

Precise ID® for Account Opening fraud suite




Fraud scores

Finding fraud is like finding needles in a haystack. If you were tasked with finding needles in a haystack, it would require an enormous amount of effort to look at every piece of hay one at a time. A better way to find as many needles as efficiently as possible would be to grab handfuls of hay, and with a quick glance at each handful decide if you think each handful contains no needles, a few needles or many needles and divide the haystack into three piles — one for each group.

When you have finished separating the hay, you'd know the best place to find needles is the pile where you put the handfuls you thought contained many of them.

A fraud score works in a very similar way. In the case of loan applications, a fraud score takes information from an application and makes thousands of comparisons to previous examples and other information to figure out if the current application looks like those that turned out to be good customers, or those that were later found to be fraudulent.



Imagine creating a 'scoring' system that would help you tell the difference between a needle and a piece of hay based on characteristics that make them different.

A needle might be 'scored' in the following way.

 $+$  $+$  $= 3 \text{ points}$

1 (hard to bend) **1** (very thin) **1** (shiny)

On the other hand, a piece of hay would more likely be


 $+$  $+$  $= 30 \text{ points}$

10 (easy to bend) **10** (much thicker) **10** (flat in color)

Now imagine you pull a fiber from the haystack. It bends easily, so you give it an 8; it's thinner than half of the other things in the pile, so you give it a 4; it's definitely not shiny, so you give it another 9. If you add those together: 8 (it bends) + 4 (it's relatively thin) + 9 (it's flat in color) = 21.

Given what you know about the fiber, it has more in common with hay than a needle. The illustration in Figure 1 gives you a visual picture of how this works. Based on its characteristics, the fiber is closer to hay than it is to a needle.

Figure 1



Precise ID® for Account Opening fraud suite


Fraud scores work the same way. The scientists who build them figure out what characteristics are different for good versus fraudulent identity uses by looking at past examples where the outcomes are known. When the model is finished, it's able to assess the characteristics of identities that have never been seen before and place it somewhere on the good/fraud spectrum. Using much more complicated math, the model does the same thing as illustrated in Figure 1 — with every calculation nudging an identity toward the “good” or the “fraud” end of the spectrum.

When this process is applied to a large number of examples, groups start emerging where the amount of goods versus frauds (or hay versus needles) becomes consistent and predictable. Think of the piece of fiber from our earlier example. Its score of 21 puts it closer to being hay than a needle, but that's far from being absolutely certain. If we had one thousand other examples that also received a score of 21, we'd find that some are actually needles. Models make mistakes.

This is the challenge of analytics. Perhaps if we had looked at additional characteristics, or done a better job of measuring the thickness of the fibers, we wouldn't have made those mistakes, but mistakes happen and it's the job of the analytical scientists to constantly revisit the process to improve it and help it adapt when conditions change. This is particularly true in fraud, where the person committing fraud is actively trying to outsmart the tools that are built to stop them. It's a cat and mouse game. As fraudsters change their technology and tactics, the tools that stop them have to keep pace.

TIP: Fraud scores NEVER provide 100% certainty that something is good or fraudulent. The goal of this type of analytical approach is to predict the likelihood that an event is one or the other. Even when confidence is really high, there's a chance the prediction isn't correct. It's the nature of statistics. This fact doesn't take away from the value of a fraud score. Despite their imperfections, model predictions are still very valuable. Using our haystack/needle example, Figure 2 (below) shows how a model used on a large set of examples would create separate piles of hay where the number of needles are consistent and predictable. Some piles have more, some less. Even though there is still a mix of hay and needles in each group, we know which pile contains the highest number of needles. In a real-world example where we are trying to find fraud among a group of good customers, this is critical. There are consequences for inconveniencing good customers when we are investigating possible fraud. Good customers might get frustrated and take their business elsewhere. Knowing the number of fraudulent vs. good customers in a particular group allows users to make informed decisions about how to handle them. This information helps businesses manage the balance between fraud losses and customer convenience.

