Socure

FRAUD INSIGHTS REPORT

Fraud, Identity Risk, and Age Evasion in Online Sports Betting and Prediction Markets

Executive Overview

By all accounts, Super Bowl 60 was a major success for online sports betting and prediction market organizations.

During the defined Super Bowl window — spanning pre-game (4:00 p.m. EST) to post-game (11:30 p.m. EST) — the organizations we monitored collectively added more than 2.5 million new customers. That represents roughly 10% of typical annual growth compressed into a single event.

This surge was driven primarily by legitimate consumers. Approval rates were strong, application quality was high, and overall fraud rates on Super Bowl Sunday were diluted by the scale of good user volume. Notably, operators that ran national commercials did not see an elevated rate of underage applicants during the game, supporting the fact that commercials ran during the Super Bowl did not overly play to younger audiences. Age mismatches, our proxy for younger populations, dropped by 3% on game day.

To summarize, a disproportionately large share of applicants were appropriately aged, applying from legal states, and presented lower fraud risk relative to typical weekly averages. Simply put, this was a **highly-effective mass onboarding event**.

That said, absolute fraud volumes did increase, as did attempts by underage individuals to unlawfully access platforms. You will see that this is a consistent theme throughout Super Bowl weekend.

We also observed several distinct behavioral patterns that are important to understand as the industry prepares for another high-volume acquisition moment in 2027.

The sections that follow summarize those patterns, break down fraud volumes and vectors, and outline recommended actions.

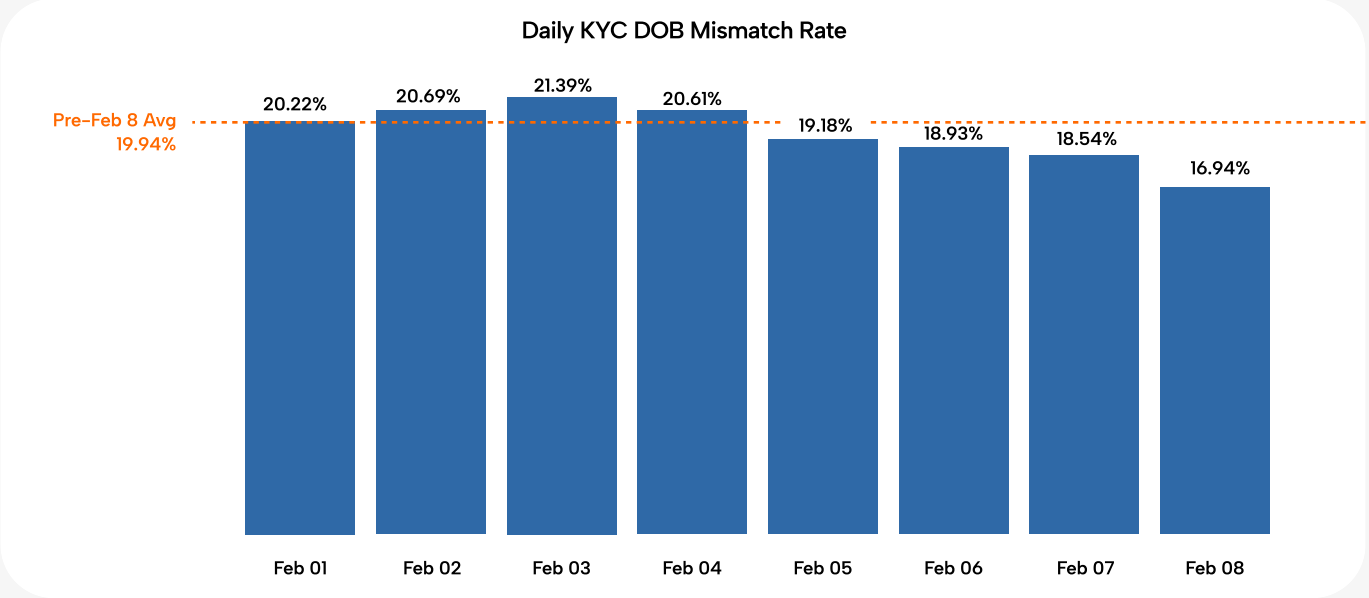

First-Party Fraud Gets an Early Start

First-party fraud (FPF) rates for new applicants were increasingly front-loaded, spiking in the days leading up to the game. For this analysis, we used a very high score cutoff of .995 for both the Identity Manipulation Model Score (identifies real applicants manipulating their contact channels to evade detection) and the Dispute Abuse Model Score (flags known or repeat dispute abusers).

