Socure

Digital Identity Fairness & Inclusion Report

A Socure whitepaper on demonstrating equitable results in consumer identity verification and financial inclusion in the US financial system

Table of Contents

Executive Summary

Socure believes that fair and equitable identity verification and fraud assessment should be a requirement of every identity solution provider in the marketplace. We apply these principles to each machine learning model solution created at the company, from start to finish. That means our approach persists from feature selection through entity resolution, as well as for investments we make in the data used to generate our best-in-class solutions. Ultimately, we aim to drive the accurate automation of 100% of all identity-related compliance and fraud management decisions made online, regardless of race, gender, age, ethnicity, or socioeconomic status.

Socure's industry-leading efforts in machine learning and AI democratize access to services and benefits by increasing approval rates and reducing unnecessary, bias-driven friction in fraud and identity scoring for thin-file, "credit invisible" consumers, fast-growing immigrant populations, and vulnerable minority groups.

To prove it, we recently partnered with two of the country's top five leading banking institutions to conduct separate analyses:

1. First, we tested the proprietary clustering and advanced entity resolution technology that powers our Consumer Identification Program (CIP) solution against a leading incumbent provider. Notably, a CIP program prescribes "the minimum standards for financial institutions and their customers regarding the identity of the customer that shall apply in connection with the opening of an account at a financial institution" and is required by Section 326 of the USA PATRIOT ACT.

   The primary goal of this exercise was to benchmark our ability to increase automated CIP decisions for the younger, 18-25 year old population who typically have limited credit and information footprints. However, we also discovered that the pioneering, patented clustering algorithms we use combined with the advanced entity resolution techniques we deploy, create major benefits in automated acceptance of every ethnicity across the board when compared to a leading incumbent solution provider.

2. Separately, we performed a unique approach to conduct a disparate impact analysis of our machine-learning Sigma Identity Fraud model to identify potential bias against key demographic groups that our signals or machine learning approach might inadvertently create. As additional context, our Sigma Identity Fraud model is used to detect potential third party identity fraud that might cause financial loss to an organization and cause harm to consumers in the form of identity theft.

Approximately **45 million** people within the U.S. aged 18+ are credit invisible[1]

Socure's industry-leading efforts in machine learning and AI democratize access to services and benefits by increasing approval rates and reducing unnecessary, bias-driven friction

1. The CFPB Office of Research, *"Data Point: Credit Invisibles"*, May 2015

The goal of this research was to evaluate how Socure's technology uniquely reduces the negative ethnic, racial, age and socioeconomic bias that can easily creep into the risk scoring of certain fraud models, resulting in unwarranted friction and reduction in the ability of these vulnerable populations to obtain services, such as lines of credit, at fair interest rate levels.

In this paper, you will find a detailed explanation of our testing methodologies and the results of these two extraordinary analyses. Furthermore, we are extremely proud to present these findings, which we consider a testament to the meticulous research and development by over 200 brilliant and deeply passionate Socure employees who make up the Data Science, Engineering, Product, and Data Acquisition teams.

Clearly, identity verification and fraud mitigation technologies play a very important role in enabling financial services organizations to reach certain populations in an increasingly digital world. The following, detailed findings offer *proof* that Socure delivers the most accurate and inclusive identity solutions available in the market today.

Some key points to call out based on our findings to help highlight the importance and potential impact of using AI/ML technologies for good:

If used exclusively across the country's financial services industry, Socure's technology could more than double the amount of automated positive approvals made in the US in any given year in the growing immigrant population, the 45 million+ group of credit invisibles, especially younger, thin-file 18-25 year olds, as well as minority populations who have historically been subject to unfair bias.

Unfair bias and the associated friction that it introduces to younger and underserved populations in high growth digital channels can result in abandonment rates of 50% or more, limiting those populations' ability to access some of the industry's most beneficial financial products as easily and seamlessly as other groups and classes.

To put our findings into more simple, everyday terms, let's use the example of a well-known social media company. Imagine they were required to rigorously validate the identities of their users as a regulated company would have to. And to satisfy this standard, they use one of the weaker, incumbent solutions in the marketplace. Inevitably, this results in significantly fewer auto-approvals and higher drop offs for users subjected to friction, with the standard dropoff rate projected to be 50% or more. Our intent is to illustrate the power of Socure's frictionless solution, over and above a typical legacy provider, based on our internal research. Using a legacy provider, the resulting, negative impact would be as follows:

- The company could lose roughly half of their 18-29 year old population, dropping their total US user base from 222M users to a under 175M, overnight.
- Instead of having 28M daily active users, the company may only see 19M logins. The diverse voice of the platform's user population would likely be silenced and left out, with 21% fewer women, 28% fewer Black, 36% fewer Hispanic, and 46% fewer Asian voices participating in the platform discussions, drastically changing the direction of public comments and debate.
- Finally, if the platform's current $10B valuation was based on its total user population, each user would be worth around $23, and the company's valuation would immediately drop by $2-4B, up to 40%!

