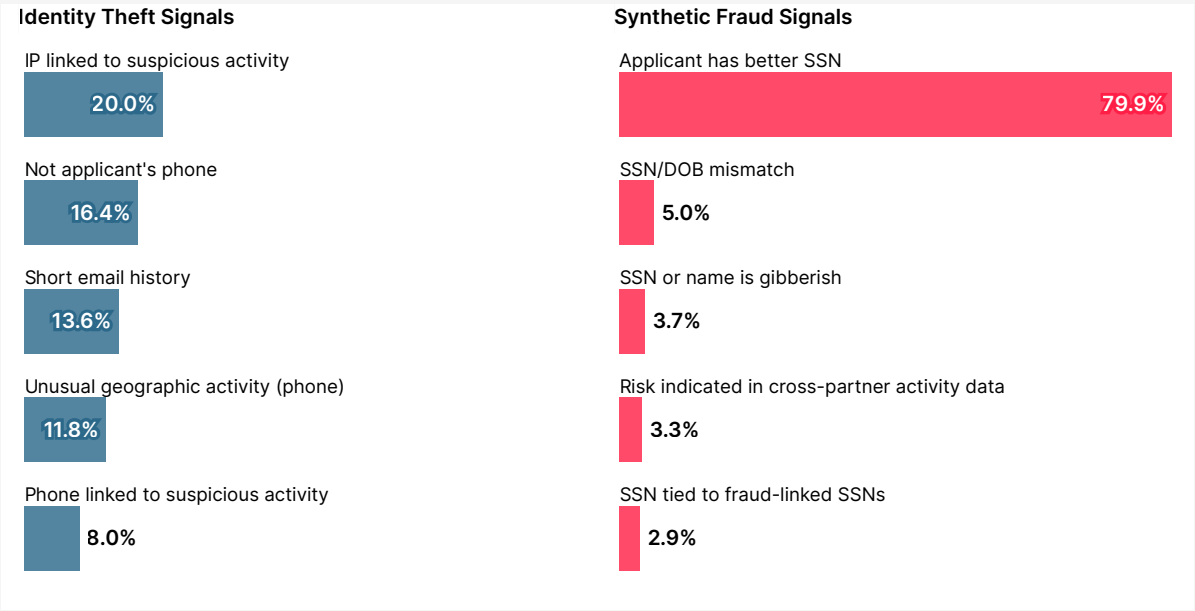

The top signals of identity theft in Other Consumer Lending are similar to what we observe across all industries, with most signals relating to two elements of a stolen identity that fraudsters frequently replace with their own: the phone number and the email.

Synthetic fraud signals by their nature are centered around issues related to SSN, although risk indications in our cross-partner activity data was also a significant signal of synthetic risk.

Telecom Identity Fraud Rates

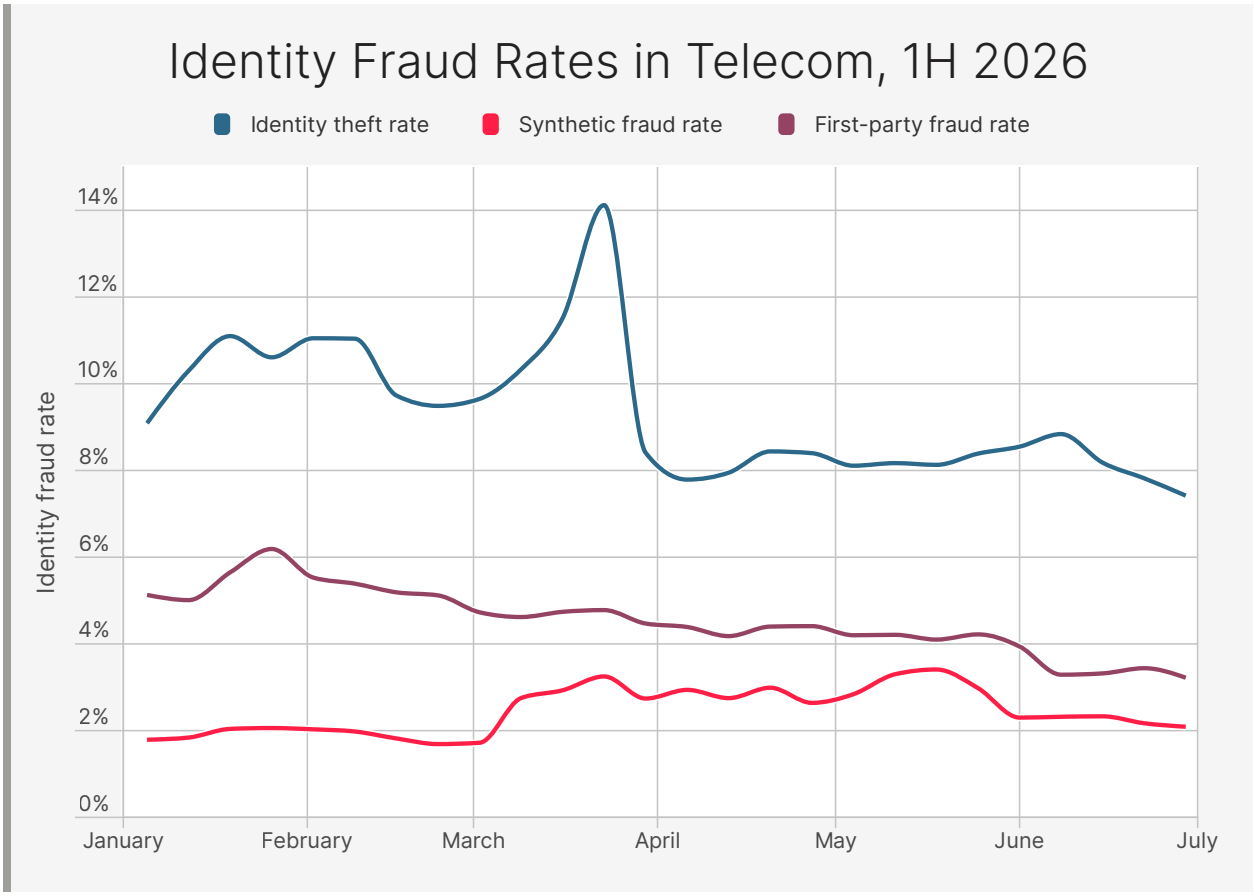

9.33%
Mean identity theft rate

2.45%
Mean synthetic fraud rate

4.53%
Mean first-party fraud rate

As we have seen in previous editions of this report, the Telecom industry produced fraud rates on par with the highest rates in other industries, and its synthetic fraud rate is again the highest of any industry we look at. Consistent with previous reports, these high rates are driven primarily by the subset of telecom companies that are mobile providers who offer smartphone and other device financing; the broadband providers in our cohort have lower rates.

Most notable this half is the massive identity theft spike in March, driven by a targeted attack at several SentiLink partners. As we've seen in many other identity theft attacks, this attack leveraged older-than-average identities (the average identity on high-risk

applications during this spike was over 50) that appear to have been sourced from data breaches connected to specific localities — more than 40% of the stolen identities listed on high-risk applications were from Maryland, and an additional 18.9% were from Pennsylvania.

