



2026

Online Fraud Insights

Insurance Industry Report

Introduction

Fraud Never Takes a Year Off

Insurance fraud costs the U.S. economy more than \$300 billion a year¹, nearly seven times the \$45 billion in insured losses UCLA Anderson estimated for the 2025 Eaton and Palisades wildfires (one of the costliest disasters in American history).² While catastrophic events trigger sudden headlines, insurance fraud operates year-round, inflicting a steady toll on honest policyholders who ultimately absorb these expenses.

Fraud often leaves observable patterns, and earlier evidence can help carriers act before exposure grows. People who exaggerate or misrepresent a claim may continue living their real lives in public, across social platforms, court records, local news, and the digital trail many leave behind.

This report shows a year of Carpe's findings from our work surfacing evidence relevant to potentially fraudulent open injury claims across auto, workers' compensation, general liability, and disability lines throughout 2025. It's also our second annual online fraud insights report, and for the first time we can compare this year's findings to historical data. Each section shows which patterns are holding, which are shifting, and where the trend lines point next.

The 2025 record, at a glance:

Hundreds of millions in modeled value identified through contradictory alerts in 2025

- Of the evidence Carpe surfaced online, 86% came from social media and 14% came from non-social sources
- 36% of contradictory-alert referrals surfaced within the first two months after loss
- Auto carried the largest share of alerts; Disability had the highest referral rate among the three compared lines, at 7.2%

86% / 14%

Social / Non-Social

Share of online evidence by source type.

36%

In the first two months

Share of contradictory-alert referrals after loss.

60.9%

Auto Alert Share

Auto carried the largest share of alerts.

Model value is an estimate, actual realized savings may differ.

7.2%

Disability Referral Rate

Highest among Auto, Workers' Comp and Disability.

Source: Carpe 2025 Assessment data.

What Changed Since Last Year?

The Biggest Shifts In Carpe's 2025 Findings.

Auto remained the largest source of alerts in 2025, at 60.9% compared with 62.0% in 2024. Workers' Comp's share grew from 23.8% to 26.8%. Delaware, which did not rank near the top last year, is now near Nevada and South Carolina for the second-highest rate of alerts per resident.

Facebook still accounts for six flags out of ten on social media, with Instagram and TikTok growing at the edges. Timing held steady too; referrals still arrive on roughly the same curve.

Metric	2024 data	(this report)	Direction
		2025 data	
Modeled value identified	\$230.6M	\$263.8M	Up materially ▲
Alert share by line of business*	Auto, 62.0%; Workers' Comp, 23.8%	Auto, 60.9%; Workers' Comp, 26.8%	Auto -1.1 pp; Workers' Comp +3.0 pp
Top state, per-100k rate	Connecticut; NV, SC, GA close	Connecticut #1; Delaware new to the top tier	Reshuffle ◀▶
Facebook share of contradictory flags	≈62%	60.5%	Down modestly, still dominant ▶
Referrals in first two months	≈35%	35.7%	Stable ▶
Social vs. non-social evidence	82.7% / 17.3%	86% / 14%	Social gained ▲
Peak age cohort	35-44	35-44	▶

*Alert shares compare calendar-year 2024 with 2025. pp = percentage points. Shares describe the recorded alert mix. Prior-year figures marked ≈ are read from last year's published report rather than a matching data export and should be treated as indicative.

The Biggest Fraud Questions

Where, What, When, And Who.

Auto remained the largest source of alerts in 2025, at 60.9% compared with 62.0% in 2024. Workers' Comp's share grew from 23.8% to 26.8%. Delaware, which did not rank near the top last year, is now near Nevada and South Carolina for the second-highest rate of alerts per resident.

Facebook still accounts for six flags out of ten on social media, with Instagram and TikTok growing at the edges. Timing held steady too; referrals still arrive on roughly the same curve.



Where

Where does injury fraud strike the hardest?



What

Lines of business take the biggest hit, and **where** are these fraudulent alerts surfacing?



When

When is fraud found throughout the life of an open claim?