The Changing US Demographic Background and the Increasing Digitization of Account Opening in the US Can Create Unfair Bias for Vulnerable Populations

Unfair bias occurs when an FI applies a practice (e.g., identity risk assessment) uniformly to all applicants while, at the same time, that practice has a discriminatory effect on certain segments of the population. An example might be a person of a certain age or demographic unnecessarily being subjected to increased friction during onboarding. Though these impacts may not be intentional, they may be the result of hidden biases in machine learning algorithms, in the data sample that is gathered and used for modeling as well as in the entity resolution techniques deployed.

15% of Black and Hispanic people are credit invisible compared to **9%** for White and Asian people[2]

Credit invisible groups continue to be negatively impacted by traditional CIP programs' primary reliance on credit header data

According to the Consumer Financial Protection Bureau (CFPB), approximately 45 million people within the US aged 18+ have limited credit history. This group of people are either not able to be scored or have no credit bureau files, rendering them credit invisible. This segment of the population is effectively shut out of the financial system that is available to the other 80% of the population.

According to a Citi Global Perspectives and Solutions economic impact report, about 15 percent of Blacks and Hispanics are credit invisible (compared to 9 percent of Whites and Asians), and an additional 13 percent of Blacks and 12 percent of Hispanics have unscored records (compared to seven percent for Whites). These differences are observed across all age groups, suggesting that they materialize early on in the adult lives of the consumer population, and persist thereafter.

The unfortunate reality of many of today's entrenched identity and fraud solution providers is that credit header data is used as the primary driver for their CIP solutions that validate consumers at account opening. Because underserved younger populations, immigrants, and certain ethnic groups may not reside in the credit reporting ecosystem, those consumers often experience substantially higher friction to receive the same services and experience reduced opportunities in credit availability as well as general access to services and benefits becomes more friction-filled.

The explosive growth of US foreign born populations will create further inequalities if legacy CIP technologies are not enhanced. According to Pew Research, the US foreign-born population reached a record of 44.8 million

2. Citi Global Perspectives & Solutions, "*Closing the Racial Inequality Gaps*", September 2020

people in 2018. Beyond having reduced information footprints, these growing immigrant groups present an identity verification challenge for digital account openings because they more often use different surname and first name characteristics as well as hyphenated names and apostrophes.

International naming conventions do not lend themselves to legacy forms of entity resolution. Immigrant populations have both a data coverage and an entity resolution hurdle that Socure has resolved to allow for higher approval rates in CIP programs, and to optimize the coverage of identity information that informs highly-accurate, machine learning fraud models for these populations.

The growth of digital channels in financial services is accelerating the problem. Today, more than ever, opportunities for consumers to engage with FIs and Fintechs occur through digital channels. Various online banking statistics show more and more individuals in the US are embracing digital banking. Post-COVID, this trend's impact may be increasingly harmful to underserved groups because many of today's most essential financial products are only available to consumers through web and mobile-based channels.

It is important to note that digital onboarding is not in itself a cause of inequality. For example, a 2021 Pew Research Center survey found that smartphone and tablet ownership is not divided across any statistically significant racial and ethnic differences. Technology access is not a driver of inequality. Rather, it is the onboarding speed (or lack thereof) and friction caused by antiquated CIP programs and biased fraud mitigation strategies that creates unequal access to financial services and benefits.

Organizations that can enable greater inclusion for credit invisibles not only promote positive social change, but tap into and build a market of long-term, loyal customers. Failure to do so means that FIs and Fintechs, for example, will not effectively reach key, growing segments of the population, resulting in missed revenue opportunities, customer growth, and development of a positive reputation within the communities they serve.