Almost all of these high-risk applications (88%) listed Hotmail email addresses, and the fraudsters behind the attack appear to have been using the same set of IP addresses repeatedly; some IP addresses appear on hundreds of different high-risk apps over the course of just a few days.

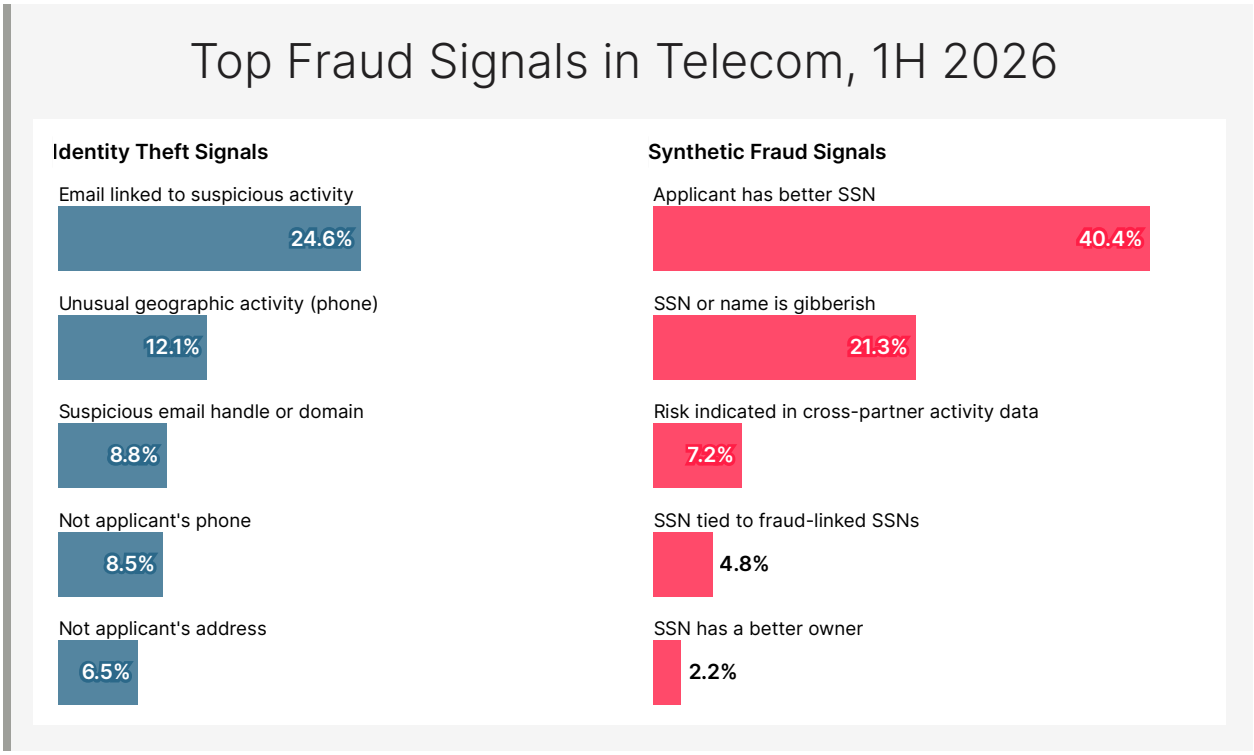

The top signals of identity theft in Telecom are similar to what we observe across all industries, with most signals relating to two elements of a stolen identity that fraudsters frequently replace with their own: the phone number and the email.

Synthetic fraud signals by their nature are centered around issues related to SSN, but it is notable that “SSN or name is gibberish” is a more common signal here than in other industries. Telecom companies appear to be one of the first types of institutions to which fraudsters apply with stolen identities; it may be that synthetic fraudsters are also using telecom applications as a testing ground for new synthetic identities and SSNs.

What Fraud Really Looks Like

The Costs of Fraud

The vast majority of fraudulent applications highlighted in this report are likely rejected by our partners because SentiLink identifies them as high risk. But the scale of this incoming fraud underscores the real stakes that financial institutions (FIs) and other lenders face: missing fraud comes at a significant cost.

Based on real-world performance data from more than 1.6 million applications submitted between 2022 and 2026 for which SentiLink also retroactively generated fraud risk scores, we were able to calculate the average charge-off amounts associated with high-risk applications by risk type:

$6,149

Overall average charge-off amount for **all high-risk fraud** applications.

$5,223

Average charge-off amount for high-risk **identity theft** applications.

$9,277

Average charge-off amount for high-risk **synthetic fraud** applications.

$8,746

Average charge-off amount for high-risk **first-party fraud** applications.

We also examined industry-specific charge-off amounts across several fraud MOs. Although the analysis was limited by data availability, we were able to compare charge-offs associated with high-risk applications against average charge-off amounts in two industries: Credit Cards and Other Consumer Lending.

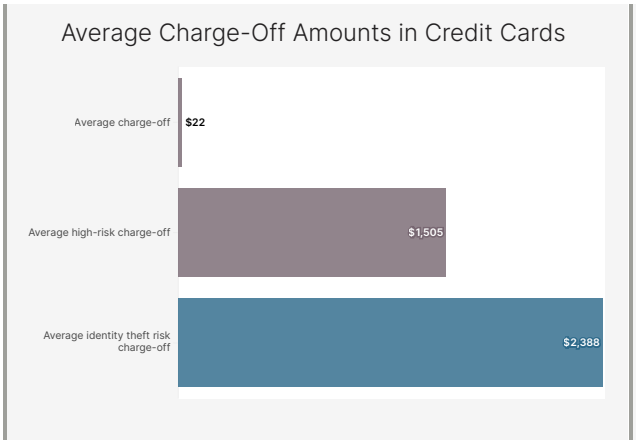

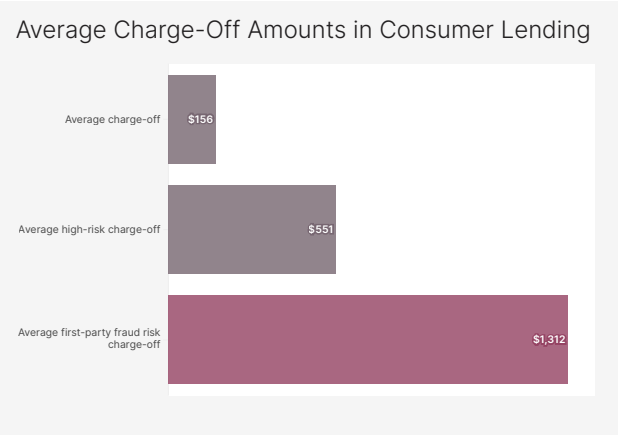

In both Credit Cards and Other Consumer Lending, charge-off amounts associated with high-risk fraud applications were significantly higher than the overall industry averages. Fraud-related charge-off amounts were **3.5 times higher** in Consumer Lending and **68 times higher** in Credit Cards.

The differences were even more stark for specific fraud MOs. In Credit Cards, identity theft was the costliest: average charge-off amounts were **108 times higher** than the average across all credit card charge-offs, including those unrelated to fraud. In Consumer Lending, first-party fraud stood out, with average charge-off amounts **8.4 times higher** than the industry-wide average.

