FrankonFraud.com

How the CFPB Complaint Database Was Overrun by Credit Washers and Cookie Cutter Complaints

A forensic analysis of 3,837,184 published consumer narratives, 2015 to 2026. Every narrative compared against every other one.

> *Three of every five narratives the bureau took in during 2025 matched another complaint word for word, and one in four narratives it has ever published came from just 204 letters.*

Frank McKenna | FrankonFraud Research | August 15, 2026

Results were prepared and written with Claude Fable 5, an AI model from Anthropic.

The finding in brief

The Consumer Financial Protection Bureau built its public complaint database so consumers could describe financial harm in their own words. This study compared every published narrative against every other one, 3,837,184 complaints received between March 2015 and July 2026, and found the record no longer works that way.

43 percent of all published narratives are word for word copies of at least one other complaint, and in 2025 the annual figure reached 59.3 percent. The most conservative tier finds 1,328,263 complaints in 6,683 template families that each appear in at least 3 states on at least 3 separate dates, and just 204 letters supply nearly one in four narratives the bureau has ever published.

The machinery is visible inside the text itself: software that rewrote one letter nine thousand ways, merge fields that broke and printed their own labels, chatbot replies pasted in whole, and a wave of identity theft block demands that grew from 52 narratives in 2015 to 140,101 in 2025. Template complaints closed with corrective relief at more than twice the rate of ordinary complaints, so the strategy pays, and the paper trail runs to specific ZIP codes, specific sellers, and prices as low as five dollars.

The letter that signed itself in

In July 2024, a complaint asked Experian to fix a credit report and closed with the words "Sincerely, Add your name here." The filer pasted a template into a federal complaint form without reading the instruction, then typed a real name directly after it (complaint 9574263).

That small accident opens a large question. Regulators use this database to spot patterns, companies use it to measure harm, researchers and journalists treat it as a public record of consumer experience, and courts and Congress cite it as evidence, so it matters whether its narratives are what they claim to be.

This report answers the question with text forensics at full scale. It measures the duplication, reconstructs the template families behind it, times their growth, tests the odds of innocent coincidence, maps the mills to the ZIP code, traces the letters across the open web, and measures whether the strategy works.

Two terms recur throughout. A cookie cutter complaint is a complaint built from shared source text rather than original writing, and credit washing is the practice of disputing accurate negative credit information as identity theft or error until it is deleted, with templated complaints as its delivery system.

Data and methods

The source is the full public export of the CFPB Consumer Complaint Database with complaints received through July 27, 2026, a snapshot of 17,119,590 complaints. Of those, 3,837,184 carry a published consumer narrative, and because the bureau collects publication consent only on its web form, every published narrative in the database is a web submission.

Each narrative was normalized before comparison: lowercased, whitespace collapsed, and the CFPB's redaction masks, dates, dollar amounts, and digit runs replaced with fixed tokens. Wording, spelling, grammar, and sentence order stayed untouched, so two narratives match only when the consumers used the same words.

Duplication was measured at two strengths. Exact groups require identical normalized text, and template families connect complaints whose narratives match at 82 percent character similarity or closer, computed with minhash candidate generation followed by direct verification of 738,823 candidate pairs.

Family analysis excludes narratives under 200 normalized characters, since short notes repeat innocently, leaving 3,389,410 comparable narratives. Three limits apply throughout: a shared template proves shared source text rather than a false claim, distinct date and state pairs serve as a conservative floor for separate filings, and 2026 figures are partial because narratives publish on a delay.

The scale of the duplication

Measure	Complaints	Share
Complaints in the database snapshot	17,119,590	
Complaints with a published narrative	3,837,184	22.4% of complaints
Narratives long enough to compare, 200 or more characters	3,389,410	88.3% of narratives
Exact copy of at least one other complaint	1,649,243	43.0% of narratives
Inside a template family at 82 percent or closer	2,057,386	60.7% of comparable
Inside a strong family, at least 3 states and 3 dates	1,328,263	39.2% of comparable
Inside a family of 1,000 or more complaints	886,847	23.1% of narratives

The headline number sits in the fourth row: 1,649,243 published narratives are exact copies of at least one other complaint. Loosen the match to 82 percent similarity and 2,057,386 complaints trace back to shared source text.