Surprisingly, the FPF rates were substantially lower on the day of the game, as you can see from the graph below. Again, Super Bowl day rates were down across the board (age mismatch, state-to-state issues, fraud scores, etc.) given the volume of "good" identity populations that attempted to sign up.

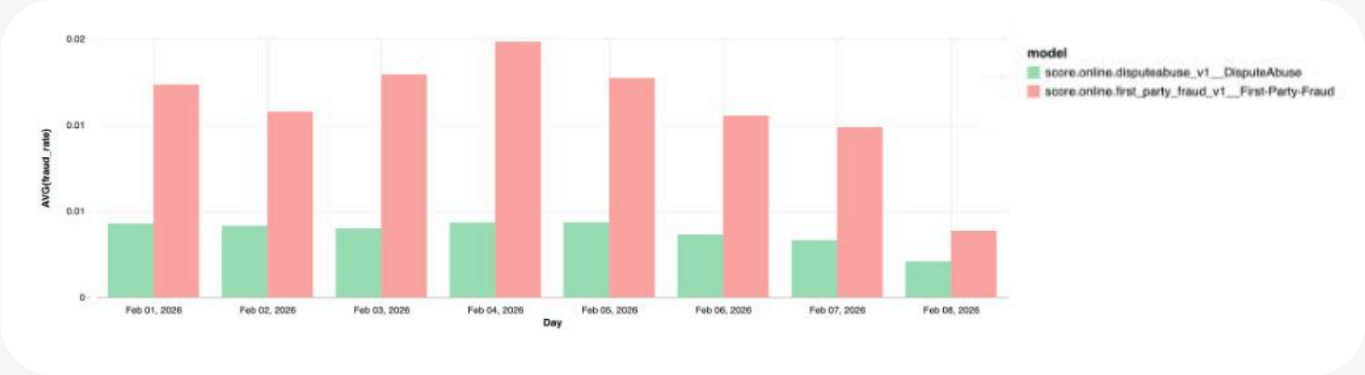

The Identity Manipulation Model, which predicts when a real identity is being manipulated (specifically, changing their contact channels) to create loss, fired at a much higher rate than the Dispute Abuse model.

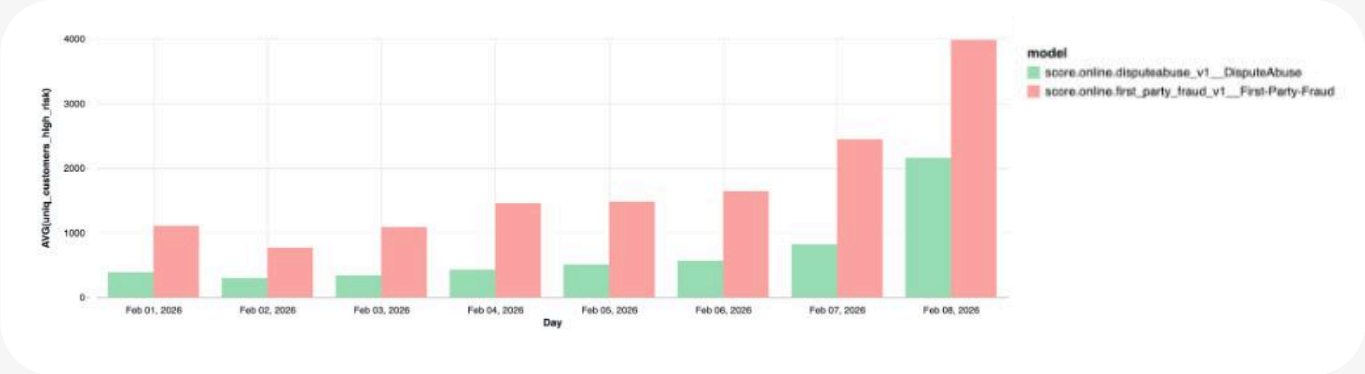

Overall, we did see a volume spike in likely first-party fraud attempts the day of the Super Bowl for aforementioned reasons (high volume of applicants). The spike in volume was substantial at 40% above even the highest of all the other days leading into the Super Bowl.

Upon further analysis, we confirmed the applicants who scored high for identity manipulation model scores were also marked as high risk across several organizations within Socure's FPF consortium. Again, for our customers not calling Socure's First-Party Fraud models, these patterns can easily be missed, as the applicants are a real person who can pass all other identity and fraud checks — yet, clearly show intent to perpetrate some form of fraud. These models should be considered for further analysis.