Who

Who is behind these fraudulent claims - are certain age groups more likely to play the system?

Sagmet shares reflect the assessed claims portfolio and available online evidence.

Source: Carpe 2025 Assessment data.

Where Contradictory Evidence Concentrates

Population Size Alone Does Not Explain Alert Concentration.


California, Florida, and Texas produce the most alerts for the same reason they produce the most claims: they have the most people. Measured as alerts per 100,000 residents, the 2025 leaders are Connecticut, Nevada, South Carolina, Delaware, and Georgia. Surprisingly, California and Texas don't even crack the top 20 when normalized for population.

Alert concentration extends well beyond the largest states and cities, with several mid-size states punching above their population weight class.

Top rankings for alerts per 100,000 residents shifted from 2025's report: Connecticut took the top spot by a wide margin, with Nevada, South Carolina, and Delaware in a virtual dead heat for second.

Delaware is the surprising newcomer, ranking nowhere near the top tier in last year's report, even normalized for population. One possible explanation is that alert concentration reflects a mix of local legal-market activity and economic pressure. The rankings do not establish causation, but shifts in the map may offer clues about the conditions surrounding potentially fraudulent claims.

Where Fraud Strikes Hardest



Alerts per 100,000 residents, 2025.

1	Connecticut	4.64
2	Nevada	3.32
3	South Carolina	3.30
4	Delaware	3.30
5	Georgia	2.91

Ranks 1-5

Ranks 6-10

Ranks 11-20

Ranks 21-50

Source: Carpe Claims monitoring data, 2025.

Where Claim Evidence Concentrates

A Closer Look At The Metros Where Contradictory Evidence Surfaced Most Often.

For the second consecutive year, Las Vegas leads the country in alert concentration, a striking result for a metro of its size. Los Angeles, San Antonio, Houston, and Jacksonville round out the top five, with Atlanta, Chicago, San Diego, Orlando, and Dallas completing the top ten. San Antonio and Jacksonville both climbed into the top five this year; the rest of the composition held steady.

Several conditions recur across these markets, including active personal-injury legal markets, tourism, and transient workforces. Carpe's data does not establish that those conditions caused the rankings.

Top Cities for Fraudulent Injuries. Metro-level fraud intensity, 2025 rankings.

01 Las Vegas

02 Los Angeles

03 San Antonio

04 Houston

05 Jacksonville

06 Atlanta

07 Chicago

08 San Diego

09 Orlando

10 Dallas

11 Brooklyn

12 Tampa

13 Philadelphia

14 Charlotte

15 Phoenix

16 Austin

17 Sacramento

18 Miami

19 Columbia

20 Buffalo


Source: Carpe Claims monitoring data, 2025.

Why Do Some Areas Run Hot?

1. **Strong Personal-Injury Legal Markets and Attorney-Driven Claims.** Areas with more personal-injury advertising may experience greater attorney-driven claim inflation and social inflation.
2. **Population Density and Organized Rings.** Dense urban zones may give coordinated schemes more room to operate unnoticed.
3. **Tourism and Transient Populations.** Visitors, seasonal workers, and short-term residents can make claims harder to verify.
4. **Economic Pressure.** Financial stress may increase the incentive to exaggerate.
5. **Legal Loopholes.** Gaps in insurance and workers' compensation statutes may create room to operate.
6. **No-Fault Collusion Dynamics.** No-fault frameworks may enable legal-medical referral mills.


These factors offer plausible explanations for why alerts cluster, though Carpe's data does not establish any one factor as the cause. Alert patterns vary by market, and continuous evidence gathering can help claims teams see what a static review may miss.

Referral Rate and Alert Share by Line of Business



Source: Carpe, 2024 published and 2025 supplied referral rates; 2025 alert distribution.

Reason-Code Shares Within Each Line, 2025



Source: Carpe, 2025 assessment data.

“Each line of business leaves its own evidence trail. In workers’ comp, it may be an undisclosed side business. Auto is more likely to surface through lifestyle, travel, or physical activity, while disability claims may reveal employment that never made it into the file. Knowing those patterns helps claims teams recognize relevant evidence sooner,”

Tom Rasmussen
VP of Claims Product for Carpe.