The strong family tier is the conservative core, letters that appear in at least 3 states on at least 3 dates, a footprint no individual produces by refiling. 90 percent of those complaints name Equifax, Experian, or TransUnion, against 40 percent for everything else, because the bureaus are where deletion lives.

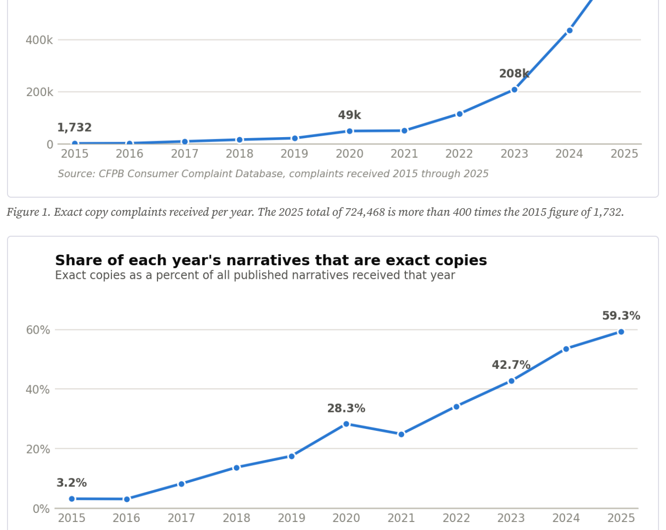

Figure 1. Exact copy complaints received per year. The 2025 total of 724,468 is more than 400 times the 2015 figure of 1,732.

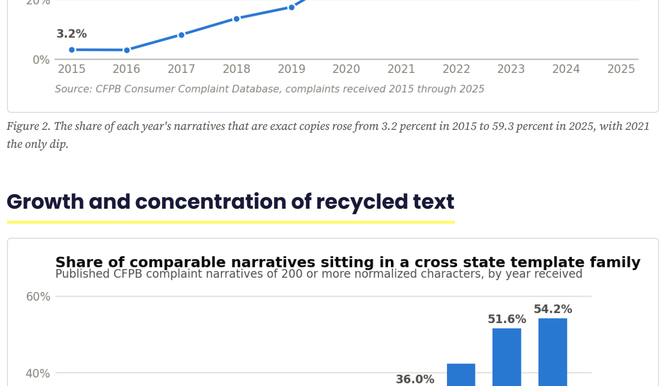

Figure 2. The share of each year's narratives that are exact copies rose from 3.2 percent in 2015 to 59.3 percent in 2025, with 2021 the only dip.

Growth and concentration of recycled text

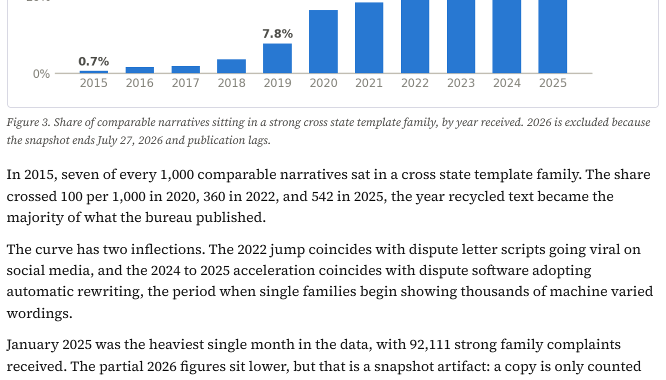

Figure 3. Share of comparable narratives sitting in a strong cross state template family, by year received. 2026 is excluded because the snapshot ends July 27, 2026 and publication lags.

In 2015, seven of every 1,000 comparable narratives sat in a cross state template family. The share crossed 100 per 1,000 in 2020, 360 in 2022, and 542 in 2025, the year recycled text became the majority of what the bureau published.

The curve has two inflections. The 2022 jump coincides with dispute letter scripts going viral on social media, and the 2024 to 2025 acceleration coincides with dispute software adopting automatic rewriting, the period when single families begin showing thousands of machine varied wordings.

January 2025 was the heaviest single month in the data, with 92,111 strong family complaints received. The partial 2026 figures sit lower, but that is a snapshot artifact: a copy is only counted once its twin is also published, and recent months lag.