In addition to the two model scores returned, customers also receive a list of 100+ raw signals and intelligence from the consortium with each response. For this pool of applicants, the dominant risk signals were primarily identity-based and consortium-driven, rather than dispute related. For example:

- Matched Fullname and date of birth (DOB) combination onto at least 10 and up to 17 other transaction consortium institutions - Extensive cross-institution identity reuse and velocity
- Found at least 2 accounts that were opened within 90 days from each other associated with the given identity
- Found at least 7 distinct emails associated with the given full name and DOB combination
- Seen input phone number in at least 13 other transaction consortium institutions
- Found at least 7 distinct emails associated with the given full name and DOB combination
- Found at least 7 distinct phones associated with the given identity
- Prior data contributions reflecting closed accounts
- Multiple emails and phones, often times recently created, tied to the same real-world identity
- Velocity of cross platform name, DOB, and SSN matching

Similar to Age Discrepancies, State Mismatches, and Sigma Identity and Synthetic Fraud Scores, we saw rates for FPF start to increase after halftime, again because good volumes tailed off during those periods.

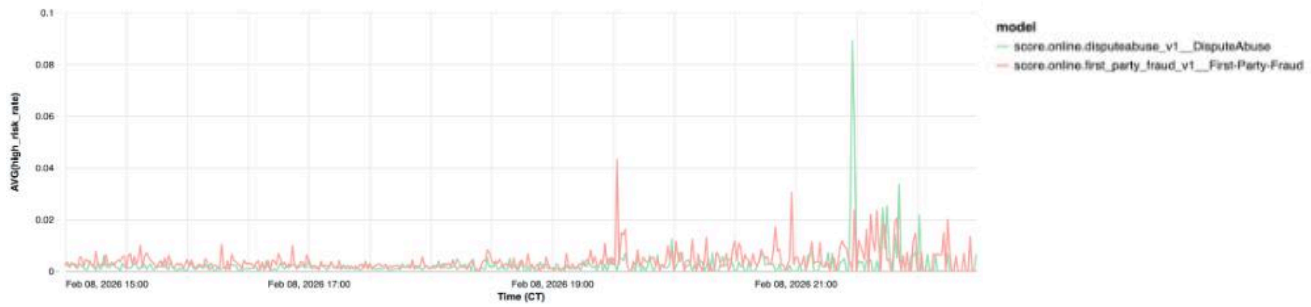

Suggested Next Steps

1. Given this behavior, we recommend testing both of Socure's First Party Fraud models; the Identity Manipulation model and the Dispute Abuse model at the point of new application. We expect this trend to repeat itself for years to come, and encourage customers to prepare accordingly.
2. Suspicious or fraudulent "disputes" used to develop the Dispute Abuse score are generally related to fraudulently-disputed ACH charges and other payments tied to retail purchases. Because of this, the standard Dispute Abuse model will become even stronger with"better remorse" tags from online gaming and predictive markets. We'd love to partner and further advance these predictive models to best serve your team and the wider gaming industry.

Identity Theft and Synthetic Identity Fraud

While overall fraud rates appeared low due to the surge in legitimate applicants, we observed a dramatic increase in the overall volume of identity theft and synthetic identity attempts.

The graph below shows fraud attack rates, measured by Sigma Identity and Sigma Synthetic models, using a high score cutoff of .99 to identify a fraudulent application.

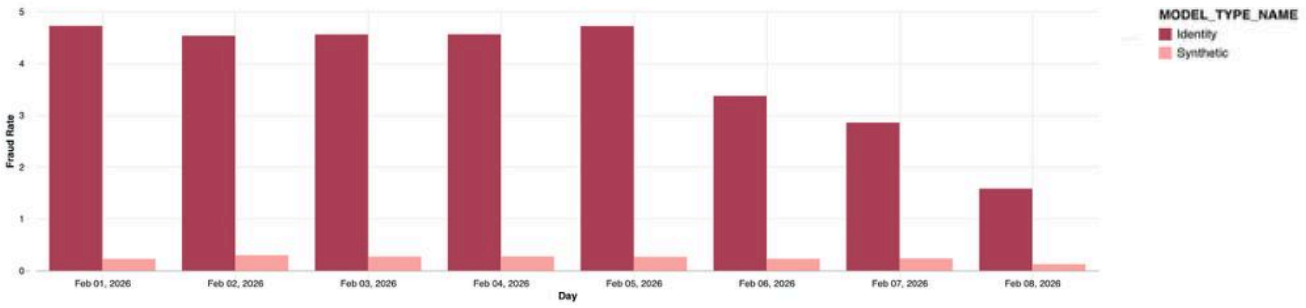

The graph above reflects overall fraud attack rates leading into the SuperBowl. The drop in fraud rates can be deceiving. Fraud rates dropped as we moved closer to the day of the Super Bowl because the amount of "good" customers applying for online gaming skyrocketed.