How Institutions Experience Fraud

Note: SentiLink does not share partner data publicly. While the charts in this section are based on real partner data and reflect real trends, we have adjusted all values using a randomly changing variable to ensure that while the charts reflect real phenomena, they do not contain real partner data.

While this report focuses on average rates and trends overall and by industry, it is important to remember that rates can vary dramatically among institutions within the same industry.

For example, the chart below depicts the identity theft rates of three partners within the same industry:

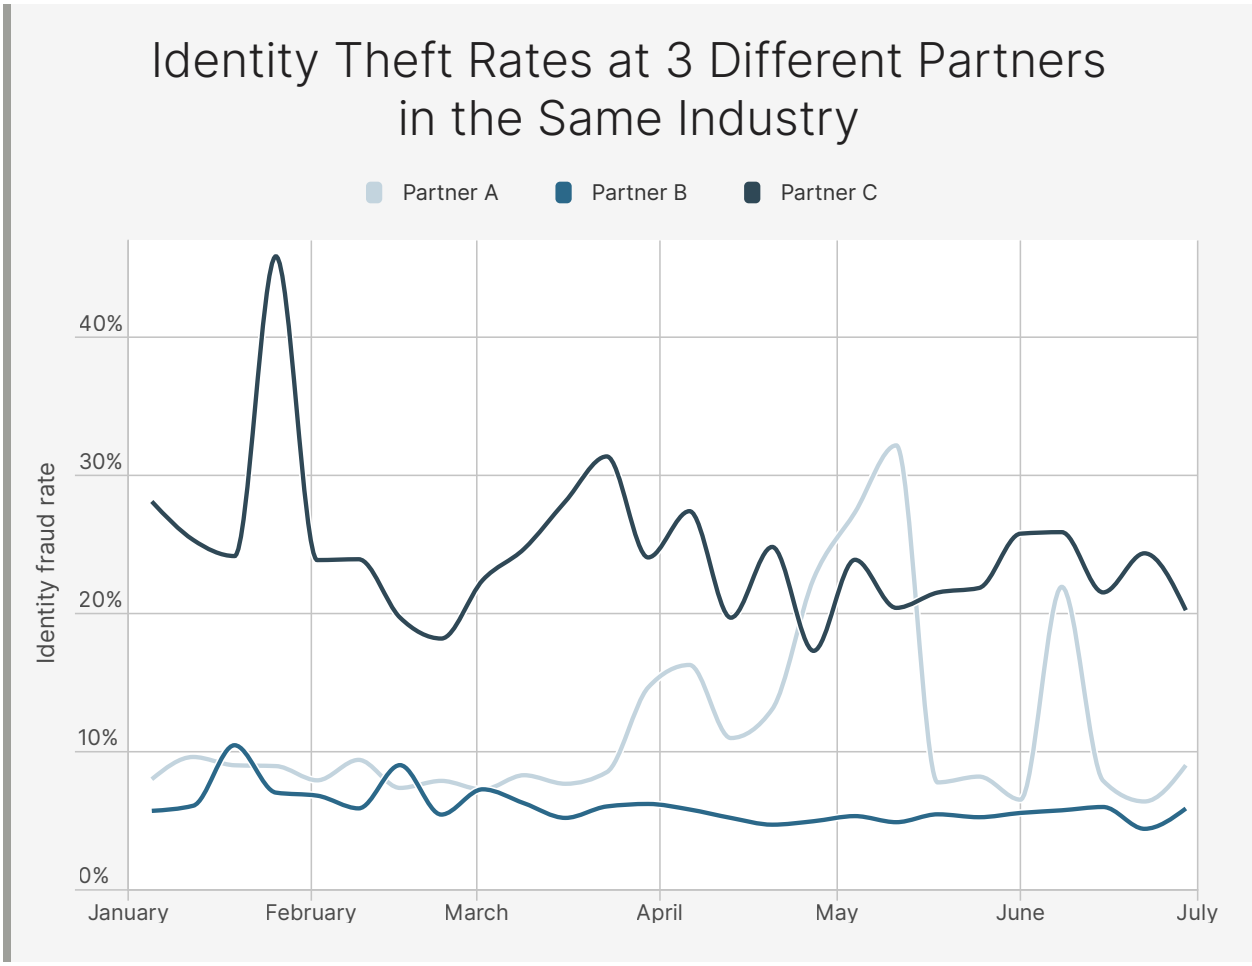

Additionally, our aggregation of the data by week makes charts more readable, but it also hides the fact that fraud rates often vary quite a bit by day. Below is an example of what a daily identity theft rate for a partner from a different industry looks like:

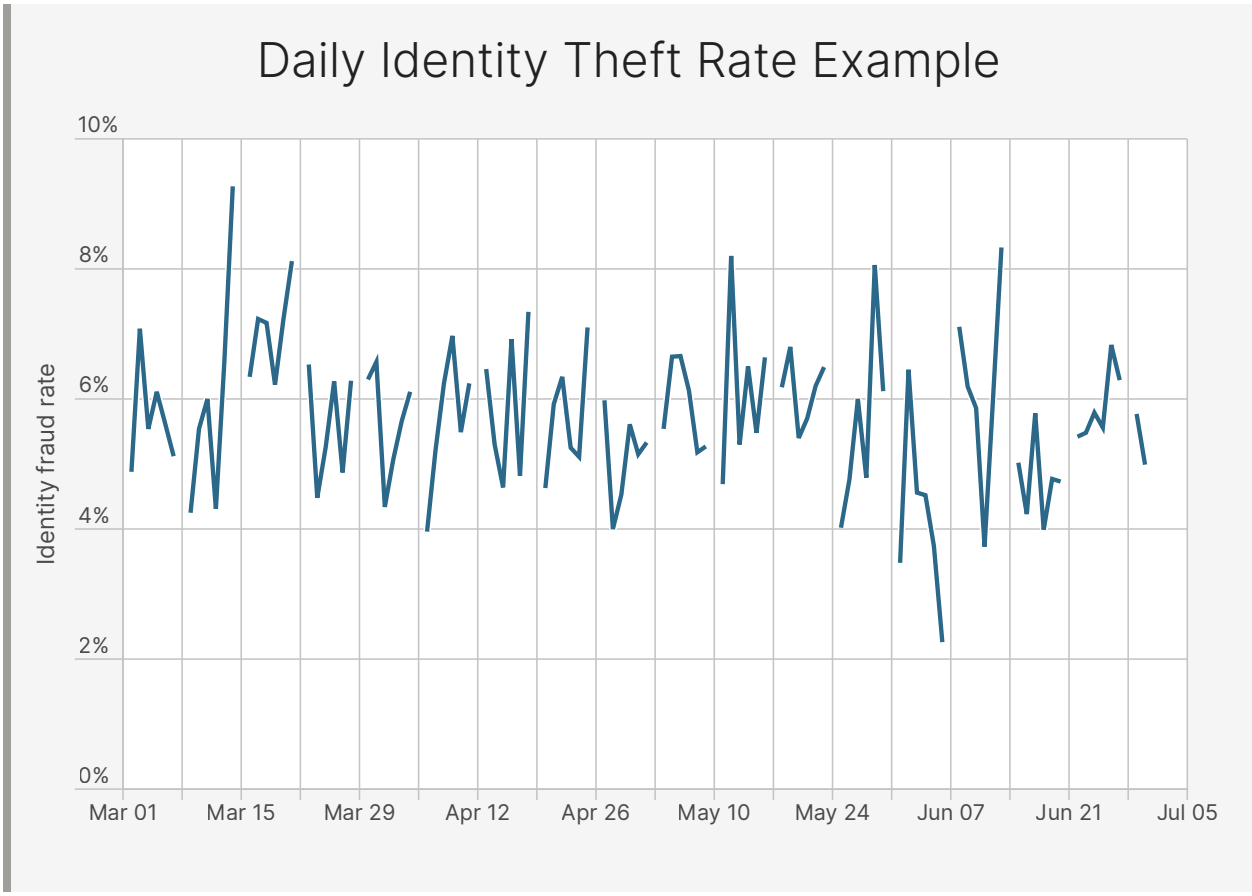

Fraud Trends and Insights

Residential Proxies Are Outpacing Detection

In two large identity theft attacks on SentiLink partners in the first half of 2026, fraudulent applications appeared to come from the victims' own neighborhoods over ordinary home internet connections. The attackers were using residential proxies, which route internet traffic through real consumer devices and make a fraudster appear to be connecting from a victim's city or even through the same internet provider.

At first, the attacks were confusing. The applications showed clear signs of coordination: stolen, long-established personal information, repeated activity across institutions, and a different IP address for nearly every attempt. But the network data looked unusually legitimate. Instead of originating far from the victims or passing through familiar VPN infrastructure, many applications appeared to come from ordinary residential connections near the victims' homes.

For years, institutions largely kept pace with this threat through commercial detection services. These companies map the IP addresses used by VPNs and proxy networks, allowing an institution to flag a connection even when it otherwise looks local and residential.

That defense now appears to be falling behind. Residential proxy networks are adding and rotating IP addresses faster than detection services can discover and classify them. SentiLink researchers saw the gap firsthand: we purchased residential proxies and checked the IP addresses we received against the commercial proxy-intelligence data used in SentiLink's screening. The addresses were not identified as proxies by an industry-standard detection service.

Several forces are accelerating the market. Proxy services have become cheaper and easier to use, while location targeting that was once a premium feature is now widely available through self-service platforms. A buyer can request an IP address in a particular city, ZIP code, or internet provider's network and rotate to a fresh one for each application.

Demand has also grown. Data companies, AI crawlers, and automated web agents increasingly use residential proxies to collect information from websites without being blocked. That legitimate market has helped finance larger device pools and better targeting tools. Fraudsters can rent the same infrastructure.

The result is an increasingly difficult race. A connection can appear local and clean, not because it is legitimate, but because the detection system has not encountered it yet. This is particularly

dangerous when fraudsters capture legacy phone numbers or emails once associated with the victim but no longer in active use (a trend we call "legacy PII," which we covered in more depth in the previous edition of this report).

Institutions, therefore, need to rely more heavily on behavioral signals, such as the same identity applying across several institutions in quick succession, or changing IP addresses with every attempt.

> Note: SentiLink's next Identity Theft Score model, available from September, contains significant improvements that allow it to flag fraudulent applications using this MO.

Mechanic's Lien Fraud

Fraud rings are exploiting mechanic's lien sales — the legal process that lets repair shops and tow companies auction a vehicle to recover an unpaid bill — to wash the titles of financed vehicles and cash them out, leaving auto lenders with a written-off loan and no collateral. SentiLink research identified hundreds of lien-sale advertisements in California community newspapers since 2024 bearing hallmarks of the scheme.

The lien sale process is vulnerable because it runs on an honor system. In California, a shop claiming an unpaid repair or storage bill notifies the vehicle's legal owner — usually the lender — who has 10 days to object before the shop may sell the car at public auction. No one verifies the shop's claim, and nearly every step can be gamed: bills padded with storage fees or repairs never performed, auctions held in remote locations or won by a single friendly bidder, notices sent to a lender's general mailbox days before the deadline — or an empty certified-mail envelope, with the delivery receipt shown to the DMV as proof of notice. Shops that value a vehicle at $4,000 or less can bypass DMV approval entirely, making such a valuation on a late-model luxury car a strong indicator of fraud.

The pattern is accelerating. Attorneys representing major auto lenders report that these cases, rare before 2020, now arrive several times per week. One Southern California body shop advertised 18 vehicles for lien sale in a recent 13-month span, including three Range Rovers and a BMW M8. Per public court filings, at least two finance companies have sued the shop, including over a nearly new leased truck the shop valued at "$4,000 or less" while billing $6,090 in storage, a vehicle the lender valued at $46,480. The filings do not establish fraud, but they fit a broader pattern of contested lien-sale activity. Criminal enforcement remains rare, as lenders' complaints are typically treated as contract disputes rather than crimes.

For lenders, the most effective defense is speed: suing and placing a DMV title hold within the objection window, before the title can be washed. The longer-term fix is regulatory — the notice rules were written for an era of local lenders and local newspapers, and haven't caught up.

Scammers Are Targeting Consumer Home Equity

For four months, SentiLink research led by Head of Fraud Insights Dr. David Maimon monitored a Telegram channel used by the Yahoo Boys, the West African networks behind much of today's romance scam activity. In it, members share instructions and materials for taking out home equity lines of credit (HELOCs) in the name of homeowners they recruit.

The posts describe the target they want: a homeowner with a credit score above 750 and a loan between $50,000 and $1,000,000, referred to as a "client." One message tells members to ask the client which bank holds their mortgage, lists roughly two dozen banks to try, and directs them to message a handler to set up the application. Other posts advertise ready-made scripts for these interactions — known as "formats," or "FMTs" — offer verified identities and bank login credentials for sale, and link to video tutorials with titles such as "How Yahoo Boys Cash Out From HELOC." The scheme runs two ways: apply in a homeowner's name using stolen personal information, or spend weeks cultivating the homeowner's trust under false pretenses before coaching them to apply and wire over the proceeds.

Separately, HELOC application data shows a pattern consistent with rising misuse. We examined roughly 87,000 HELOC applications submitted to SentiLink partners during 2025 by applicants aged 55 or older, then filtered for identities tied to three or more applications to different partners within the seven days before the HELOC attempt — a velocity consistent with coordinated identity misuse rather than legitimate borrowing. That filter surfaced 761 applications, and the monthly trend is pronounced: 28 in January, climbing unevenly through the summer, then 140 in October, 113 in November, and a year-high 199 in December — a 161% rise from September to December. In some cases, a single identity applied to several HELOC providers at once, burst activity consistent with an operation testing multiple lenders against the clock.

