

Whitepaper

First-Party Fraud in Deposit Accounts




Table of Contents

| | |
|----|--|
| 03 | Fighting First-Party Fraud in DDAs with Targeted Scores and Flags |
| 05 | SentiLink's DDA First Party Fraud Scores and Flags: Under the hood |
| 08 | Evaluating Performance: DDA FPF Scores and Flags |
| 11 | API Response |
| 12 | Conclusion |


Fighting First-Party Fraud in DDAs with Targeted Scores and Flags

Conventional wisdom, along with the opinion of Federal banking regulators, has long characterized deposit products as the lowest-risk financial offerings.¹ While not immune from fraudulent activity over the decades, these products are not extensions of credit or loans on a balance sheet, so scams like “check kiting” could be considered as a cost of doing business to be absorbed. In recent years, however, that has begun to change. First-party fraud affecting Demand Deposit Accounts (“DDAs”) has grown into a material problem for financial institutions in the U.S. Many large FIs have shared with us that they lose tens of millions of dollars per year to this.

At SentiLink, we are focused on building solutions that outperform and overcome the limitations of other approaches in the marketplace to help our partners address their first-party fraud challenges. Through our research we have identified nine key signals that — when ingested via our new DDA First Party Fraud Scores, as one or more individual Flags, or as model inputs — provide predictive and actionable insights into fraudulent first-party behavior.

Why a different approach to tackling first-party fraud M.O.s in DDAs is needed

Two common ways fraudsters scam financial institutions (“FIs”) using their own identity are via ACH fraud and check fraud. With ACH fraud, a consumer opens a



¹ See, for example, the [FFIEC BSA/AML examination manual](#) which ranks “Resident Consumer Account (DDA, Savings, Time, CD)” as presenting the lowest customer risk.

DDA, funds the account via ACH from another bank (that is, they debit the other account), spends or exits the funds, and then disputes the ACH transfer at the bank from which they transferred the funds, claiming that the transfer of funds was unauthorized.

With check fraud, the consumer deposits a check that will ultimately be returned as altered or fictitious, then exits the funds before the nature of the check is discovered. A common variant of this is known as “card cracking,” which involves the consumer explicitly making a claim that the check was deposited due to their account being taken over. Thus, these two M.O.s — ACH fraud and check fraud — often share the characteristic of involving false claims of account takeover.

These problems persist despite consortia, new and old, that attempt to address this problem. Consortia generally work by bringing together a set of FIs willing to share with each other labels of individuals they’ve concluded have committed first-party fraud. These conclusions can be reached after manual review or after observations indicative of one of the fraud M.O.s above. Such observations could include a high number of customer-initiated ACH returns or bounced checks, a high dollar amount for a single customer-initiated ACH return or bounced check, or other signs of first-party fraud on DDAs, like mule activity. Once a member FI flags one of these cases, other members will be able to see that the case is marked as “bad” by another member, along with the reason it was flagged.

Clearly, such indicators are valuable to participating members. Nevertheless, with this approach, any individual consortium will have an incomplete picture of the universe of first-party fraudsters. There will naturally be some fraudsters that have not opened accounts with any member of a given consortium, creating blind spots. In addition, some fraudsters may have accounts with many consortium members but have not yet gone bad on any of those open accounts.












For a consortium to be able to flag an individual as bad, that individual must both commit fraud at a consortium member FI and be successfully flagged as having committed that fraud by that FI.

SentiLink's DDA First Party Fraud Scores and Flags: Under the hood

SentiLink's new DDA First Party Fraud Scores — with two scores separately targeting ACH fraud and check fraud — and suite of associated flags offer FIs flexibility and optionality to best fit their needs. These signals provide coverage on all U.S. adults who apply for financial products, not only consumers flagged by a consortium. We look for historical patterns and current signals that indicate whether a presented identity is more likely to be associated with first-party fraud schemes on DDAs as described above. We further validate the performance of these patterns by comparing them to a curated set of first-party fraud labels developed in conjunction with our partners for DDAs. Performance is measured by the precision and recall of each pattern.