The volume of fraud attacks on Super Bowl Sunday holds an inverse relationship to fraud rate. The graph below shows the overall daily fraud attack volumes, again measured by the Sigma Identity and Sigma Synthetic models.

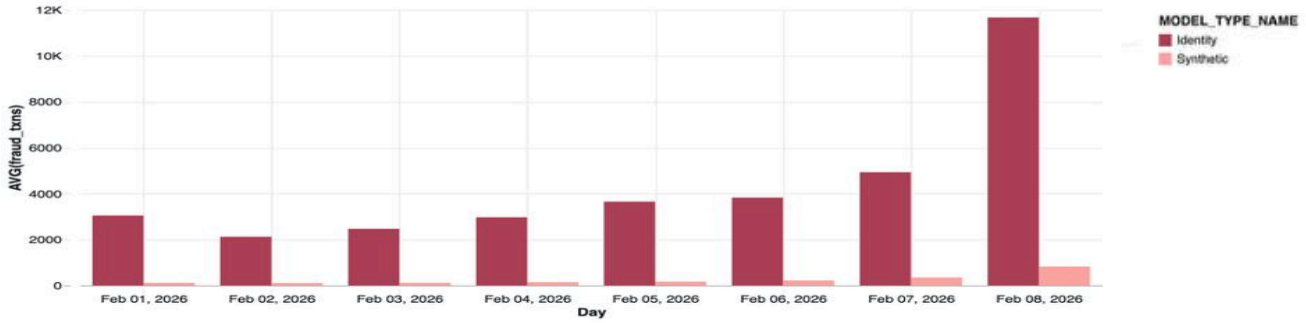

The graph above reflects overall fraud attack volumes leading into the SuperBowl.

Fraud volumes grew 376% in the week leading up to the game, nearly 4.8 times the week prior. Identity theft played a larger role than synthetic fraud during the Super Bowl period.

Both identity theft and synthetic identity fraud rates rose in the third and fourth quarters and peaked during the post-game window as legitimate volume tapered.

Under normal conditions, identity fraud represents approximately 3–5% of incoming applications. During the Super Bowl window, that rate increased substantially, peaking at 21% during post-game analysis. Synthetic identity fraud followed a similar pattern, rising alongside identity fraud as application volume remained elevated.

Identity Fraud Rates

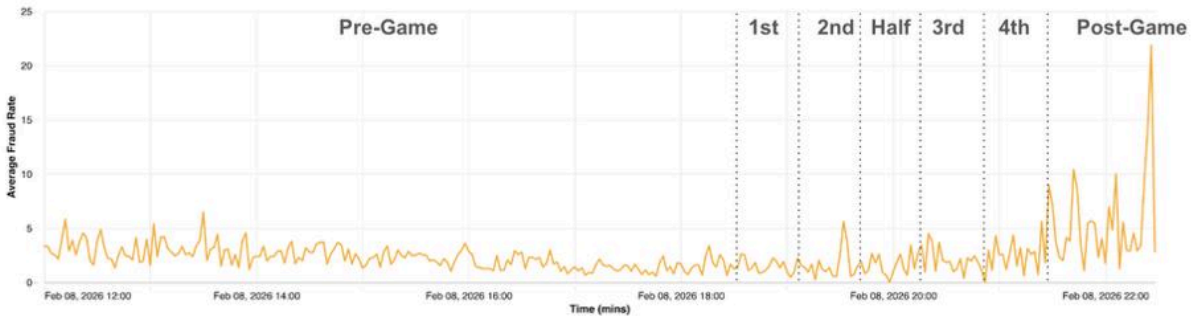

Reflects the average Identity Fraud rate witnessed during the game time that was analyzed. The average rate increased from 2.5% pre-game to 4.7% post game. It is important to note that during the game and all the way through post-game analysis the volumes dropped substantially, which helps to reflect the high percentage of those applicants who were attempting to gain access to an online gaming account using someone else's identity.