What Lines Take the Biggest Hit?

Disability Retained The Highest Referral Rate.

Disability continued to have the highest referral rate of the three lines, with 7.2% of monitored claims referred in 2025, followed by Workers’ Comp at 6.2% and Auto at 5.3%. Rates dipped across all three from 2024.

Referral workload also depends on the size of the book. A large Auto portfolio can generate more referrals than a smaller Disability portfolio, even at a lower rate. Looking at alert volume, Auto accounted for roughly six in ten of Carpe’s alerts in 2025, with Workers’ Comp contributing just over a quarter.

The evidence behind those alerts gives claims teams a useful point of comparison for their own caseloads. Work-related posts may raise questions about claimed lost wages, while physical activity may conflict with reported limitations. Carpe’s Reason Code breakdown shows which types of evidence prompted alerts in each line, with full definitions in the glossary.

Workers’ Comp: Association with a business drives 42.7% of WC flags, more than twice any other reason. A common Workers’ Comp contradiction is a claimant who reports being unable to work but appears to be working.

Evidence In The Wild: The Football Coach. An alert on a workers’ compensation claim showed the claimant coaching football while collecting benefits for a disabling injury that should have prevented him from performing duties as a coach. The finding prompted an online search and uncovered evidence that later anchored mediation, where the claim settled for one-third of the original demand. These findings resulted in a **\$177,000 reserve reduction**.

Auto: Lifestyle activity leads Auto alerts at 28.6%, with travel (24.7%) and physical activity (24.6%) close behind. Claimants continue to document daily life online. The average adult spent nearly two and a half hours a day on social media in 2025³, leaving ample opportunity for contradictory evidence to surface.

General Liability: Activity Beyond the Claim. Lifestyle activity accounts for 31.6% of General Liability alerts, followed by travel at 27.3% and physical activity at 22.3%. Together, these categories show how daily activity can create evidence relevant to an alleged injury or limitation.

Disability: The Long-Game Problem. Association with a business leads at 33.4%, followed by pain and illness posts at 24.6% and travel at 20.1%. Disability claims run long, which gives claimants more time to contradict themselves.

Disability Evidence In The Wild: A disability claimant disclosed employment history ending in 2010. Professional-profile evidence identified two additional employers spanning 2013 to the present. The evidence contributed to a **\$423,000 reserve release**. Employment representations are among the most testable claims a file contains, because careers leave public records.

Where Online Evidence Appears

Facebook Still Leads, But The Evidence Trail Keeps Widening.

Facebook remains the largest source of social-media evidence, although meaningful signals now appear across a wider group of platforms.


Facebook remains the evidence leader. It accounted for 60.5% of contradictory flags in 2025, barely off the roughly 62% in last year's report and still more than the next four platforms combined. Its durability makes sense: the cohorts filing the most claims are the cohorts that still live on Facebook, and both of those factors are slow to change.

Instagram accounted for 13.8% of flags, with visual posts continuing to surface relevant activity. The platform built for showing off remains an excellent place to show off exactly the types of activities an injury claim says cannot be performed.

LinkedIn generated 7.8% of all flags, concentrated almost entirely on associations with businesses. For workers' comp and disability claims, it may be the highest-signal platform per flag in the entire dataset. People curate LinkedIn for employers, not for their insurance carrier.

TikTok accounts for 3.2% of flags today. Its role may grow as younger cohorts age into peak claim-filing years. Fraud evidence turns up wherever people spend their time online. Facebook leads today and is likely to remain the largest source for years to come. But a program tuned only to Facebook leaves four flags in ten outside its view, and as other platforms grow in popularity they become more valuable over time.

Share of All Contradictory Flags by Source, 2025.



Source: Carpe, 2025 assessment data.

Social Media Is Only Part of the Picture

Fourteen Percent Of Evidence Came From Beyond Social Platforms.