Immigrants today account for **13.7%** of the U.S. population, nearly triple the share (4.8%) in 1970[3]

Closing racial gaps can increase the U.S. GDP by $5 trillion over 5 years[4]

3. Pew Research Center, *"Key findings about U.S. immigrants"*, August 2020

4. Citi Global Perspectives & Solutions, "*Closing the Racial Inequality Gaps"*, September 2020

The Socure Advantage in Regulated Customer Identification (CIP) Programs: How ML-Driven Clustering and Advanced Entity Resolution Technologies Increase Fairness and Access for the Underserved

Well governed models are key to achieving highly accurate identity outcomes. While this white paper is not a formal outline of Socure's response to regulatory requirements, the company cares a great deal about fairness in the credit and banking approval processes faced by credit invisibles and protected class segments. And even though we do not offer credit solutions that can be used for adverse action, our machine learning models and CIP solutions can be used to verify whether a credit applicant is who they say they are, and therefore perform many of the same tests that are included in CFPB examination procedures to ensure our machine learning algorithms and CIP compliance solutions do not negatively impact these vulnerable populations. This gives them access not previously available from a fully automated solution.

Socure has aligned its model governance efforts with US regulatory bodies and supports its customers with:

- Model risk management programs
- Sound model development
- Validation practices as well as guidance on the use and implementation of our models

Socure's data science team regularly assesses its fraud models and input variables for any adverse effect on protected classes in the form of disparate impact. We employ methods that encompass proxies for protected class status, implicit bias of non-representative data or algorithmic assumptions. We do this to avoid any correlation with subsequent performance (i.e., over- or under-predicted fraud outcomes) for each protected class.

Socure Disparate Impact Analysis Testing

Socure builds and offers classification models that employ hundreds (out of eight thousand available) of independent variables called "predictors" to determine the likelihood that a presented identity actually belongs to the entity presenting that identity. Socure conducts analysis on these parameters using a variable neutralization process. The variable neutralization methodology measures the impact of a single variable on the accuracy of the model output in addition to any disparate impact on a selected population.

The company's methodology includes established and academically accepted research methods to infer gender and race from provided names and location of residence. Socure uses an analysis algorithm that leverages multiple historical sources for nuanced and reliable outcomes, recognizing and allowing for certain limitations, such as relying upon a government-defined, binary gender. Our analyses rely on common research methods such as the Bayesian Improved Surname Geocoding (BISG) proxy methodology commonly used by the CFPB to determine race.

Rigorous model governance in combination with the methods and practices used to develop and test Socure's models ensures that the company has the most accurate and inclusive models of any identity verification provider in the industry.

Proving Fair and Equitable Outcomes for All

Test 1: CIP Auto Acceptance: Socure vs. A Legacy Incumbent Provider

Socure recently ran a head-to-head accuracy and auto-acceptance rate comparison against a legacy identity verification provider on behalf of a top five US bank. To evaluate objectively the lift that Socure's advanced CIP solution would provide over the incumbent provider, Socure ran the test based on the exact CIP rules used by the incumbent's solution. This established a fair comparative analysis of auto approval rates across key demographic segments, and the ultimate findings were compelling.

As the results illustrate, the impact of Socure's solution was positive and meaningful for every tested demographic. However, the increases in auto-acceptance rates in the Gen Z, Asian, Hispanic, and female populations were the greatest. These specific increases are due to focused combinations of data assets and clustering technologies deployed by Socure to achieve superior results in these critical segments.

As established previously, younger populations have a smaller information footprint. Therefore, there are immense challenges in validating the Gen Z segment via CIP technologies that are overly-focused and reliant on traditional credit header data. Generally, younger populations do not become credit-active until their mid or late 20's. Socure invests heavily in finding and evaluating unique but reliable data sources that maintain accurate data on younger populations.

Socure's deliberate choice of entity resolution technologies and the addition of clustering algorithms reduce bias for these protected class segments. Furthermore, we observed that the increased lift in auto-acceptance for the 62+, Black, and female populations can be attributed to these advanced technologies.

Going deeper, the disparity in performance of legacy providers can be explained through over-reliance on static edit distances, soundex phonetic encoding, and rigid lexical rules that have long been favored in entity resolution problem solving. However, when stress tested and researched by the Socure team, we uncovered positive bias in Anglo-Saxon names in addition to a negative bias related to Latin, Middle Eastern, and Southeast Asian names due to typical hyphenation or short string length. To combat this inequality, an ensemble of phonetic, edit-based, token-based, and sequence-based algorithms were deployed to decrease bias and increase performance.