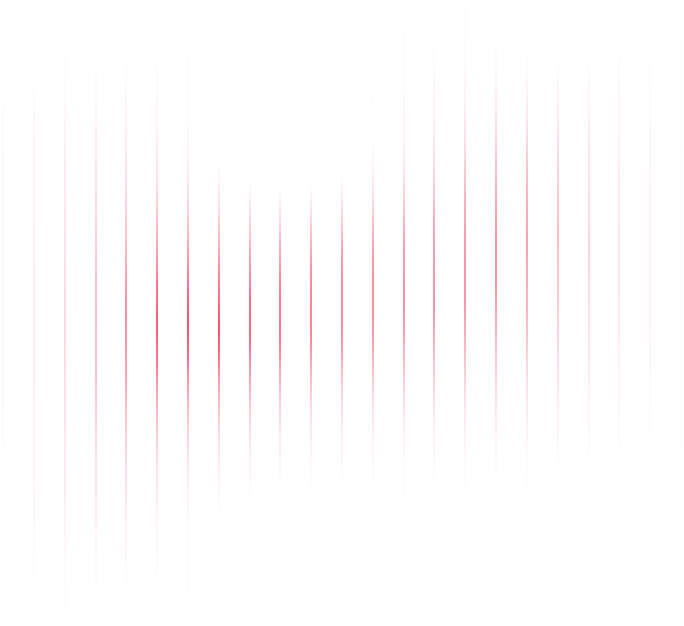

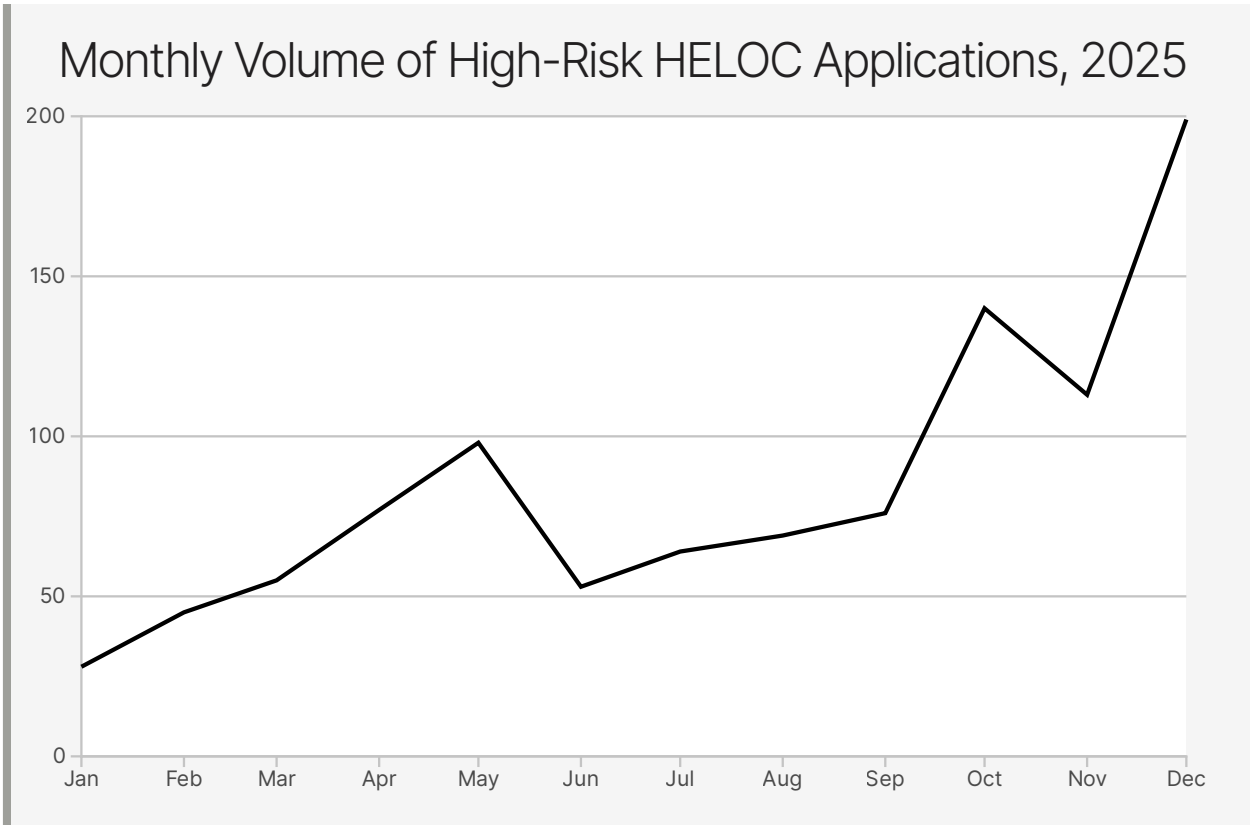

Network and device intelligence on these applications found heavy use of high-risk VPNs. But a meaningful subset came from residential or local IP addresses that the screening did not flag. That is consistent with several explanations: residential proxies the detection tooling missed — the commoditized, hard-to-detect supply discussed in the Residential Proxy section — a compromised device, or the homeowner completing the application in person under a scammer's coaching. Identity verification signals split the same way, some consistent with deepfake impersonation, others with the genuine homeowner interacting directly with the lender under manipulation.

We do not tie these applications to the actors organizing in the Telegram channels; the two are independent signals that happened to surface over the same months. For lenders, the signals that separate this activity are behavioral rather than identity-based — multi-partner velocity, simultaneous applications, VPN or residential proxy use — because in the victim-assisted variant, standard identity checks pass. Additional details appear in our HELOC blog post.

Yahoo Boys Are Coming for Retirement Accounts

In a separate investigation into the Yahoo Boys networks, SentiLink's Fraud Intelligence Team found vendors in a Telegram channel selling "formats" written specifically to talk victims into liquidating their retirement accounts. Where the home-equity formats went after a homeowner's largest asset, these target the other large pool of money many older victims hold: their 401(k)s.

One FMT script reviewed by SentiLink opens as relationship talk and pivots to an "agent" described as a "former Trump administration official" with "a knack for tripling" retirement accounts through exclusive access to alternative investments. From there, it anticipates every objection in the fluent voice of a financial advisor, walking the victim through risk tolerance, diversification, age, time-to-retirement, and ongoing monitoring, before arriving at the ask: the agent "will need access to your 401k account" and "might initiate some withdrawals."

The pitch is engineered for this moment. The median household nearing retirement holds about $185,000 in retirement savings per the Federal Reserve's most recent survey, and most Americans believe they haven't saved enough. An offer to triple the account speaks directly to that anxiety. The claimed political connection lands in a news environment that abounds with headlines of insider dealing and well-timed trades in Washington.

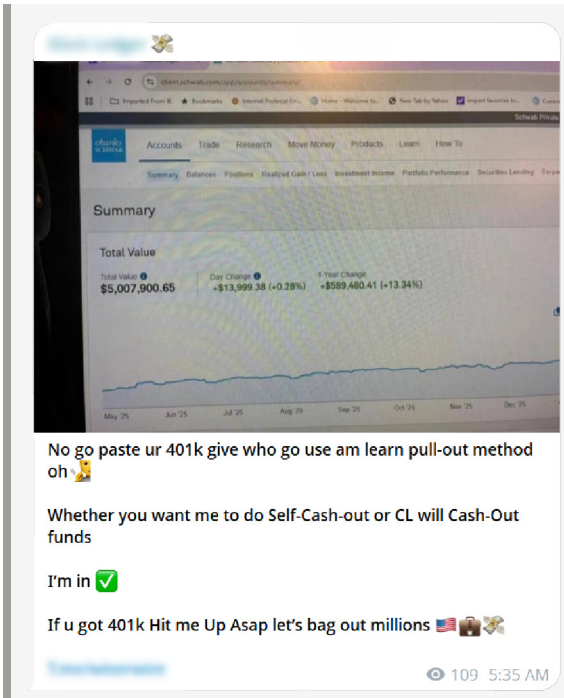

The Telegram channel lays out best practices to pull off the scheme; the execution and laundering pipeline is the established one. Buyers are advised to control victim accounts through remote desktop sessions on compromised computers in the victim's own region, so logins appear local — the same geographic camouflage discussed in the Residential Proxy section of this report. Proceeds move through drop accounts opened with stolen SSNs, runner networks that take 10–20% commissions, and conversion to cryptocurrency, including privacy coins.