Of the evidence Carpe surfaced online in 2025, 86% came from social media. The remaining 14% came from non-social sources: local news, court records, criminal and booking records, ecommerce listings, public databases, and thousands more. Non-social evidence includes some of the most decisive in the dataset, and a carrier that equates online evidence with social media is leaving the strongest material on the table.

Evidence In The Wild: The Booking Record. A booking-record alert revealed that a workers' compensation claimant had been incarcerated following a parole violation. The claim was closed, and the carrier recorded a **\$55,528 reserve reduction**. No social platform surfaced this evidence and the carrier confirmed it would not have known without a full online investigation.

Year over year, the split leaned further toward social, from 82.7% in last year's report to 85.6% this year. The direction may change as platforms and privacy habits evolve, but the operational lesson holds: effective detection has to follow evidence across the open web.

Evidence Found on Social Media vs. Non-Social Sources.



Source: Carpe, 2024 published report and 2025 source distribution.

When Is Contradictory Evidence Found?

The First Two Months Are Crucial, But Nearly Two-Thirds Of Evidence Is Discovered Later.


The largest concentration of referrals appears early. In 2025, 35.7% of referrals arrived within the first two months after loss, split almost evenly between month 0 (18.1%) and month 1 (17.6%). Contradictory evidence often appears quickly, often before the file has finished its initial review (or before the claimant hires an injury attorney).

Evidence continues to surface well past the first two months. Nearly two-thirds of contradictory evidence arrives after the window in which most claims receive whatever scrutiny they will ever get, with referrals documented all the way out to month 24. A one-time investigation forfeits the long-tail value of ongoing assessment, where long-term lines like disability and workers' comp surface most of their contradictions.

Assess Early And Often. Early assessment covers the biggest single window for contradictory evidence. Continued monitoring catches later contradictions that a one-time search cannot.

Early Referrals Lead; Evidence Continues To Surface

Share Of Referrals By Month After Loss, 2025.



Source: Carpe, 2025 assessment data.

What Happens When a Profile Goes Private?

A Privacy Change Can Become Part Of The Claim Timeline.

Most public-content referrals arrive early, but access can change later in the claim.

Carpe added the **Public Profile Made Private (PPMP)** reason code in 2025 to detect changes in access. Unlike content-based reason codes, PPMP records when a social profile moves from public to private during an open claim. This non-content indication can put the file back on the claims team's radar.

Across the initial PPMP data, the median last-confirmed-public observation occurred 273 days after loss, while the median first-confirmed-private observation occurred 346 days after loss. Together, those observations place the typical transition window roughly nine to eleven and a half months after loss. The exact change may have occurred at any point between the two observations.

That window falls well beyond the timeline for initial claim review, and a single point-in-time search cannot show that a profile later became private. Continuous assessment can surface the change while the claim is still active, giving the claims team a reason to review what was already captured and decide whether further investigation is warranted.


In 15.5% of PPMP alerts, another URL carried a different reason code, meaning the privacy change formed part of a broader evidentiary picture in this subset of cases. PPMP records behavior rather than intent, and claims teams determine what, if anything, the change means for the file.

Which Signals Appear by Age?

How Alert Patterns Vary Across Age Groups.

The peak alert cohort in 2025 was ages 35 to 44, carrying 25.2% of alerts, with the distribution falling away on both sides: 45 to 54 at 22.2%, 55 to 64 at 18.7%, 25 to 34 at 18.3%, 65+ at 10.4%, and 18 to 24 at 5.1%. The shape of the curve matches last year's report: the working-age middle of the claims population produces the most evidence.

Who Raises the Red Flags?



Source: Carpe, 2025 assessment data.

What These Signals Tell Us

Physical-activity flags peak in

the 18-to-24 cohort and decline with every decade of age.

Younger people tend to exercise, play sports, and post about those activities. Instagram and TikTok surface a larger share of evidence among the youngest cohorts.

Work-related evidence is most common among mid-career claimants.

Association-with-a-business flags peak in these cohorts, consistent with side businesses and undisclosed employment. The gap between claimed work status and actual work status is one of the most common contradictions in the dataset.