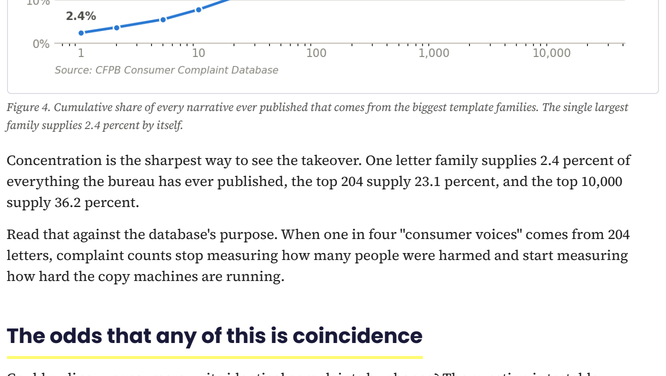

Figure 4. Cumulative share of every narrative ever published that comes from the biggest template families. The single largest family supplies 2.4 percent by itself.

Concentration is the sharpest way to see the takeover. One letter family supplies 2.4 percent of everything the bureau has ever published, the top 204 supply 23.1 percent, and the top 10,000 supply 36.2 percent.

Read that against the database's purpose. When one in four "consumer voices" comes from 204 letters, complaint counts stop measuring how many people were harmed and start measuring how hard the copy machines are running.

The odds that any of this is coincidence

Could ordinary consumers write identical complaints by chance? The question is testable, because written language is predictable, and its predictability can be measured from this corpus itself.

A character model trained on 150,000 distinct narratives needed 2.44 bits of information per character on held out text. At that measured rate, the odds of a second consumer independently reproducing an average length narrative of 963 normalized characters are about 1 in 10 to the 707th power.

Handing the coincidence theory every possible break changes nothing. At a floor of 1 bit per character, below what Shannon measured for humans predicting English, the average narrative still comes out at 1 in 10 to the 290th power, roughly the odds of winning the Powerball jackpot 34 times in a row.

Narrative	Copies	Odds of one chance match, measured	Odds at the most generous floor
Privacy letter citing 15 USC 1681, 843 characters	30,110	1 in 10^586	1 in 10^254
"My credit reports are inaccurate," 268 characters	24,246	1 in 10^188	1 in 10^81
Unauthorized accounts letter, 500 characters	18,436	1 in 10^371	1 in 10^150
Average published narrative, 963 characters		1 in 10^707	1 in 10^290

For scale, the observable universe contains about 10 to the 80th atoms, so even the floor odds for the 268 character letter resemble two strangers picking the same single atom out of all of them. The database's 7.4 trillion possible narrative pairs give chance about 10 to the 13th tries, a budget that only produces identical pairs on texts shorter than about 43 characters.

Every letter in the top 150 is longer than that, most by hundreds of characters. Chance is not a candidate explanation for anything in this report.

The template economy

Rank	Complaints	Active dates	What the letter says	Sample ID
1	93,287	Apr 2022 to Jul 2026	Privacy letter reciting 15 USC 1681 sections 602 and 604, demands deletion	5430101
2	48,786	Feb 2022 to Jul 2026	Same privacy letter with the account list typed inline	5238646
3	25,909	Jan 2025 to Jun 2026	Cash App complaint citing the CFPA and the CFPB action against Block	11557688
4	24,558	Oct 2021 to Jul 2026	"My credit reports are inaccurate," with the misspelling "inquires"	4772310
5	20,113	Mar 2022 to May 2026	"I'm really not sure what happened," plus "There is no third party involved"	5300347
6	19,601	Jun 2023 to Jun 2026	"The banking system is dependent upon fair and accurate credit reporting"	7173688
7	19,543	Apr 2025 to Jul 2026	"You have reported inaccurate and unauthorized accounts"	12870952
8	19,310	Dec 2023 to Jan 2026	"My account was never late," blames a billing system error	8033808
9	14,717	May 2024 to Jul 2026	Banking system letter again, in a quoted variant	9021528
10	13,972	Aug 2023 to May 2026	Identity theft letter promising an attached Word document	7407332

The top family shows what modern filing machinery looks like. Its 93,287 complaints arrive in 32,617 slightly different wordings, the signature of mail merge software wrapping a fixed skeleton around each customer's account list, and one copy opens "Hi I am submitting this CLIENT NAME without any influence and this is not a third party" (complaint 7201493), the vendor's merge field left unfilled.