Non-Domiciled CDL Holders Are a Prime Fraud Target — and Now Some Are Busting Out Themselves

People who hold non-domiciled commercial driver's licenses (CDLs) have long been a favored target for fraudsters, and in the first half of 2026, our Fraud Intelligence Team found the credential drawing a second kind of exploitation — this time from the holders themselves. A non-domiciled CDL goes to a foreign national who is legally present in the U.S. but not a permanent resident; the Federal Motor Carrier Safety Administration (FMCSA) estimates about 200,000 people hold one.

The cohort is valuable because to earn it, a holder has to be work-authorized, which means they have a genuine Social Security number. Most then build an established credit file and often a registered trucking business by operating as drivers. Their identities are real and frequently deemed creditworthy by FIs — they clear KYC and eCBSV checks and win the credit, accounts, and loan approvals that thinner-file immigrants and temporary residents rarely get. They are also hard to flag: the "non-domiciled" marker, though printed on the license face, is frequently left out of the barcode, so an automated ID scan usually can't tell one from an ordinary CDL. And because these holders are here temporarily and often leave the country, misuse of their identities tends to go unreported.

That has long made the population a target for fraudsters. The reuse of departed immigrants' identities is Assumed Identity Abuse (AIA), which we described as "growing and maturing" in our 2H 2025 report. In the CDL variant, newly formed trucking LLCs finance vehicles across many lenders, make a few payments, then go silent, sometimes followed by credit "washing" through false identity theft disputes.

What's new this period is that policy is turning some of these holders into first-party fraudsters. A federal rule that took effect in March 2026 restricts the credential to a handful of visa categories, and FMCSA projects that roughly 194,000 holders will exit the freight market as their licenses expire or are revoked. The rule deports no one; what it does is take away the livelihood of a large, credit-active population on predictable dates — the expirations printed on their licenses — along with much of their reason to protect a credit file they may never use again. Set against an enforcement environment of asylum denials and removals that is already pushing many holders out of the country, it creates a straightforward incentive: run up credit, then leave. In recent casework, we have begun to see holders doing this themselves — running up credit on their own genuine identities before going silent, rather than falling victim to a takeover.

The limits of the data are worth stating plainly, too: nothing we've seen shows that non-domiciled CDL holders commit fraud at higher rates than anyone else. Some who leave are victims of identity takeover, and some who stay will simply default because they've run out of money, which is ordinary credit risk, not fraud.

Even so, the setup warrants attention: a large pool of credit-active identities set to lose their income on a known schedule, with rings already organized around them. A further rise in both AIA and departure-timed bust-out looks likely, though it will probably be gradual, since licenses expire at renewal rather than all at once.

Conclusion

The fraud ecosystem continues to evolve and mature. In this report, we have detailed a drop in fraud rates — particularly the identity theft rate — over the course of the half. That decline is real, but it should also be understood in context: even at its *lowest* point in the half, the identity theft rate never dipped below 5%. Last year, the first-half identity theft rate was 4.81%. The threat is still there. Rates may have lowered, but they are not *low*.

At the same time, both the marketplaces serving the fraud ecosystem and fraud attacks themselves have become more sophisticated. The days of scammers shaking you down for Google Play gift cards are gone; now they're targeting vulnerable identities to extract value via HELOCs and 401(k)s. The days of identity thieves you can catch easily are going, too; the most advanced fraud rings are now using legacy PII and residential proxies to counteract fraud-fighting measures and make their applications look indistinguishable from the real thing.

While we've been tracking first-party fraud for too short a time period to draw conclusions about its trajectory, it is clear that it represents a serious threat. And given our findings that the average identity age on first-party fraud applications is younger than what we typically see on applications overall, there may be some truth to the idea that it is being fueled in part by viral social media posts that present first-party fraud MOs to users as "hacks" or "glitches," not crimes.

The intent of this report, however, is not to scare but to share knowledge. While fraud is evolving quickly, it can still be detected. AI might be able to beat a liveness check, but it can't rewrite history, or change the network-wide signals that models like our fraud scores leverage to assess risk.

The fight continues, but the better we understand fraud, the better our chances of defeating it.

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 other hard-to-detect forms of identity 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.

Headquartered in San Francisco, the company was founded in 2017 by Naftali Harris and Max Blumenfeld and has raised $85M to date from investors including Andreessen Horowitz, Craft Ventures, and NYCA Partners, among others.

Proudly serving 500+ partners

Including:

13 of the top 15 banks

8 of the top 10 credit unions

2 of the top 3 telcos

Plus: Government, Tax Prep, Education, Utilities, and more.

Every day, SentiLink verifies 3M+ identities and flags 150k+ high-risk applications to help our partners reduce fraud losses and verify more legitimate customers.

Catch more fraud →

Appendix

Methodology

Industry fraud rate methodology (identity theft and synthetic fraud rates)

We began by categorizing a set of roughly 200 out of SentiLink's 500+ partners whose application data allows us to label their applications by industry, and who fit the industries we've identified for inclusion in this report: DDA, credit card lending, auto lending, telecom, and other consumer lending (a category that includes consumer lenders that don't fall into any of the above categories, such as Buy Now, Pay Later providers and AltFi lenders).

For those partners, we pulled data on all applications created between 1/1/2026 and 6/30/2026 (inclusive) that were scored using SentiLink's Synthetic Fraud Score version 1.8.1, 1.8.2, or 2.0.0, or SentiLink's Identity Theft Score version 1.7.2, 1.9.0, or 2.1.0. These model versions were the most widely used by our partners in 1H 2026. SentiLink maintains granular precision tables for these models derived from real-world application labels, which allow us to realistically estimate the probability that a given application is genuinely fraudulent based on its score.

Using the scored applications and the corresponding precision tables, we assigned each application a probability of synthetic fraud and a probability of identity theft.

Then, looking at all of the applications associated with each partner, we calculated per-partner fraud rates for each fraud type using the following formula:

```
sum of fraud probabilities / total number of applications = fraud rate
```

Before calculating fraud rates, we excluded partners with insufficient data volume (fewer than 1,000 total scored applications for either score) and manually dropped partners with anomalous, misleading, or incomplete data, such as those with bulk uploads of risky applications that could artificially inflate rates, or partners with significant time gaps in their application volume. We also manually removed some applications from some partners where they could be confirmed to be test applications submitted by the partner themselves, rather than real applications from a consumer.

To account for differences in partner size while still reflecting fraud activity across the network, we grouped partners into three tiers by application volume for each industry. This approach prevents fraud rates from being driven entirely by a small number of large partners, while also reducing

the influence of volatility from very small partners. We then assigned higher weights to higher-volume partners and lower weights to lower-volume partners, and used per-partner fraud rates to calculate a weighted average for each use-case category. This weighted average represents our best estimate of industry-specific fraud rates for 1H 2026.