The impact of Socure's solution was positive and meaningful for every tested demographic

Primary goal was to benchmark our ability to increase automated CIP decisions for the 18-25 year old population. However, we also discovered that the pioneering, patented clustering algorithms we use combined with our advanced entity resolution techniques, create major benefits in automated acceptance of every group across the board when compared to a leading incumbent solution provider

Auto-Acceptance Rates

Socure: **94.10%**
Legacy Provider: **47.40%**

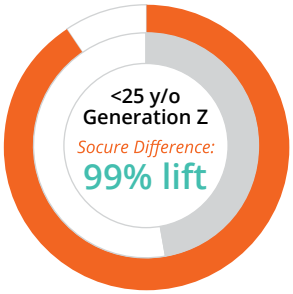

Socure: **93.40%**
Legacy Provider: **80.70%**

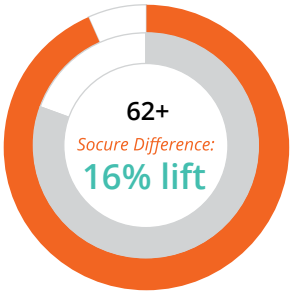

Gen Z increases are a testament to Socure's data acquisition investments for the 18-25 year old population

Socure: **92.10%**
Legacy Provider: **76.20%**

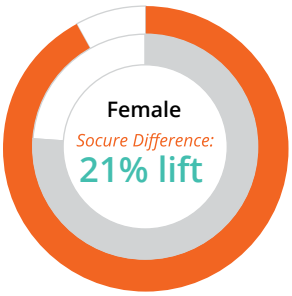

Socure: **88.40%**
Legacy Provider: **75.00%**

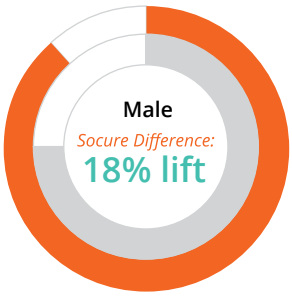

Groups including non-white ethnic groups, ages 62+, and female populations positively benefit from Socure's overall technology and deep expertise in machine learning

Socure: **89.80%**
Legacy Provider: **76.80%**

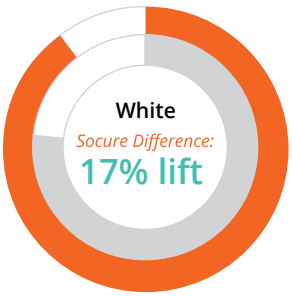

Socure: **91.10%**
Legacy Provider: **66.80%**

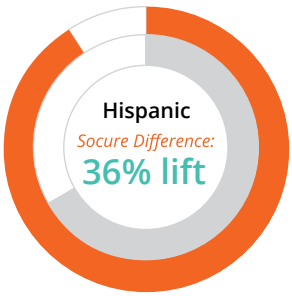

Socure: **92.40%**
Legacy Provider: **63.50%**

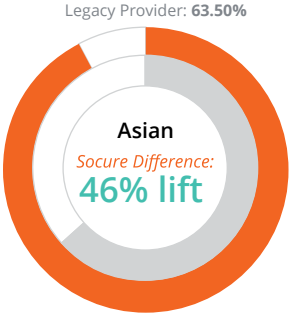

Socure: **92.20%**
Legacy Provider: **72.20%**

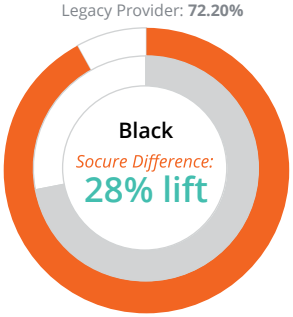

TEST 2: Disparate Impact and Bias Assessment of the Socure Sigma Identity Fraud Model

In a separate test of Socure's ability to mitigate bias and deliver fair and equitable results, we recently completed a disparate impact analysis of our machine learning Sigma Identity Fraud model on behalf of another top five US bank. The purpose of the test was to identify any potential bias that our signals or machine learning approach might generate. As additional background, our Sigma Identity Fraud model is used to identify potential third party identity fraud that might cause financial loss to an organization and cause harm to consumers in the form of identity theft.

The Bank's Four Primary Goals:

- Enable the bank to implement fraud models that satisfy requirements for bias testing prior to delivery
- Use bias testing that is more accurate than its incumbent bias testing capabilities
- Assess populations beyond those currently accounted for within the CFPB model
- Establish the bank as a leader on bias testing within the industry

What Socure Tested:

- Hundreds of thousands of records
- Evaluated top 200 predictors
- All variables were tested across multiple demographics
  - Race
  - Ethnicity
  - Sex
  - Age
  - Socioeconomic class
  - Immigration status

Disparate Impact & Bias Test Analysis Results

There was no negative bias identified with any predictor, across all population segments tested. While there is some variation in variables for certain populations, as indicated in the analysis graphs, none of the variables had a statistically significant negative impact on model performance. As a result, Socure was able to clearly demonstrate fairness in our models for this customer.