Figure 2



100 Needles
High Risk

Low Fraud Risk Score



50 Needles
Medium Risk

Medium Fraud Risk



10 Needles
Low Risk

High Fraud Risk Score


Model performance

The key term used to describe model performance is “false positive.” A false positive is a good customer who’s accidentally classified as high risk. In terms of business impact, a false positive results in wasted costs to investigate a case where fraud isn’t occurring. It can also result in the loss of a good customer if they become frustrated with the delay or a request for additional information in order to be approved. The term “false positive” is widely used to discuss model performance because it describes the cost and customer impacts clearly.


The example in Figure 3 revisits our last illustration of needles in a haystack. This time, we’ve noted the number of pieces of hay and needles in each stack. Now we can see the number of pieces of hay in relation to the number of needles in each stack. The ratio of hay to needles is the false-positive ratio.

Here’s what these numbers tell us: If we look at every item in the stack on the right to find the needles, we’ll look at 1,000 total items and find 10 needles — or 100 items per needle. If these were loan applications, that would mean spending money and impacting 100 customers in order to stop a fraud. Compare that to the stack on the left, where only 10 reviews are needed in order to stop a fraud.


Figure 3



100 Needles
900 pieces of Hay
= 1,000 total items where
100 are needles
1000/100
False Positive = 10 to 1



50 Needles
950 pieces of Hay
= 1,000 total items where
50 are needles
1000/50
False Positive = 20 to 1



10 Needles
990 pieces of Hay
= 1,000 total items where
10 are needles
1000/10
False Positive = 100 to 1

If there were only one way to handle fraud risk, a lender would have obviously wanted to focus on the stack where the false-positive ratio is lower (10:1). Fortunately, there are many ways to handle fraud risk, and they vary in terms of their costs and how much they impact customers. Those that are low cost, fast and invisible to the customer can be used to weed out fraud even where the false positives are higher.

TIP: False positive is an important measurement because it tells us so much about the costs and customer impacts of fraud prevention efforts, but it needs to be considered as part of a bigger picture. Specifically, we need to consider how much fraud we're actually stopping. For example, in the stack on the right in Figure 3, we see a false-positive ratio of 10:1. Imagine if that stack could be broken down further. What if we could create a smaller stack that had a false-positive ratio of 3:1? Initially, that sounds terrific. We'd definitely want to focus on the stack where we only spend money on working three cases to stop a fraud instead of 10, right? But what if that stack only has three items in it and one fraud. By focusing only on the false positive alone, we've lost sight of the other 99 frauds. It's a balancing act. Discussions about false positives should also address how big the population is, and how much fraud is in it.

TIP: Be skeptical of fraud tools that tout low false positives without talking about how much fraud they catch or the size of the population they impact. It's easy to find the small number of really obvious frauds with low false positives. The trick is finding a significant amount of fraud and still keeping false positives low.

What to do when fraud risk is detected

The Precise ID for Account Opening fraud scores give you the ability to divide your haystack into smaller piles knowing which piles will have the most needles. It requires some additional action to complete the task of separating good customers from the fraud. You need a plan to know when to take action and what action to take.

You must decide when the level of risk is high enough to take action. Usually, that's based on a combination of how much might be lost when fraud occurs and the cost of handling the risky cases. It's a balancing act. The work to find fraud costs money, and sometimes the verification process chases away good customers. Together these impacts add up, and at some point, they can be as expensive as the fraud itself. The following section will help you create a simple, low-cost process that protects good business opportunities and pushes away the fraud.