These qualitative and quantitative assessments have led us to focus on the following nine Flags, all of which can be returned via API at the time of application.

- | | |
|--|--|
|  Ties to risky PPP loans |  Ties to identity theft via address |
|  Previous synthetic fraud |  Ties to identity theft via phone |
|  Velocity of DDA applications |  Ties to fraud-furnished phones |
|  Short history phones |  Phone ties to other identities |
|  Compromised identity | |

The following definitions give detailed descriptions of the flags and the context in which they are useful:



Ties to risky PPP loans: This flag indicates whether an individual is tied to PPP loans, based on an exact match on name and address, processed by certain lenders who gave out a high volume of fraudulent loans when the program was in place; these lenders were ultimately called out in the congressional report on PPP fraud.² Flagged individuals were often recruited via Telegram channels to apply for these loans. These same dark web channels are also used to recruit individuals to participate in first-party fraud schemes aimed at DDAs.



Previous synthetic fraud: This flag indicates whether the individual applying has committed synthetic fraud in the past, which highlights the individual's increased propensity to obfuscate parts of their identity and participate in fraud schemes more generally.



Velocity of DDA applications: This flag indicates whether the applicant has an unusually high number of DDA applications prior to their current application. There are very limited reasons why someone needs many DDAs; however, having many accounts can allow someone to more easily engage in the schemes described above. Moreover, having DDAs at multiple FIs makes it easier to leverage the blind spots of existing consortia and commit fraud at different institutions. An individual could take advantage of those established accounts in succession before bad behavior at one FI is reported to other members.



Short history phones: This flag indicates whether the applicant has an unusually high number of phones used for a short period of time prior to their application. Using many phones for a short period of time can be an indicator of participation in money mule or other fraud schemes.



Compromised identity: This flag indicates whether an unusually high number of identity theft events have been tied to the applying identity (name, DOB, and SSN) recently prior to the application. Individuals whose identities have recently been used for identity theft present not only elevated risk of this type of victimization, but also elevated risk of participating in other related fraud or money mule schemes.

² [“We are Not The Fraud Police: How Fintechs Facilitated Fraud in the Paycheck Protection Program.”](#) Staff Report of the Select Subcommittee on the Coronavirus Crisis. December 2022.

- **Ties to identity theft via address:** This flag is similar to the compromised identity pattern above, except that it checks whether the application address has been tied to an unusually high number of identity theft events recently prior to the application.
- **Ties to identity theft via phone:** This flag is similar to the compromised identity pattern above, except that it checks whether the application phone has been tied to an unusually high number of identity theft events recently prior to the application. This flag often detects instances where the applicant is, in fact, an identity thief but in this instance is applying using their own true identity along with a phone that they've used previously in applications using stolen identities.
- **Ties to fraud-furnished phones:** This flag indicates whether the individual has been tied to a phone that has previously been used in the course of identity theft events to apply for financial products on behalf of many identities. Such individuals may have actively participated in money muling schemes in the past and thus may be susceptible to be recruited for such schemes in the future.
- **Phone ties to other identities:** This flag indicates that the phone on the application has been recently tied to other unrelated identities, which indicates their elevated risk for participating in check and ACH fraud schemes in the future.

Targeted choices for FIs

SentiLink's DDA First Party Fraud solutions can be consumed by our partners in three ways:

- As separate Scores targeting ACH and check fraud, which are based on machine learning models trained using numeric features associated with each individual flag as inputs;
- As a set of discrete Flags that can be used directly and to build decision logic;
- As feature inputs to proprietary in-house models or rules. We return via API not only flags, but the numeric features as well to support model development and use in rules.

For the last use case, partners should note that these new signals will soon be integrated with our Facets attribute solution and available as a bundle, if desired.

Evaluating Performance

We often refer to a few key terms in our research into identity crimes and first-party fraud to measure the impact of a solution: recall, general hit rate, and relative likelihood. Recall refers to the amount of true positives captured. General hit rate refers to the amount of cases that are flagged (whether a true positive or not) out of the entire data set. Relative likelihood refers to the ratio between recall and general hit rate.