Socure uses a proprietary approach, leveraging the variable neutralization methodology to identify negative effects of individual variables on the accuracy of the model output for each key demographic.

The analysis for each population segment is visualized on a graph. The graph has three sections to track the implications of disparate impact. Each section represents a different combination of predictiveness and disparate impact. Each model feature is represented as an 'x' on the graph.

The next two pages show the results for each population segment we tested.

Results Legend

AREA	PREDICTIVENESS VS. IMPACT	JUSTIFICATION
Green	Low disparate impact across a range of predictive quality	Low disparate impact means that there is no fair lending discrimination concern
Yellow	Significant disparate impact but also significant predictiveness	The fact that a policy or practice creates a disparity on a prohibited basis is not alone proof of a violation. When a lender's policy or practice has a disparate impact, the next step is to determine whether the policy or practice is justified by "business necessity." The justification must be manifest, and may not be hypothetical or speculative
Red	Low predictiveness; high disparate impact	There is no justification for predictors in this quadrant.

The x-axis (AUC) represents the degree of predictiveness, while the y-axis (ARG) represents the degree of impact on protected classes. Using this method, Socure data scientists can visualize the effect that any one predictor has on the model for specific population segments. As a result, Socure data scientists, and our customers, can easily identify features that may be creating bias for a selected group of people. A feature (x) on the graph falling into the yellow or red areas of the graph may potentially indicate bias and trigger a deeper analysis. Socure data scientists employ a process to evaluate the predictiveness and impact of any variable that may be indicating bias. The goal of this process is to identify and remove variables that both indicate significant disparate impact and predictiveness.

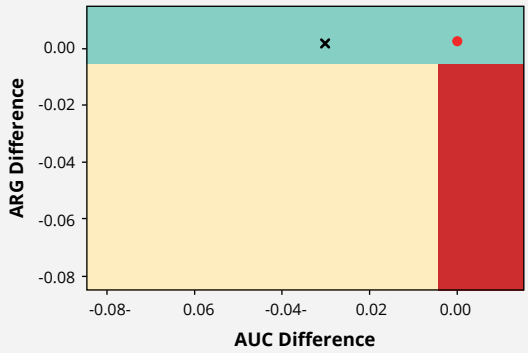

Definitions:
AUC: Area Under the ROC Curve (AUC) represents the probability that a random positive example is classified positively as opposed to that of a random negative example.

ARG: Acceptance Rate Gap (ARG) : Acceptance rate gap is the acceptance rate difference between protected class and base class.

Disparate Impact & Bias Test Analysis Results

In the graphs, the x-axis represents the AUC difference and the y-axis represents the ARG difference. Features, represented by X's, may display moderate bias variation along the Y-axis; however, if they have low predictive value, then the feature does not present a disparate impact. A feature with both significant predictiveness AND disparate impact would be seen in the yellow section of the graphs. **The Sigma Identity Fraud model did not demonstrate or reflect any bias for the protected classes analyzed in the study.**