Synthetic Fraud Rates

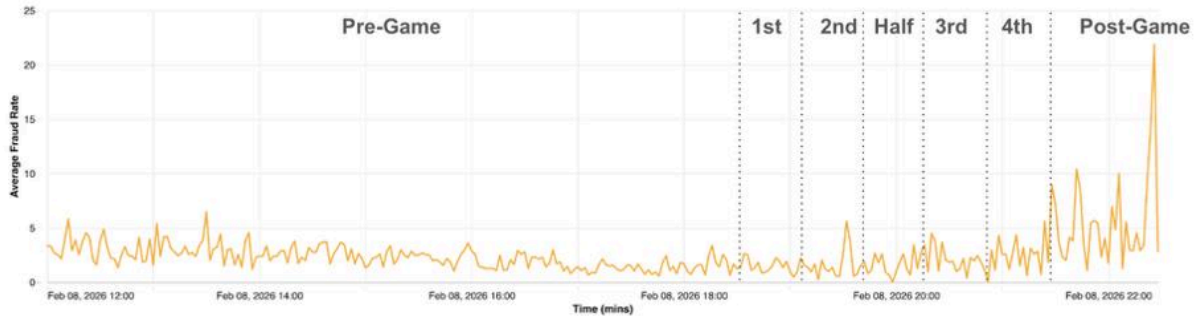

Reflects the average Synthetic Identity Fraud rate witnessed during the game time that was analyzed. The average rate increased from 2.1% pre-game to 6.0% post game. It is important to note that during the game and all the way through post-game analysis the volumes dropped substantially, which helps to reflect the high percentage of those applicants who were attempting to gain access to an online gaming account using someone else's identity.

Common characteristics for applications that scored above the threshold included:

- Use of stolen — but structurally accurate — PII such as name, DOB, SSN and address
- Manipulated contact layers including new emails and substituted phones
- Newly registered or obscure domains used for email provisioning
- Proxy or VPN routing and IP to address distance anomalies
- Automation assisted account creation

While the overall fraud attack rate across all players on Super Bowl Sunday was low, there were specific Socure customers that faced attempted fraud attack rates as high as 40% for periods of time pre- and post-game. We believe this is more of a reflection of which fraud ring is focused on what sports betting or predictive market provider. The fraud ring that hit this customer used bots and automation to speed rates, and likely had a higher amount of uniquely developed and accurate identities that could have been saved for attacks specifically on Super Bowl Sunday.

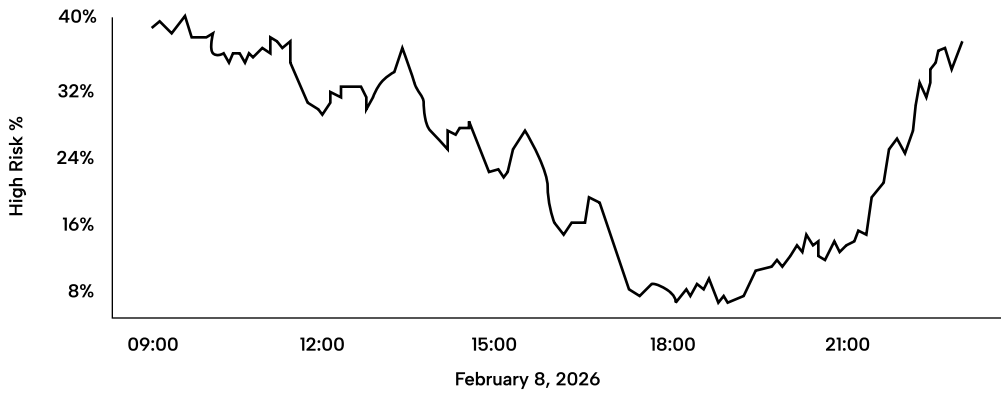

Graph reflects fraud attack rate for a single Socure customer who faced the largest fraud attack.

We also identified what appear to be international attacks targeting two specific participants in the analysis. The attacks were distinct and followed completely different patterns:

- One ring relied almost exclusively on newly created Gmail accounts, all having similar user name patterns.
- The other used newly created email addresses tied to .us domains purchased from a small set of fraud-leaning domain registrars.

These cases illustrate how identity fraud and synthetic identity fraud are closely linked. Stolen identities often serve as the raw material. Over time, that data is blended, reused, and modified to create synthetic personas that do not belong to a real person but can be durable enough to pass automated checks.