Older claimants surface different signals entirely.

Financial-concern flags skew toward older cohorts, and possible-fatality flags concentrate heavily among claimants 65 and older, a signal with obvious relevance for long-duration disability claims. Facebook's share of evidence rises steadily with age and dominates the oldest cohorts.

Who's Not Being Honest?

How Different Generations Raise Different Red Flags.

Fraud Found by Age and Domain

	18–24	25–34	35–44	45–54	55–64	65+
Facebook	24.5%	52.3%	68.4%	77.8%	83.8%	86.0%
Instagram	49.5%	27.5%	17.5%	11.3%	7.3%	5.1%
LinkedIn	10.5%	13.1%	9.5%	7.8%	7.6%	8.4%
TikTok	14.1%	6.6%	4.2%	2.9%	1.2%	0.3%
X	1.5%	0.5%	0.5%	0.3%	0.1%	0.1%

Source: Carpe, 2025 social-platform flags for records with known age.

Reason Fraud is Flagged by Age

	18–24	25–34	35–44	45–54	55–64	65+
Physical activity	31.1%	19.1%	18.3%	15.8%	13.5%	10.3%
Lifestyle activity	15.9%	16.6%	20.5%	24.6%	24.9%	21.7%
Travel	16.9%	15.0%	15.2%	17.8%	20.4%	18.1%
Business association	15.2%	25.0%	21.2%	19.2%	18.9%	14.6%
Pain / illness / medical	1.8%	3.2%	4.8%	6.6%	8.2%	8.7%
Financial concern	2.0%	2.2%	2.6%	2.4%	2.3%	2.5%
Possible fatality	2.1%	0.8%	0.7%	1.1%	1.8%	13.1%
All remaining codes	15.1%	18.1%	16.7%	12.5%	10.0%	11.0%

Source: Carpe, 2025 assessment data.

Predictions for 2027 and Beyond

Where The Trend Lines Point Next.

Taken together, this year's findings point to five shifts carriers should prepare for next:

1. Facebook Holds While the Edges Grow

Facebook's share held steady, but Instagram and TikTok are gaining with younger claimants. As those cohorts age into the peak 35-to-44 alert band, evidence will keep spreading across platforms.

2. Auto May Remain the Largest Source of Alerts

Auto accounted for 60.9% of alerts in 2025. Whether that share persists will depend on the mix of claims assessed and how referral rates evolve. Its evidence spans lifestyle, travel, and physical activity, rewarding broad detection over narrow searches.

3. Fraud Detection Moves Upstream

Carriers will increasingly use AI to identify which claims need attention earlier, combining internal risk signals with external evidence before exposure compounds. The opportunity is faster evidence review and prioritization, with claims professionals retaining the decision.

4. Evidence Availability Becomes a Lifecycle Issue

Referrals continued through month 24, and the median last-public and first-private observations occurred 273 and 346 days after loss, respectively. Long-duration claims will require attention well beyond initial review.

5. Fraud Gets Easier to Create and Harder to Verify

Classic exaggeration will remain, while generative AI makes fabricated images and documents cheaper and more convincing. Synthetic content raises the risk of both deception and false positives, so tools such as Carpe Vision will play a growing role in helping claims teams identify AI-generated or altered imagery. Provenance, timestamps, original-source links, and corroboration across independent sources will become more important.

"Online fraud signals are increasingly spread across multiple platforms as well as across time, and they're becoming harder to verify as synthetic content improves.

Carriers will need to combine early risk signals with continuous, source-level evidence across every line of business."

Tom Rasmussen
VP of Claims Product for Carpe.

Methodology

This report is built on Carpe's continuous assessment of open injury claims across auto, workers' compensation, general liability, and disability lines during calendar-year 2025, combined with qualitative insight from Carpe's claims product team. We segmented the data across multiple dimensions, including age group, social media platform, state, city, line of business, reason code, and time from first notice of loss.