Protected Class: Female
Base Class: All other Classes

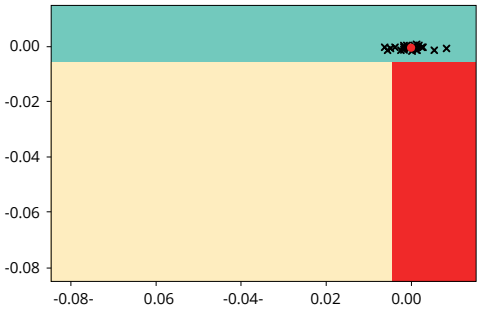

Protected Class: Black
Base Class: All other Classes

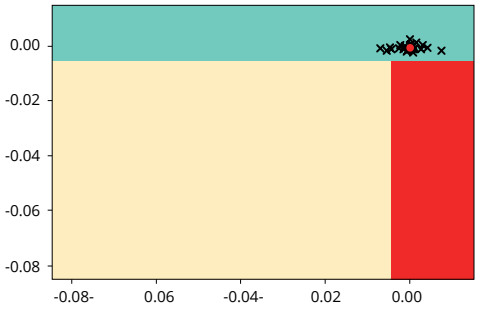

Protected Class: Generation Z
Base Class: All other Classes

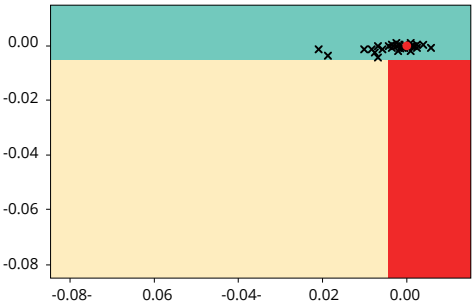

Protected Class: Hispanic
Base Class: All other Classes

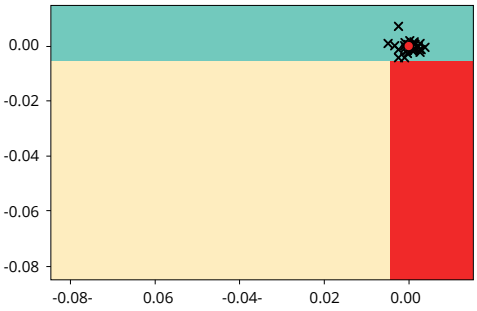

Protected Class: Low Income
Base Class: All other Classes

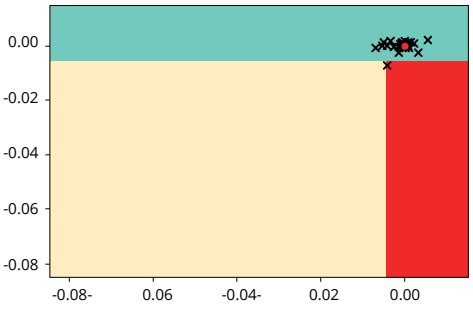

Protected Class: Asian
Base Class: All other Classes

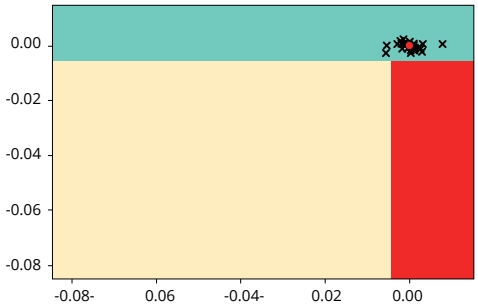

Protected Class: Immigrants
Base Class: All other Classes

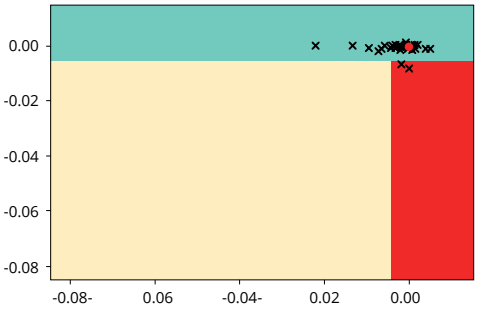

Conclusion

Through empirical data analysis, the Socure team firmly believes it is advancing our mission and commitments while remaining at the forefront of industry in using AI for good. Our solutions enable our customers to identify and serve more hard-to-reach populations than any other identity provider. We take great pride in striving for our goal to auto-approve 100% of identities in the digital space, helping to positively impact millions of people along the way.

Financial institutions and Fintechs that are concerned about employing fraud and CIP models that are fair, transparent, and explainable are welcome to contact us. Let us run a proof-of-concept test for you, like the ones described in this paper, to see how Socure can help you improve accuracy and inclusivity in your fraud and CIP programs.

If you are a member of a supervising regulatory body, we welcome the opportunity to walk you through how we mitigate disparate impact and bias in our models.

Socure offers identity and fraud models that produce higher accuracy and greater inclusion to enable financial services organizations to leverage *AI for Good*

Schedule a demo

Socure

Learn how Socure can help you grow and power financial inclusion. Connect with us today

Socure is the leading platform for digital identity verification and trust. Its predictive analytics platform applies artificial intelligence and machine learning techniques with trusted online/offline data intelligence from email, phone, address, IP, device, velocity, and the broader internet to verify identities in real time. The company has more than 750 customers across the financial services, gaming, healthcare, telecom, and e-commerce industries, including four of the top five banks, seven of the top 10 card issuers, three of the top MSBs, the top payroll provider, the top credit bureau, the top online gaming operator, the top Buy Now, Pay Later provider, and over 100 of the largest fintechs.