Suggested Next Steps

1. While it would be easy to predict that identity theft will remain the dominant attack vector in 2027, that is not guaranteed. We recommend maintaining identity fraud and synthetic fraud models if you are using them today (and strongly considering adding them if you are not). **These controls are not only important for reducing financial loss.** We observed underage applicants using what appears to be a parent's identity to establish accounts. Younger users are technically savvy and share tactics quickly within their friends' group. We expect identity theft, familiar fraud, and synthetic identity strategies to increase in underage attempts. These models, combined with document verification and selfie validation, materially reduce that risk.
2. One operator experienced a more concentrated fraud attack than others. Because they process through RiskOS®, we were able to work together with the customer to implement short term strategies focused directly on the behavioral patterns of that attack. This limited additional exposure. RiskOS, provided at no cost to Sigma and RiskScores customers, enables rapid deployment of temporary controls to curb fraud spikes, particularly when patterns remain consistent within a single attacking entity.
3. We expect higher fraud rates, as well as date-of-birth mismatches, following halftime and into the post game window to persist in 2027. Historically, legitimate volume declines after halftime, increasing relative fraud concentration. Customers of Sigma Identity and Sigma Synthetic could take advantage of the lower volumes and reduce score thresholds to capture a higher percent of those likely fraudulent applications following halftime through post-game.
4. The use of Socure's Device Intelligence and robust device and location signals produced are strong additions to our Sigma models and RiskScores and can be used to further hone fraud strategies to focus on increased fraud detect rates, lower friction and reduced false positives, or both.

Coordinated Bonus Abuse Rings

Super Bowl promotional incentives created ideal conditions for organized bonus abuse. To execute at scale, these actors must first successfully establish accounts through identity theft, synthetic fraud, or first party fraud, often supported by mule networks.

One customer alerted us to a large fraud ring attempting coordinated bonus abuse. The actors were using stolen consumer identities to gain account access. Where successful, the accounts completed required bets or predictions to qualify for bonuses. Once credited, the promotional funds were withdrawn to bank accounts, where we assume those fraudulently obtained funds were quickly moved elsewhere.

Tactics included:

- Automated account creation at scale, likely bot driven
- Automated, volume engineered trades to unlock rewards
- Synchronized withdrawals once incentives were credited
- Convergence of payouts to a limited number of endpoints
- Concentration in newly registered .us domains
- Heavy use of VPN services, specifically Norton VPN
- Structured and repeatable identity construction patterns

These behaviors were repeated so consistently in this customer's application flow that adding a single targeted rule materially increased fraud detection at the same false positive rate.

A second customer saw an attempted attack using recently-developed and unique gmail emails. Prior research conducted by Socure shows that organized fraud rings are becoming more sophisticated in suppressing linkages between identity elements such as email to phone and identity to address, reflecting an increasingly sophisticated approach to identity creation.

While this pattern is concerning in the long run, today, the newly created identity elements score high in the Sigma models, and Socure's email and phone RiskScores and Correlation values trigger additional risk of the identity.

Suggested Next Steps

1. Leverage Sigma and RiskScores to identify applicants likely to engage in bonus abuse. More importantly, use onboarding signals and the scores themselves within transaction monitoring. Early lifecycle signals, especially for marginal populations, add meaningful precision to account level models. Industry estimates suggest up to a 25% increase in fraud capture among new accounts and more than a 70% reduction in coordinated fraud ring impact when upfront signals are integrated into transaction monitoring.
2. Use Socure's Account Intelligence, which validates the ownership of the bank account and also assesses if the account is open and active. This is increasingly important to identify bonus abuse attacks, where fraudulent players often send promotional dollars to bank accounts where the owner of the bank account does not correlate to the identity that was used to open an account.
3. Use RiskOS to enable immediate and temporary fraud strategy changes that can be removed once the threat subsides.

Age Verification and Underage Access Attempts

Super Bowl Sunday revealed a significant surge in the volume of DOB mismatches during new account creation.

DOB mismatch rates increased from 12% pre-game to 26% post-game, a 244% surge from the beginning of the game through post-game. Application volumes increased 300% in the 24-hour period leading into the game. This likely represents tens of thousands of underage individuals attempting to gain access during the event window.