Shared mistakes travel with the letters and work like fingerprints. Family four repeats the misspelling "accounts and inquires" across 24,558 complaints, and other families carry "Its possible that there is an error," "pressume" for presume, and "me personally who is filling this complaint," none of which appears anywhere on the open web.

Some machinery is cruder. One identity theft letter opening "This is my third endeavor to tell you" appears 8,397 times in wording scrambled by a synonym spinner, which turned plain English into "the records I am questioning connect with no exchanges acquiring any possession of goods" (complaint 5613551).

And some of it writes prose. The strangest letter in the database was filed 5,097 times word for word, and it reads like the first page of a novel:

> *In the quiet humdrum of everyday life, a sudden jolt shattered the tranquility : I had become a victim of identity theft. Like a thief in the night, unseen hands had stolen my personal information, plunging me into a bewildering maze of uncertainty. With each fraudulent transaction and unauthorized account, the breach of trust cut deep, leaving me grappling with the unsettling reality of a violated identity. Yet, in the midst of chaos, a glimmer of resilience emerged, a resolve to reclaim what was rightfully mine and restore the harmony stolen by unseen adversaries. These are the following accounts need to be removed :*
>
> Complaint 8445776, filed 5,097 times word for word, February 2024 through December 2025

Look at the seam. Four sentences of AI novelist prose collapse into "These are the following accounts need to be removed :", the broken grammar of the merge point where software pastes in each customer's account list.

The letter's numbers match its machinery: 55 percent of copies came from Texas, its heaviest ZIPs sit in the Houston suburbs, its copies arrive almost exactly 3.1 to a filing session, one consumer to all three bureaus at once, and 60.1 percent of them closed with relief. It appears nowhere on the open web.

The TikTok letters

Two of the largest templates trace to the social media letter economy, and together they show both of its modes: the slow viral burn and the overnight spike.

The slow burn is the number one letter in the database, here in full:

> *In accordance with the Fair Credit Reporting act. The List of accounts below has violated my federally protected consumer rights to privacy and confidentiality under 15 USC 1681. 15 U.S.C 1681 section 602 A. States I have the right to privacy. 15 U.S.C 1681 Section 604 A Section 2 : It also states a consumer reporting agency can not furnish a account without my written instructions 15 U.S.C 1681c. (a) (5) Section States : no consumer reporting agency may make any consumer report containing any of the following items of information Any other adverse item of information, other than records of convictions of crimes which antedates the report by more than seven years. 15 U.S.C. 1681s-2 (A) (1) A person shall not furnish any information relating to a consumer to any consumer reporting agency if the person knows or has reasonable cause to believe that the information is inaccurate.*
>
> Complaint 6584982, filed 30,110 times word for word, February 2023 through June 2026, with "furnish a account" frozen into every copy

This letter is taught across TikTok dispute videos and sold in packs on Gumroad and Etsy for 5 to 15 dollars, with a mainstream forms retailer carrying a version as well. It was filed from 52 states and territories, aimed 99.9 percent at consumer reporting agencies, and won: 61.1 percent of its copies closed with non monetary relief, triple the baseline.

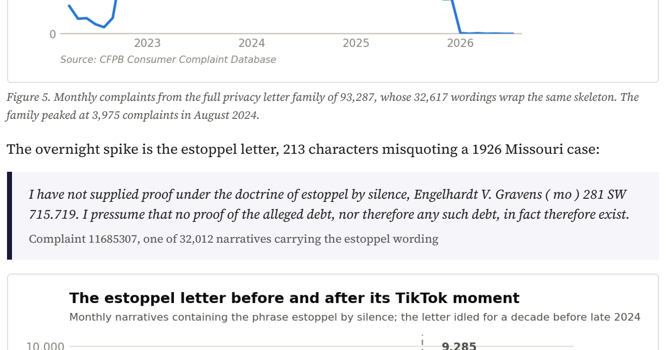

Figure 5. Monthly complaints from the full privacy letter family of 93,287, whose 32,617 wordings wrap the same skeleton. The family peaked at 3,975 complaints in August 2024.