For example, imagine a set of DDA applications, some of which were onboarded and resulted in confirmed first-party fraud. A signal with 2% recall and 0.25% general hit rate would capture 2% of all of the cases of first-party fraud, and all of the applications captured would make up 0.25% of all applications. A flagged application would be 8X more likely to result in first-party fraud.

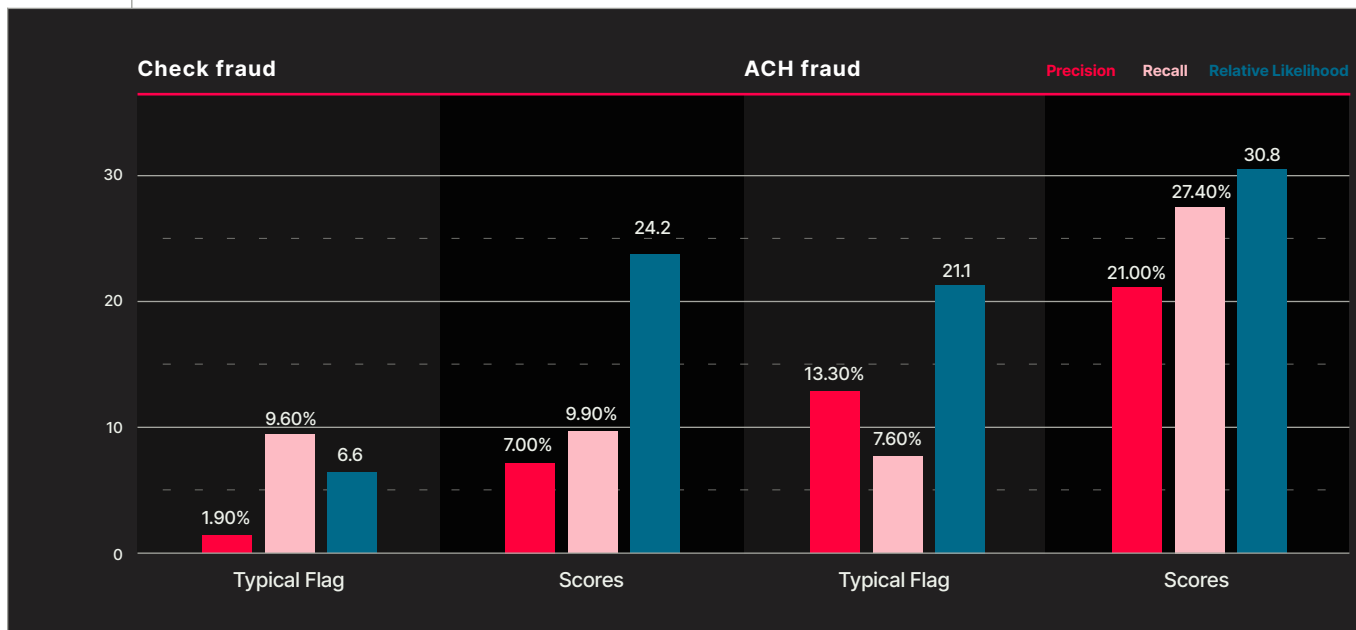
To evaluate the performance of each Score and set of Flags, we utilized two case study datasets of applications: one from a depository institution facing check fraud, and one from a depository institution dealing with ACH fraud issues. We present below the performance of scores, then the performance of flags on the same data sets for comparison.

Performance of scores

As with our industry-leading Synthetic and ID Theft Scores, we expect many will wish to receive intelligence on suspected first-party fraud in DDAs via scores. For such FIs, the most relevant metric will be precision and recall at a score cutoff that balances false negatives and false positives well for that FI.

As might be expected, these scores — which use the flags and accompanying numeric features as inputs — are able to predict first-party fraud in DDAs with greater accuracy than the flags themselves.