Actions to take

When action does need to be taken, it's usually divided into three categories based on how much involvement or difficulty it presents for the customer:

- **No impact:** As the name suggests, no-impact actions are invisible to the customer. It's the first level of action and its goal is to quickly look at an additional piece of available information that helps you confirm that the application is legitimate. No-impact actions should be fast and low cost. The goal at this level is to keep the good applicants moving and weed out the handful of those that need a bit more scrutiny.
- **Mild impact:** Mild impact typically requires some additional involvement from the customer. While it's an extra step for them, it should be something they can respond to immediately. It might be as simple as asking to reconfirm application data, provide a previous address, or confirm information about a previous or other account with you. As a rule of thumb, the customer should be able to do what they're asked without leaving their chair.
- **TIP:** Assume every customer is a GOOD customer and treat them accordingly. Remember that fraud is relatively rare and even groups with high fraud risk will have several good customers for every attempted fraud.

- **TIP:** Avoid making customers feel like they're being singled out. No one wants to think they failed a fraud check. Let the applicant think a request for more verification is a normal part of the application process for everyone.
- **High impact:** High-impact treatments are a last resort. This level of verification requires the customer to take some action that's likely outside of the application process. Because of the inconvenience, some good customers won't comply and will be lost. High-impact actions are reserved for the very highest risk situations, where the risk of fraud outweighs the value of having a good customer. Examples of high-impact verification might require providing a copy of an identity document, proof of income or visiting a branch to complete the application process.

TIP: Whatever type of verification you choose for high-impact verification should give you high confidence that customers who satisfy it are acceptable. At the end of the waterfall, there are only two groups of customers — those you accepted and those who didn't respond. The right set of verifications allow good customers to pass and push the fraud to abandon the process.




The waterfall approach

Take actions using a waterfall approach like the one pictured in Figure 4. Notice that the mild-impact and high-impact actions are used on two groups — those that received a score that leads to the action and those that received a less risky score but had trouble getting through the verification that was requested.

TIP: We use the term waterfall to describe an applicant who might cascade from one verification request to the next until they pass or confront a request that they choose not to comply with.

Figure 4



TIP: The actions that you take must be consistent. Knowing the action taken will be critical later, and if they are applied consistently, that's easy to figure out. You might later want to modify your process if too many customers aren't passing — or if fraud starts slipping through. It's easier to figure out what to change if you know what you did.

Practical use

Now that you know the basics, it's time to put the Precise ID for Account Opening fraud suite to work. The Precise ID for Account Opening fraud suite gives you three different scores. Each one is built to detect a specific type of fraud. The descriptions below explain each type and why they have to be treated differently.

Identity theft: Identity theft is what most people think of when the topic of fraud comes up. Identity theft is the simple act of using someone else's identity. As more and more information becomes available online, it's become easier for criminals to impersonate a victim. Information about all of us is readily available through public data sources, through our social media presence and through the troves of digitally stored data that's been stolen and made available for sale.

Type of treatment needed: Since identity theft involves a victim, the key to stopping it is to find the real owner of the identity. If they can be found, they'll confirm if they're aware that their identity is being used. The challenge is that perpetrators will try to make contacting the identity owner difficult. An identity thief will provide contact information (phone number, email, etc.) that leads back to the thief and not the victim.

TIP: Make sure you're confident that information you use to contact the applicant is correct. Don't rely on the information from the application unless you verify through another source.

First-party fraud: First-party fraud refers to a person using their own information to commit fraud. First-party fraud overlaps considerably with credit risk for that reason. In many cases, lenders simply consider these cases to be like any other customer who can't repay and defaults on their obligation; however, more and more lenders are realizing that some of these defaults seem deliberate. Whether it's intentional or not might be debatable, but when a person rapidly acquires access to credit, uses it quickly, then ceases paying on every obligation at the same time, there's reason to think it might have been intentional. If nothing else, the behaviors that lead to this kind of default can be seen and the outcome can be predicted, so it's gaining acceptance as a type of fraud.