The overnight spike is the estoppel letter, 213 characters misquoting a 1926 Missouri case:

> *I have not supplied proof under the doctrine of estoppel by silence, Engelhardt V. Gravens (mo) 281 SW 715.719. I pressume that no proof of the alleged debt, nor therefore any such debt, in fact therefore exist.*
>
> Complaint 11685307, one of 32,012 narratives carrying the estoppel wording

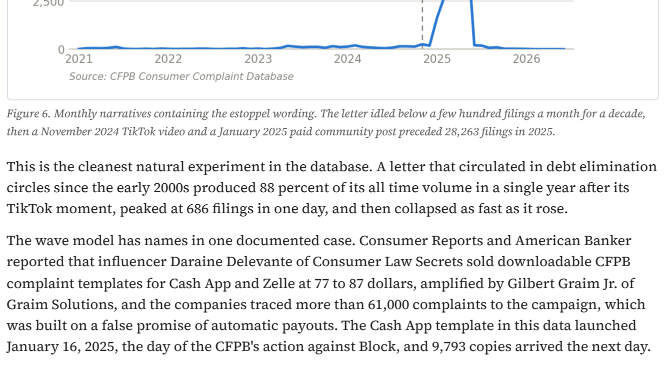

Figure 6. Monthly narratives containing the estoppel wording. The letter idled below a few hundred filings a month for a decade, then a November 2024 TikTok video and a January 2025 paid community post preceded 28,263 filings in 2025.

This is the cleanest natural experiment in the database. A letter that circulated in debt elimination circles since the early 2000s produced 88 percent of its all time volume in a single year after its TikTok moment, peaked at 686 filings in one day, and then collapsed as fast as it rose.

The wave model has names in one documented case. Consumer Reports and American Banker reported that influencer Daraine Delevante of Consumer Law Secrets sold downloadable CFPB complaint templates for Cash App and Zelle at 77 to 87 dollars, amplified by Gilbert Graim Jr. of Graim Solutions, and the companies traced more than 61,000 complaints to the campaign, which was built on a false promise of automatic payouts. The Cash App template in this data launched January 16, 2025, the day of the CFPB's action against Block, and 9,793 copies arrived the next day.

Credit washing and the 605B machine

Section 605B of the Fair Credit Reporting Act requires a bureau to block reported information within four business days once a consumer submits an identity theft report. It is the strongest tool in the statute, and it is the mechanism credit washing abuses, because a blocked account vanishes whether or not the theft was real.

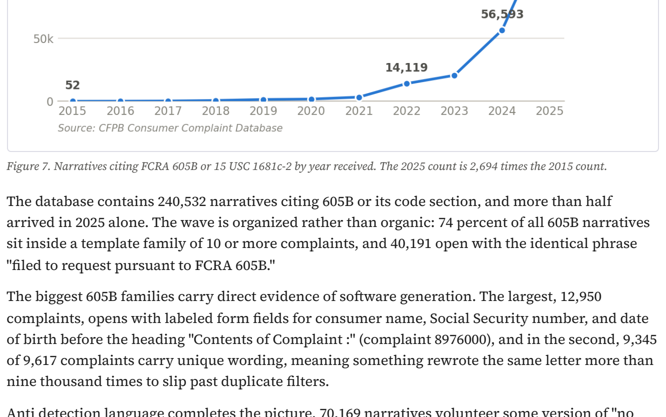

Figure 7. Narratives citing FCRA 605B or 15 USC 1681c-2 by year received. The 2025 count is 2,694 times the 2015 count.

The database contains 240,532 narratives citing 605B or its code section, and more than half arrived in 2025 alone. The wave is organized rather than organic: 74 percent of all 605B narratives sit inside a template family of 10 or more complaints, and 40,191 open with the identical phrase "filed to request pursuant to FCRA 605B."

The biggest 605B families carry direct evidence of software generation. The largest, 12,950 complaints, opens with labeled form fields for consumer name, Social Security number, and date of birth before the heading "Contents of Complaint :" (complaint 8976000), and in the second, 9,345 of 9,617 complaints carry unique wording, meaning something rewrote the same letter more than nine thousand times to slip past duplicate filters.

Anti detection language completes the picture. 70,169 narratives volunteer some version of "no third party" involvement, 106 repeat "This is from me and not a credit repair company" word for word, and 16,581 call a bureau's credit repair suspicion a "stall tactic," recycling a reply letter published by Strike Debt in 2012 and credited to Carreon and Associates.

Does the strategy work? The outcome data says yes.