- **Check Fraud:** For the check fraud dataset, at the score cutoff that maximized the F1 score,³ the **recall was 9.9%** - slightly higher than a typical flag. **Relative likelihood was 24.2X**, meaning an applicant scoring above the cutoff was 24.2X times more likely to deposit a check that would later return as altered or fictitious compared to a typical approved applicant, and about three times as precise as a typical flag. On this particular dataset, **precision was 7.0%**.
- **ACH Fraud:** For the ACH fraud dataset, we see a similar performance boost from the score compared to the median flag, though this is driven more from recall than precision. At the score cutoff that maximized the F1 score, the **recall was 27.4%**, **relative likelihood was 30.8X**, and **precision was 21.0%**. This is a stark improvement in performance compared to the median flag, which is outlined in the next section.



³ That is, a cutoff with the best balance between false positives and false negatives, assuming the cost of a false positive equals the cost of a false negative.

In addition to targeting check fraud and ACH fraud, we have observed similarly high levels of performance of our Scores where first-party fraud in DDAs presents different challenges. For example, on a dataset from a large fintech facing high levels of transaction disputes due to fraud, a score cutoff that would achieve 20.6% recall had a relative likelihood of 13.9X. A custom scoring model built using the same features achieved similar recall of 20.2%, but with an even higher relative likelihood of 29.4X.

Performance of flags

The performance of the nine flags varies across FIs and fraud types; FIs may target different customer bases, or fraudsters may target different FIs at different times.

- **Check Fraud:** For the same check fraud dataset, the median recall of the flags was 9.6%, with a median relative likelihood of 6.6X, meaning **an individual flagged was 6.6 times more likely to deposit a check that they would later return as altered or fictitious** compared to a typical approved applicant. The median precision was 1.9%.
- **ACH Fraud:** For the same ACH fraud dataset, the median recall of the flags was 7.6%, with a median relative likelihood of 21.1X, meaning **an individual flagged was 21.1 times as likely to initiate an ACH transaction that would result in a consumer-initiated return** compared to a typical approved applicant. The median precision was 13.3%.

API Response

The following example API response shows an API call that includes two flags and their numeric metadata. The specific flags an FI wishes to receive can be custom configured by SentiLink, and FIs can receive any of the nine flags noted above:

API Response: Flags

```
[
  {
    "flag_name": "sentilink_dda_velocity_flag",
    "flag_version": "1.0.0",
    "flag_value": true,
    "metadata": {
      "dda_application_velocity_120d": 3,
      "dda_application_velocity_180d": 4,
      "dda_application_velocity_360d": 4,
      "dda_application_velocity_720d": 6,
      "dda_application_velocity_1080d": 6,
    }
  },
  {
    "flag_name": "sentilink_many_short_history_phones_flag",
    "flag_version": "1.0.0",
    "flag_value": false,
    "metadata": {
      "num_short_history_phones_1y": 0,
      "num_short_history_phones_2y": 0,
      "num_short_history_phones_3y": 0,
      "num_short_history_phones_5y": 1,
      "num_short_history_phones_7y": 2
    }
  }
]
```

API Response: Scores

The following example API response shows an API call for each score:

```
[
  {
    "sentilink_first_party_check_fraud_score": {
      "version": "1.1.0",
      "score": 912
    }
  },
  {
    "sentilink_first_party_ach_fraud_score": {
      "version": "1.1.0",
      "score": 570
    }
  }
]
```

Conclusion

First-party fraud in DDAs has become a significant pain point for many financial institutions that is hard to detect and leads to significant losses. SentiLink has developed solutions to address this problem that do not rely on a consortium model, and instead defines a set of patterns that indicate whether a presented identity is more likely to be associated with first-party fraud on a DDA in the future. These signals can be consumed in several ways: as targeted scores aimed at each of the previously mentioned fraud M.O.s separately; as flags to incorporate into decisioning processes; or as features to be fed into a model as well as building rules. Furthermore, for FIs willing to provide large numbers of labels (i.e., tags of individuals who have committed first-party fraud on their DDAs), SentiLink is able to train and host custom models designed to produce scores predictive of that specific FI's observed fraud behavior, with these DDA First Party Fraud Flags and their metadata as inputs. Contact your SentiLink Partner Success Manager to receive the complete API docs for this product and learn more.