Type of treatment needed: First-party fraud can't be handled the same way as identity theft because the person perpetrating it is the owner of the identity. The information they provide is correct, and if contacted, they'll approve the activity. First-party fraud risk is also unique because it can change over time. For example, a person who takes a new job and relocates may show some temporary signs of risk as they open new accounts and change their spending to finance their new lifestyle as they get settled. When first-party fraud risk is present, the best approach for lenders is to exercise care in offering additional financing to the customer until they're established and their accumulation and use of credit begin to stabilize. People who may initially appear to have risk of first-party fraud can become valuable customers later on. It's often worth accepting them and monitoring them to see if they become more or less risky over time.



Synthetic identity fraud: Synthetic identity fraud involves the creation of a fake identity that's used to commit fraud. Once created, the identity is "planted" across a wide variety of consumer data sources to appear more legitimate with the ultimate goal of establishing a consumer credit profile. Real people can be duped, paid or may be actively involved in the scheme to add the synthetic identity to existing accounts that they own. Perpetrators also seek out lenders with weak processes that allow synthetic identities to open accounts on their own. In both of these scenarios, synthetic identity can be reported to credit reporting agencies. Once they have a foothold, more and more lenders are drawn into the scheme as the identity looks more and more like a real borrower. In the end, synthetic identity fraud behaves like first-party fraud. Balances rapidly rise and payments stop suddenly. The difference between synthetic identity and first-party fraud is that the intent behind synthetic identity is 100% criminal. Where a true person may temporarily look risky and later stabilize to become a great customer, synthetic identity ends in losses every time. There's also no recourse for lenders when synthetic identity occurs. There's no one to pursue for collections, litigation or any other means to recover any portion of the debt.

Type of treatment needed: Because synthetic identities always go bad, the stakes are very high. Synthetic identity fraud is a risk lenders simply can't take. The trick to treating synthetic identity is to apply high-impact treatments as the first step and require identity documentation that would be difficult for a synthetic identity to obtain or create. Fortunately, the fraud model in the Precise ID for Account Opening fraud suite is extremely effective. It finds synthetic identity with very little effort.

TIP: It's important to understand the three types of fraud and have a plan to handle each type. Some types of verification work great for one type of fraud and poorly for others. When you know what you're facing, you can use the best approach.



Workstation cheat sheet

The Precise ID for Account Opening fraud suite delivers three separate fraud scores that will indicate the level of identity theft, first-party and synthetic identity risk for each identity for which you request the service.

Each score was developed by Experian's team of analytical scientists who've spent decades developing these products using hundreds of thousands of examples of fraud and hundreds of million of examples of good borrower behavior. While many customers think of Experian® only as credit reporting agency, it's important to recognize that Experian is a global data services company with data assets that go far beyond credit profiles. As individuals, each of us leaves a footprint in the world that's much larger than our credit history. As we pursue our interests and engage with the world around us, a much more robust picture is created. Experian's broader business focuses on this larger picture. This provides Experian with fantastic insight into identity that we leverage to find the patterns that drive our analytics. The fraud scores in the Precise ID for Account Opening fraud suite are the product of our analytical expertise combined with our extensive data assets that go beyond consumer credit information.

You'll find a workstation cheat sheet on the last page. It's intended to be used as a quick reference for people assigned to assess the scores and take action when fraud risk is detected. The cheat sheet will help ensure that actions are applied consistently.

Suggestions on when to take action are provided. Reviewers can quickly see if a score falls within a range where action needs to be taken. The suggested ranges are intended as a starting point. They can be adjusted so the level of effort fits the resources you have available.


Suggested actions are provided. For each score range, you'll find one or more suggested actions. You may have access to tools, resources or information that can be more easily used and that's encouraged.

False positive estimates are provided. Based on the research done to build each model, an estimate gives you a sense of the work, costs and customer impact for stopping each fraud. Your specific experience will vary, but this provides a starting point. If the score ranges for taking action need to be modified, this information helps set expectations for what will happen in terms of increased/decreased work and fraud risk.




Precise ID® for Account Opening Fraud suite workstation cheat sheet


Identity fraud score



First-party fraud score



Synthetic ID score



Note: If multiple types of fraud are detected, address the one that requires the **HIGHER** impact action.