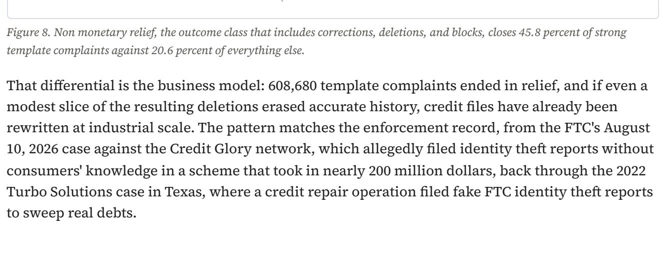

Figure 8. Non monetary relief, the outcome class that includes corrections, deletions, and blocks, closes 45.8 percent of strong template complaints against 20.6 percent of everything else.

That differential is the business model: 608,680 template complaints ended in relief, and if even a modest slice of the resulting deletions erased accurate history, credit files have already been rewritten at industrial scale. The pattern matches the enforcement record, from the FTC's August 10, 2026 case against the Credit Glory network, which allegedly filed identity theft reports without consumers' knowledge in a scheme that took in nearly 200 million dollars, back through the 2022 Turbo Solutions case in Texas, where a credit repair operation filed fake FTC identity theft reports to sweep real debts.

Machine fingerprints inside the narratives

The strongest evidence of automation is the machines telling on themselves. Twenty four complaints across ten states contain the string "Dispute2025" fused into the letter, as in "thatDispute2025 $ you -- credit reporting agency" (complaint 24506748), and the artifact appears inside two different letter templates between October 2025 and July 2026, so one generation system carrying a label named Dispute2025 produced both.

AI assistants leak into the record as well. 103 narratives keep the chatbot's opening line, among them an Illinois complaint beginning "Sure, I can help you with that. Here 's a sample credit report complaint" (6816070) and a Florida complaint reading "ChatGPT said : ChatGPT Certainly! Here 's a draft letter you can use" (11625721).

Other filings kept the blank spots: "[Your Name] [Your Address] [City, State, ZIP Code]", "[Insert Date]", and three identical complaints closing "Sincerely, Add your name here." followed by a real name. One letter filed 909 times even carries its own confession, "Please review this letter meticulously as it is submitted on my behalf," third party drafting language inside a first person complaint channel.

The geography of the mills

Matching 92.6 percent of all complaints to a full 5-digit ZIP and dividing by population turns the database into a map. Nationally the record works out to 5.48 complaints per 100 residents since 2011, 1.23 narratives per 100, 0.53 exact copies per 100, and 0.43 strong template complaints per 100.

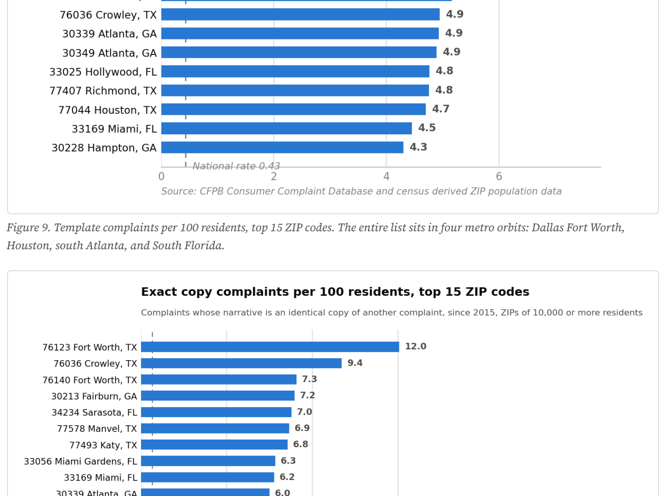

Figure 9. Template complaints per 100 residents, top 15 ZIP codes. The entire list sits in four metro orbits: Dallas Fort Worth, Houston, south Atlanta, and South Florida.

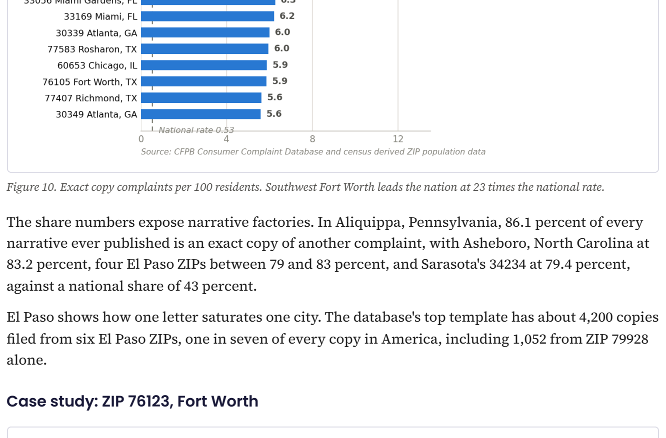

Figure 10. Exact copy complaints per 100 residents. Southwest Fort Worth leads the nation at 23 times the national rate.