Date of Birth Mismatch Trend

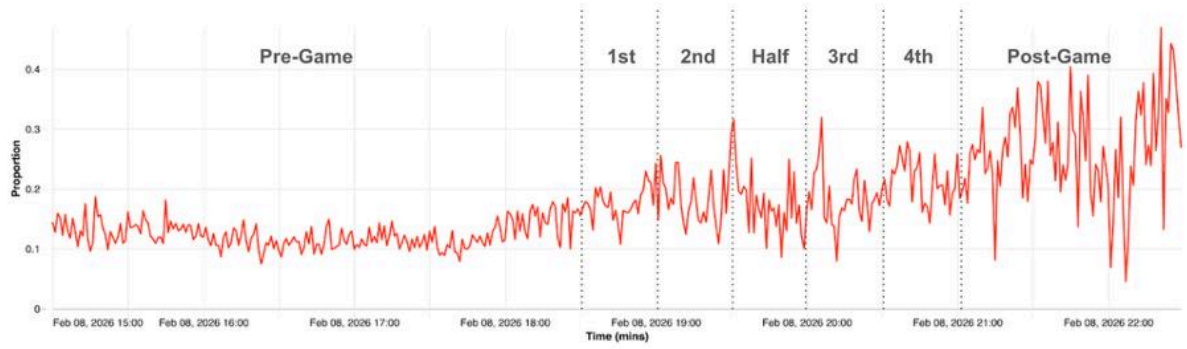

Reflects the total percentage of applications with a date of birth (DOB) asserted on the application that did not match the DOB found in Socure's consortium and/or public records. The average rate increased from 12% pre-game to 26% post game. As you can see, as the game started and went through post-game timing, we saw an increase in the number of applicants not providing their true DOB. It is important to note that during the game and all the way through post-game analysis the volumes dropped substantially, which helps to reflect the high percentage of those applicants who were trying to gain access to gaming accounts using a DOB that is not their own.

It is important to restate that the rate of DOB mismatch was lower than normal during pre-game, where there was a heavy concentration of advertisements from these industries. We believe this creates some form of proof that online betting and predictive markets that ran commercials during that time were not targeted at the underaged population.

These DOB manipulation behaviors fall into two primary categories:

1. Age washing, where applicants slightly adjust their DOB to meet eligibility thresholds

2. Familiar fraud, where older friends or family members' verified identity details are used

At identity verification step-ups, manipulation and attempts to overcome age validation persisted. We observed a 300% increase in applicants whose asserted DOB at ID check did not match extracted document data or corroborated device signals.

Effective controls for age assessment included:

- Use of a passive fraud scoring and KYC (Verify) identity verification solutions
- Device and digital footprint analysis
- Phone tenure and email maturity checks (to catch those underage applicants using their parents form of ID
- Automated escalation to government ID verification AND live selfie

The Super Bowl demonstrated that while minors are increasingly digitally sophisticated, layered identity systems can detect subtle inconsistencies across PII, device, and behavioral signals. This is particularly relevant as age gating becomes a broader regulatory focus beyond gaming.

Document Validation and Age Mismatches: A second bite at the apple

High-volume events like the Super Bowl compress onboarding timelines and increase pressure to approve accounts quickly. While identity graphing and digital intelligence provide strong passive signals, document validation remains a critical escalation tool when risk thresholds are exceeded.

- Government ID verification identifies ID theft and synthetic attempts
- Deepfake detection and spoof resistance spots AI-generated efforts
- Document to device correlation helps to assess fraud risk and assesses age
- Real-time decision automation through RiskOS allows for quick fraud strategy changes to impact fraud ring spikes

During the Super Bowl surge across our online gaming clients, we observed multiple instances of underage individuals attempting to bypass age restrictions by misusing close relatives' identities (our assumption).. In several cases, these applicants had already been flagged and escalated to document validation and, in some instances, live selfie assessment. What we learned is simple: Underage applicants are persistent and will attempt multiple times to overcome age controls.

In a few cases, minors used a parent or sibling's ID while submitting their own selfie. This clearly illustrates the importance of live selfie controls in detecting repeat behavior across platforms.

We also observed several attempts to manipulate the selfie verification process itself by taking a photograph of the picture on a driver's license and submit that picture as the selfie.

Minors also use pictures on their phone to represent a selfie of a drivers license from a parent or sibling.

In other cases, minors used physical photos to attempt bypass.

These patterns reinforce two themes:

- Underage users are actively experimenting with creative workarounds during high traffic events.
- Liveness detection, combined with selfie verification, is critical to preventing image reuse and impersonation, particularly during event-driven spikes like the Super Bowl.

Overall, this activity reflects coordinated and repeated attempts rather than isolated mistakes. It reinforces the need for layered identity verification controls across gaming platforms to support effective age assessment and compliance.

State Mismatch: Crossing State Lines Without Moving

As kickoff approached and the game moved into halftime, a second pattern became more pronounced. We saw an increase in the rate of state mismatches during new account creation, where applicants claimed residency in a state that conflicted with independent identity, device, and network signals indicating they were located elsewhere. The mismatch rate was substantially lower pre-game at 3.1%, but more than doubled post-game to 7.9%.