How To Read The Findings

These findings describe claims assessed by Carpe in 2025 and may not represent the broader U.S. claims population. Results reflect both claimant behavior and evidence available from public online sources during the assessment period. Segment and geographic comparisons are descriptive; differences in claim mix and small subgroup sizes may influence the results. Because raw counts are withheld to protect carrier confidentiality, findings for smaller segments should be interpreted directionally. Segment shares do not measure the probability of fraud within each line of business, age group, or platform because assessed-claim exposure differs across segments.

State-level figures are expressed as alerts per 100,000 residents to reduce the effect of population size. They are not claim-level fraud rates and may also reflect differences in Carpe's assessed-claim and carrier mix by state. City figures are published as simple rankings and should be interpreted as alert concentration within Carpe's 2025 data.

How referral rates and alert shares differ. Referral rate is the number of referred claims divided by the number of monitored claims within the same line of business and year, expressed as a percentage. Its numerator is referred claims, not alert counts. Alert share is a line's alert count divided by the total alert count across all lines. The year-over-year referral-rate comparison covers Auto, Workers' Comp, and Disability, for which both years' rates are available. Referral rates do not establish confirmed fraud prevalence.

What counts as a contradictory alert. An alert is generated when posted content contradicts the claimed injury, limitation, or work status, or surfaces material information about the claim. Every alert is tied to its original source. The evidence is stored, timestamped, and delivered with the reasoning behind the alert.

How modeled value is calculated. Modeled-value figures estimate the potential loss-cost impact of contradictory alerts. The model uses carrier-reported outcomes where available and applies standard reserve-impact assumptions when carriers do not provide final disposition data. Because carriers do not consistently report outcomes, these figures represent modeled estimates. Actual realized savings may differ.

Track record. Since inception, Carpe has reviewed millions of claims and identified nearly \$1 billion in total value, with about 5% of all claims reviewed revealing evidence of fraud that had not previously been found.

Notes & Sources.

1. Coalition Against Insurance Fraud, "Who Me?" Who Commits Insurance Fraud and Why, produced in partnership with Verisk, https://helpstopfraud.org/wp-content/uploads/2024/01/WHO_ME_STUDY_REPORT.pdf.

2. Zhiyun Li and William Yu, "Economic Impact of the Los Angeles Wildfires," UCLA Anderson Forecast, updated March 3, 2025, <https://www.anderson.ucla.edu/about/centers/ucla-anderson-forecast/economic-impact-los-angeles-wildfires>.

3. Naveen Kumar, "Average Time Spent on Social Media Per Day (2026 Stats)," DemandSage, June 11, 2026, <https://www.demandsage.com/average-time-spent-on-social-media/>.

Glossary of Reason Codes

Association with a Business.

For workers' comp and disability claims, the subject is associated with a business not in the carrier's file. For other lines, the subject's employment is physically demanding.

Travel

The subject posted about traveling more than 160 miles from the provided address.

Lifestyle Activity

Evidence of the injured party returning to regular, elective activity inconsistent with the claim.

Physical Activity.

The subject was found being physically active.

Pain, Illness, or Other Medical Issues.

The subject posted about being in pain, medical procedures, or medical check-ins.

Potentially Relevant Information.

The subject is associated with a potentially relevant event, such as an additional accident or property information.

Financial Concern

Financial troubles such as bankruptcy or personal fundraising, or the sale of high-value items.

Potentially Unlawful Activity

The subject was found to have a criminal background or an association with potentially unlawful activity.

Information Related to the Claim

Content published directly about the claim, such as mentions of filing or managing it, or how the incident happened.

Possible Fatality

Evidence suggesting the subject has passed away.

Public Profile Made Private (PPMP)

A monitored social profile changed from public to private during the life of an open claim.

Large Group Exposure

Evidence indicating the subject's claim connects to a broader group event or mass exposure relevant to the claim.



The Shortest Path From Data to Decision.

Carpe gives insurers the shortest path from data to decision, transforming online data and public records into decision-ready intelligence for fraud detection, claims management, and underwriting.

[Learn more at carpe.io](https://carpe.io)