The share numbers expose narrative factories. In Aliquippa, Pennsylvania, 86.1 percent of every narrative ever published is an exact copy of another complaint, with Asheboro, North Carolina at 83.2 percent, four El Paso ZIPs between 79 and 83 percent, and Sarasota's 34234 at 79.4 percent, against a national share of 43 percent.

El Paso shows how one letter saturates one city. The database's top template has about 4,200 copies filed from six El Paso ZIPs, one in seven of every copy in America, including 1,052 from ZIP 79928 alone.

Case study: ZIP 76123, Fort Worth

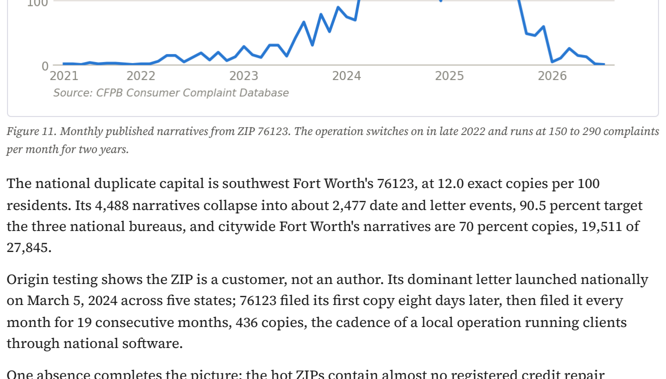

Figure 11. Monthly published narratives from ZIP 76123. The operation switches on in late 2022 and runs at 150 to 290 complaints per month for two years.

The national duplicate capital is southwest Fort Worth's 76123, at 12.0 exact copies per 100 residents. Its 4,488 narratives collapse into about 2,477 date and letter events, 90.5 percent target the three national bureaus, and citywide Fort Worth's narratives are 70 percent copies, 19,511 of 27,845.

Origin testing shows the ZIP is a customer, not an author. Its dominant letter launched nationally on March 5, 2024 across five states; 76123 filed its first copy eight days later, then filed it every month for 19 consecutive months, 436 copies, the cadence of a local operation running clients through national software.

One absence completes the picture: the hot ZIPs contain almost no registered credit repair storefronts, and Texas requires credit services organizations to register with the Secretary of State. The filer profile is the modern home based operator, licensed dispute software, marketed through social media, invisible to directories, a model with local precedent in United States v. RMCN Credit Services, the McKinney, Texas company that paid a 2.35 million dollar judgment in 2014 after sending more than a million templated dispute letters containing fabricated statements.

Where the letters come from

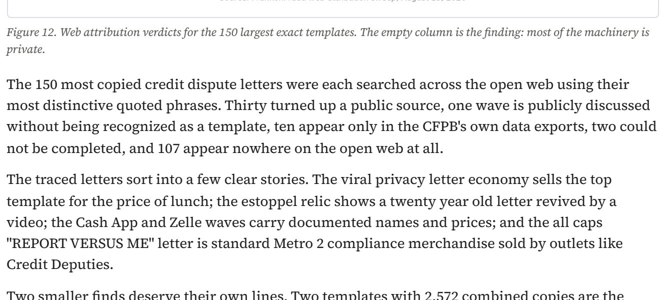

Figure 12. Web attribution verdicts for the 150 largest exact templates. The empty column is the finding: most of the machinery is private.

The 150 most copied credit dispute letters were each searched across the open web using their most distinctive quoted phrases. Thirty turned up a public source, one wave is publicly discussed without being recognized as a template, ten appear only in the CFPB's own data exports, two could not be completed, and 107 appear nowhere on the open web at all.

The traced letters sort into a few clear stories. The viral privacy letter economy sells the top template for the price of lunch; the estoppel relic shows a twenty year old letter revived by a video; the Cash App and Zelle waves carry documented names and prices; and the all caps "REPORT VERSUS ME" letter is standard Metro 2 compliance merchandise sold by outlets like Credit Deputies.