Industry fraud rate methodology (first-party fraud)

First-party fraud is a broad term that encompasses a wide variety of underlying identity fraud risk patterns. SentiLink's First-Party Fraud Score (on which this analysis is based) reflects the likelihood of first-party fraud risk across a wide range of first-party fraud vectors, though not all of them. The identity fraud risk signals that SentiLink has developed have been shown to be performant in tackling ACH fraud, check return fraud, credit card fraud, and several other first-party fraud vectors.

Due to data limitations, we approached the calculation of the first-party fraud (FPF) rate and all subsequent FPF analyses differently from the other analyses in this report. We took a random sample of applications from each partner in the fraud report cohort over the time period, weighted by volume such that higher-volume partners contributed more applications to the sample.

We also manually filtered out two partners whose first-party fraud rates we believed would have applied misleading skew to the final result. In both cases, these partners had first-party fraud rates *higher* than the average rate, so removing these partners resulted in a reduction of our reported first-party fraud rates in Auto Lending and overall.

We then analyzed this sample using a conservative score threshold — only very high-scoring applications were counted as fraud — and then calculated fraud rates for each partner, as well as weighted use-case and overall rates, using the following formula:

```
(sum of applications >= score threshold) / total number of applications
```

Weighting for the use-case rates was based on partner application volume; weights for the overall rates followed the same methodology described in the Overall fraud rate methodology section below.

Methodology for the supplemental first-party fraud analyses, such as the ITIN and time-of-day analyses, was the same as we used for those analyses elsewhere in the report, and is detailed further in the sections below.

Overall fraud rate methodology

Our overall fraud rates apply a weighted average by industry classification. Each industry's weight reflects its estimated share of the 2024 domestic market, based on total revenues. For Financials verticals, total revenues are deemed to be net interest income plus relevant fees and non-interest income streams; for Telecoms, our weighting specifically considers wireless equipment and related device financing revenues, as SentiLink products are primarily used in equipment/device financing applications.

Financials revenue figures used in these weighting calculations are sourced from McKinsey & Company's U.S. Banking Revenue Pool analysis and adjusted to account for non-banking revenue contributions to the total market size. To validate these estimates, we conducted an internal bottom-up analysis to size interest and non-interest incomes by industry. We estimated interest income by aggregating consumer debt balances held by banks and non-banks across industries and applying corresponding net interest spreads sourced from publicly available market data.

Telecom revenue figures are sourced from publicly available SEC filings for core market participants, including, but not limited to, Verizon Communications (VZ), AT&T (T), T-Mobile US (TMUS), and Charter Communications (CHTR).

Sources used include McKinsey & Company's U.S. Banking Revenue Pool analysis, FI Navigator, the Federal Reserve Bank of New York's Center for Microeconomic Data, Cox Automotive, the Philadelphia Federal Reserve's Consumer Finance Institute - Special Report, and public company securities filings.

Geographic data methodology

For geographic analysis, we used the same scored applications from the fraud rate analyses and grouped them by the zip code provided on each application. For each zip code, we calculated the fraud rate using the following formula:

```
(sum of fraud probabilities for applications in zip code) / total
number of applications in zip code
```

We used the same weighting method used for our overall fraud rates here.

We then aggregated these zip code-level calculations to the county level. For zip codes spanning multiple counties, we assigned a primary county based on population density. We used public U.S. Census Bureau Gazetteer files to perform this county-level assignment and mapping.

For this analysis, we included only zip codes with at least 500 total applications during the studied period and at least 50 applications from each use case.

Cost estimate methodology

Average charge-off estimate methodology

To estimate the costs associated with fraud in various industries, we examined a total of 13 datasets that contained application information and performance data relating to applications across credit cards, auto lending, banking, and consumer lending, spanning the time period of 2022 to 2026.

From these datasets, we sampled 1.6M+ applications that had been scored by SentiLink as part of our retrostudy process — in other words, these are applications that were *not* scored by SentiLink at the time of application, and were retroactively scored later using as-of scoring. This allowed us to flag applications that were approved but subsequently charged off and use our scores to determine which of these accounts likely represents fraud loss, and further determine which specific type of fraud it most likely represents. (For example, a charged-off account that SentiLink's Identity Theft Score model scored at 950 is highly likely to be identity theft; in this analysis, we attributed the charge-off amount associated with that account to identity theft.)

We broke those 1.6M+ applications down by use case and calculated average charge-off amounts in this dataset for all applications, as well as applications that SentiLink had retroactively scored as high-risk. We also broke out average high-risk charge-off amounts for industries for which we had sufficient data (credit cards and consumer lending) as well as average high-risk charge-off amounts associated with a specific score (i.e., the average charge-off amount for an application that scored as high risk for identity theft).

Additional analysis methodology

First-party synthetic and third-party synthetic analysis

To analyze the distribution of first-party and third-party synthetic fraud, we looked at high-risk synthetic applications that could be classified with high confidence into one of the two categories based on our model scores (in addition to our general synthetic fraud model, SentiLink also offers specific First-Party Synthetic and Third-Party Synthetic scores). Specifically, we included only applications with first-party and third-party scores that scored as very high risk for *either* first-party or third-party fraud, but not both.

This analysis was conducted on a reduced sample size of 124,773 applications due to scoring availability constraints. Partners that receive only an overall synthetic fraud score — including all telecommunications partners — were excluded, as their applications could not be confidently classified as first- or third-party synthetic based on that score alone. Restricting the analysis to applications that scored as high risk for first-party or third-party synthetic fraud — but not both — reduced the pool further.

ITIN analysis

We compared fraud rates between applications submitted with ITINs and those submitted with SSNs using the same fraud rate methodology applied elsewhere in this report. We performed this analysis only for identity theft and first-party fraud. Since SSNs on synthetic fraud applications are by definition not the applicant's real SSN, we do not perform it for synthetic fraud.

Note: unlike our regular fraud rates, the ITIN and non-ITIN rates in these analyses are not weighted by partner volume. They also analyzed only applications that included a valid nine-digit entry in the SSN field, which is a subset of the applications used to calculate our overall and per-industry rates as not all SentiLink partners collect SSNs. For these reasons, the rates in the ITIN analyses are not directly comparable to the overall and per-industry fraud rates elsewhere in this report.

Fraud signals analysis

SentiLink's identity theft and synthetic models provide scores based on the data included in an application, but they *also* provide reason codes to indicate which specific element of the application most influenced the model's score, and a direction parameter that indicates whether the element in question moved the application's score up or down. Here's an example of what this looks like:

To analyze fraud signals for this report, we captured the reason codes associated with high-risk applications (applications with an 80% probability or higher of being fraudulent). We then analyzed these to determine which reason codes — i.e., which signals — appeared most frequently as the rank-one signal of risk in high-risk applications.

Day of week and time of day analysis

To analyze fraud rates by day of week and time of day, we used the same methodology used to calculate our industry fraud rates. The time data used here comes from our partners, who typically provide the date and time an application was submitted.

Note that we have taken a slightly different approach with our identity theft and synthetic fraud data as compared to our first-party fraud data.