State Mismatch Trend

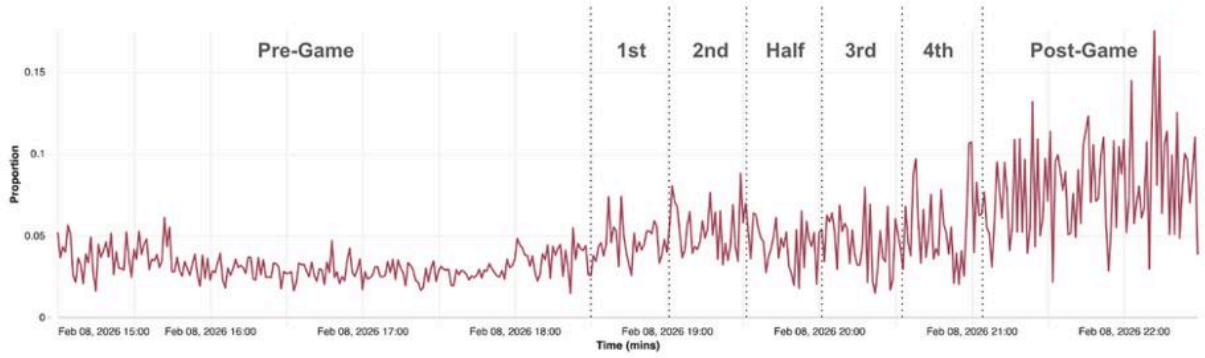

Reflects the total percentage of applicants whose physical address and state asserted on their application were not where they lived. The average rate jumped from 3.1% pre-game to 7.9% post game. Similar to Figure 2 above, it is important to note that during the game and all the way through post-game analysis the volumes dropped substantially, which helps to reflect the high percentage of those applicants who were likely in a state that does not allow online gaming.

This behavior aligns with what is referred to as jurisdictional evasion. Applicants attempt to bypass state-by-state gaming regulations by misrepresenting where they live, or physically are, at the time of account sign up. In practice, this often takes the form of geo-laundering, where users digitally reposition themselves across state lines without physically moving.

Recent regulatory disclosures and investigative reporting have highlighted this trend across the online betting ecosystem. Common tactics include the reuse of out-of-state addresses, borrowed identity details tied to permissive jurisdictions, and attempts to appear located within a legal betting state using IP proxies and VPNs, even when they are not.

We will drill further into this behavior further because it requires some fairly complex analysis to get down to solid numbers of people trying to actively fake their location to open accounts illegally. But the early takeaway is clear: Geography remains one of the most actively-tested boundaries in online gaming, and state mismatch is a critical first sign signal for understanding compliance risk during major events.

Suggested Next Steps

This next step is all Socure's. During our analysis, we identified how difficult it is to accurately measure how many occurrences of illegitimately hiding locations to overcome state-by-state voting regulations. We also could not see what the state-to-state patterns were, and are very interested to learn how to calculate these numbers accurately. For instance, certain address mismatch, IP proxy, and location information patterns could represent fence jumping by real consumers, legitimate travel and/or identity theft. Because of this, we need to drill down further to better identify and report back to you what we are seeing related to people trying to avoid state's regulations.

Key Takeaways

From where we sit, we see the Super Bowl as a successful endeavor for the online sports betting, predictive markets, the Seattle Seahawks, and Bad Bunny.

Good identities came in droves. Fraud, age discrepancy and state mismatch rates were historically low.

However, because of the massive volumes of “good” customers, our customers did also see a spike in fraud attempts, underage consumers trying to fraudulently gain access to accounts, and state-by-state matches that we do not fully understand yet.

Additionally, international fraud rings did attempt sizable efforts to open accounts for the use of promotional bonus abuse.

Our recommendations are included under each individual section. And by utilizing RiskOS, organizations are enabled to react immediately when necessary and gain a broader understanding of what is happening in real-time.

As a final reminder, Socure’s full-suite approach with our online gaming and predictive market customers addresses all of the issues seen during the Super Bowl with multiple layers, which include:

- Sigma Identity evaluates the cohesiveness and historical integrity of identity elements across name, DOB, SSN, address, email, phone, device, and behavioral signals.
- Email, Address and Phone RiskScores and Correlation values assess the risk of each individual identity element and creates additional signal that can be used to develop customer risk strategies
- Sigma Synthetic isolates fabricated or manipulated identities by identifying structural construction patterns and proof of life gaps across the Identity Graph.
- Digital Intelligence surfaces device level signals that expose emulation, automation/bot activity, and infrastructure reuse. Device is also an important linking variable.
- Identity Graph connects and ranks all the data points around a person’s PII and the devices they use into a single, persistent view of identity, which powers more accurate fraud detection, better customer recognition, and higher conversion.
- Account Intelligence and Sigma First Party Fraud extend protection to the payout moment, validating bank account ownership and identifying repeat bad actors who may be using their own real identities but have abusive intent.

We greatly appreciate the opportunity to serve you, and look forward to answering any questions related to our Super Bowl analysis, the recommendations and ways to optimize Socure’s full solution suite to limit compliance, fraud and first-party fraud risk.