Two smaller finds deserve their own lines. Two templates with 2,572 combined copies are the Federal Trade Commission's own sample dispute letter with one word changed, "dispute" becoming "delete," and a credit repair software vendor, Client Dispute Manager Software, ranks first on exact quoted searches for three late payment templates totaling about 4,600 complaints, a lead one manual check short of confirmed.

The empty column may matter most. Letters filed tens of thousands of times with zero public footprint were never meant to be seen, which points to private infrastructure: dispute software letter libraries, client only downloads, and AI generation inside tools, exactly what the machine varied wordings predict.

What the findings mean

For the database, this is a data integrity problem of the first order. A majority of what the bureau now publishes as consumer voice is recycled text, and any analysis that counts complaints as independent observations, from academic research to company scorecards to regulatory attention, is counting the output of copy machines.

For the credit system, the outcome differential is the alarm. Template campaigns close with corrective relief at more than twice the baseline rate, so the complaint channel functions as a deletion engine, and every wrongly deleted account degrades the accuracy of the credit files lenders price against.

For consumers, the picture is honestly double sided. Real identity theft victims and real dispute rights live inside this same channel, many filers are themselves customers of services that promised legal deletion, and the FCRA protects the right to dispute either way, which is why this report measures text sharing rather than accusing any individual filer.

The bureau has started to respond. In June 2026 the CFPB rebuilt its complaint portal with identity verification, a truthfulness attestation, and a waiting period, citing abuse by credit repair firms and influencers filing mass generated disputes after complaint volume hit 6.6 million in 2025.

What should happen next

For the CFPB: flag narrative reuse across filers in the published data, publish template family identifiers alongside complaints, and use submission metadata, the sessions, addresses, and timing this public export cannot see, to separate consumer filings from bulk channels.

For the bureaus and furnishers: dispute logs already record submission channels, and matching them against the template families and ZIP concentrations in this report, starting with the Fort Worth corridor and the Dispute2025 engine, would convert statistical patterns into named operators.

For investigators and researchers: the highest value threads are the 605B mail merge families, the FTC letter mutation, the software vendor lead, and interviews with filers in the factory ZIPs, since the people who paid for these letters know who sold them. Every figure in this report can be reproduced from the complaint IDs in the companion workbooks.

Conclusion

The CFPB complaint database was built to let consumers speak and, for its first years, it mostly did. The record now shows something else: 43 percent of its narratives are copies, 59 percent of a recent year's intake matched another complaint word for word, one in four published narratives comes from 204 letters, and the fastest growing genre is the identity theft block demand that credit washing requires.

The evidence of machinery is not inferential; it is written into the complaints themselves, in merge fields, chatbot preambles, unfilled brackets, and a letter that admits it was submitted on someone's behalf. The database remains a rich record, but of a different thing than intended: it is now the best public measurement of the credit repair industry's output, and it should be read, published, and policed as such.

Methods and reproducibility

Pipeline: streaming extraction of 17,119,590 rows, narrative normalization, exact grouping by hashed text, minhash candidate generation with direct verification at the 82 percent threshold, cohesion checks on the largest families, ZIP matching against a census derived population table, and a 150 template web attribution sweep run August 15, 2026.

Companion workbooks: CFPB Template Inventory, Exact Match Odds, ZIP Hotspot Analysis, Duplicate Narrative Hotspots, Template Attribution Search volumes 1 and 2, Fort Worth Credit Repair Landscape, and the memo CFPB Narrative Template Forensic Review.

Sources

CFPB Consumer Complaint Database, full public export, complaints received through July 27, 2026

Federal Trade Commission press release, August 10, 2026, action against the Credit Glory network

Consumer Reports, September 11, 2025, and American Banker, October 7, 2025, on templated Zelle and Cash App complaints

American Banker, June 2026, on the CFPB complaint portal overhaul

United States v. RMCN Credit Services, Department of Justice, 2014; United States v. Turbo Solutions Inc., Department of Justice, 2022

Consumer Federation of America, September 9, 2025, on Navy Federal overdraft complaints

Strike Debt, The Debt Resisters Operations Manual, Appendix A, 2012, sample letters credited to Carreon and Associates

Federal Trade Commission sample letter for disputing credit report errors, updated April 2024

Population data: census derived ZIP table distributed with the uszipcode open data project