In identity theft, the listed address typically corresponds with the victim's location, and isn't indicative of what time zone the fraudster is in. In synthetic fraud, the listed address may correlate with the fraudster's general location, but it may also be a mail-receiving location the fraudster has access to, or it may reflect some other phenomenon. For example, we highlighted Sheridan, WY as a synthetic fraud hotspot in our 2H 2024 Fraud Report, but not because fraudsters are actually located there — rather, they're taking advantage of a local regulation that makes Sheridan an attractive "virtual" base for fraud schemes.

So in the charts for those fraud MOs, since we can't know what time zone the fraudster is actually in, we based our analysis on U.S. Eastern time. In the continental United States, the Eastern time zone has the largest population, so it seemed the most logical choice for normalizing this data.

While IP address would theoretically be a better metric for determining the *fraudster's* actual time zone, unfortunately, most of our partners do not pass us IP address data. Also, the use of technologies such as VPNs and residential proxy IPs makes IP address information unreliable.

In first-party fraud, though, the fraudster is submitting an application that contains mostly their real PII, typically including their real address, so the application address *does* give us a good indication of the fraudster's actual location. For this reason, in this analysis we've charted specifically applications from fraudsters in the U.S. Eastern time zone to present a cleaner picture of when they're operating throughout the day.

Age analysis methodology

To analyze the average ages associated with identity theft applications, we sorted our dataset, filtering for only applications with a high probability of being fraud (75% or greater) and applications with a low probability of being fraud (5% or lower). We then calculated the average ages of these two subgroups, both per use-case and overall.

Fraud trend and case study methodology

The fraud trends and case studies in this report come from the observations of SentiLink's Fraud Intelligence Team, who manually review and label hundreds of thousands of applications every year, as well as from the research of SentiLink Head of Fraud Insights Dr. David Maimon and his team, who monitor and collect data from a variety of online fraud marketplaces on the clear and dark web.

Sample Sizes

Identity theft rates, synthetic fraud rates, and most other analyses	
Credit card applications	20,376,434
Auto lending applications	22,943,617
Other consumer lending applications	101,729,415
DDA applications	10,182,693
Telecom applications	15,524,177
Total applications	**170,756,336**

First-party fraud rates	
Total applications	31,513,338

First-party synthetic vs. third-party synthetic rates	
Total applications	124,773

Limitations

The numbers in this report reflect our estimate of the fraud rates experienced by our partners in the first half of 2026. However, there are some limitations to our approach that should be considered when interpreting the report, including but not necessarily limited to:

- Real identity fraud is a lagging indicator and may take months or years to confirm. Many instances of fraud will never be confirmed. The rates in this report are, necessarily, estimates, and they reflect identity fraud *attempts* rather than the rate of successful identity fraud.
- Our data reflects only what we can see from the applications of SentiLink partners, and only a subset of SentiLink partners and SentiLink partner applications were used in the analysis.
- Due to the number of steps we took to filter the data and the way we calculated the fraud probabilities, we are likely undercounting the true fraud rate by a small amount. We feel this is preferable to overcounting and exaggerating it.
- Geographic data is based on the zip codes provided in applications, which do not necessarily reflect the real-world locations of fraudsters or their victims.
- Zip codes do not always correspond to municipal borders, and in cases where a single zip code exists across multiple cities or counties, this can impact our assessment of the county's overall fraud proportion.
- Our first-party fraud analysis, as noted in that section and in the methodology section, is based on more limited data and should be considered preliminary and indicative rather than definitive.

Comparisons with Previous Fraud Reports

While the methodology we used in this report is mostly the same as what we used in the previous report, there are some significant differences that should be considered when making comparisons between the fraud rates in this report and previous reports.

Specifically, the list of partners whose application data is included in this analysis has changed for a variety of reasons, including:

- SentiLink's list of partners is always expanding, and this report includes data from newly signed partners for whom we did not have data in 2H 2025.

- This report includes applications scored using our latest Identity Theft Score, Synthetic Abuse Score, and First-Party Fraud Score models, none of which were used in analysis for previous reports.

- Partner application volume patterns may shift for a wide variety of reasons.

- Our filtering process for this report includes filtering out data from partners whose monthly application volume is too small, whose volume patterns suggest irregularities such as the bulk uploading of risky cases, or whose data we've determined to be irregular via other forms of manual review. As partner volume and use patterns vary, these filters result in a different final cohort of partners being included in the analysis for each report.

It is also important to note that both the model version and the methodology used to calculate our first-party fraud rate have changed slightly, so those numbers are not directly comparable to the first-party fraud results in the previous report.

Definitions

Synthetic Fraud involves creating a fake identity or combining a real identity with fake PII and using it to perpetrate fraud. Specifically, we define a synthetic identity as a name, DOB, and SSN that don't all belong to the same cohesive, real person. For more information, check out our blog post: What is Synthetic Fraud?

Subtypes of synthetic fraud included in this report's synthetic fraud rate analysis include:

- **First-Party Synthetic Fraud** describes when a fraudster uses their true name and DOB, but an SSN that doesn't belong to them. This is also sometimes called identity manipulation.

- **Third-Party Synthetic Fraud** describes when a fraudster invents a synthetic identity wholesale, combining a name, DOB, and SSN that are not associated with a real person. This is also sometimes called identity fabrication.

- **Synthetic Identity Theft** describes when a fraudster is using a real base identity that is *not* their own, but has inserted an SSN that doesn't belong to that stolen identity.

- **Friendly Fraud Synthetic** describes when a fraudster submits an application that contains a mix of their own PII and an SSN associated with a friend or relative — often either a deceased relative or a child.

SentiLink's fraud taxonomy also categorizes a variety of other types of synthetic fraud. For more information on how we categorize synthetic fraud and its subtypes, please refer to our whitepaper: Defining Synthetic Fraud.

Identity Theft describes a wide variety of fraud types that involve using another person's identity.

Subtypes of identity theft included in this report's identity theft rate analysis include:

- **Identity theft** describes when a fraudster uses stolen PII from an individual with whom they have no personal connection.
- **Friendly fraud** describes when a fraudster steals and uses the complete identity of a relative or associate.
- **Inmate identity abuse** describes when a fraudster with no connection to the victim steals the identity of an incarcerated person.
- **Assumed Identity Abuse (AIA)** describes when a fraudster is leveraging the U.S. identity of an immigrant who is no longer in the United States.
- **Scam victim** describes when the identity listed on an application *is* the person behind the application, but they are likely to be the victim of a scam. (SentiLink's model does not explicitly label scam victims; however, our research into scam victims has shown that applications from scam victims often do contain elements that result in the application being scored as high-risk for identity theft.)

First-Party Fraud (in the context of this report) describes applications in which the applicant is using their own core PII elements, but the application has an elevated risk of inaccuracies in other aspects of the credit history and identity information.

Credits

Lead author and researcher: Charlie Custer

Design: Michael Purcell

Additional contributions: Burt Helm, SentiLink Fraud Intelligence team, SentiLink Partner Success team, SentiLink Marketing team, Paul Byrne, Kelley Byrnes, Cam Dufty, Larry Gomperts, Jason Kratovil, David Maimon, Andrew Villacis, Mike Yu, Lucas Zhang

Consulting and review: Naftali Harris, Sam Mulholland, John Leitner


SentiLink