EPISODE 20 | Case File

BNPL Fraud: How Crime Rings Get Caught

MAY 26, 2026

This transcript was auto-generated and may contain errors or inaccuracies.

Jeff Welcome to "What the F Happened? Fraud and Financial Crime Deconstructed," a DEFEND podcast where we break down what's actually happening across fraud scams, AML and financial crime.

Emily Each episode cuts through the noise to explain the tactics, trends, and real world impact behind the headlines so you're better prepared for what comes next. Let's get into it.

Jeff Picture this you have been eyeing a high end espresso machine for weeks.

Emily Oh, I know that feeling, right?

Jeff It is a little speed to drop the cash all at once, but you finally get to the checkout page. And there it is. The buy now, pay later button.

Emily The famous BNPL option.

Jeff Exactly. You click it, the screen loads, and in that tiny fraction of a second before the little green check mark appears, you probably assume the only thing happening is like a quick ping to your bank.

Emily Yeah, that's what most people think. Just checking if you have funds.

Jeff But you'd be totally wrong because whenever you click that button, an absolutely massive, highly sophisticated war against organized cyber criminals is taking place. And the battlefield exists entirely in the milliseconds of that loading screen.

Emily It is wild to think about. Welcome to today's deep Dive. I am your expert co-host, ready to get into the weeds of this high speed war.

Jeff And I'm your host. So our mission for this deep dive is to understand a very specific, notoriously difficult engineering problem. Basically, how to balance a completely seamless, zero friction customer experience with rigorous ironclad fraud prevention.

Emily Yeah. And for a long time, the industry considered this an impossible trade off. To really grasp the engineering challenge, we first have to look at the economic reality of the BNPL industry, right?

Jeff Because consumer adoption has just skyrocketed.

Emily It has. And whenever transaction velocity increases, organized fraud rings naturally follow the money. I mean, the core challenge for any provider is stopping these highly organized, well-funded rings without adding security friction.

Jeff Because friction kills the sale.

Emily Exactly. If a legitimate shopper has to jump through a dozen multi factor authentication hoops, or upload a photo of their driver's license for a fifty dollars purchase, they don't just get annoyed.

Jeff No, they abandon the shopping cart entirely.

Emily Yep. They just leave.

Jeff Which introduces a metric from the sources that perfectly illustrates how bad this problem can get. It is called the hurt ratio.

Emily The hurt ratio is arguably the most critical metric for a consumer facing fintech company. Mechanically, it's just a measure of false positive.

Jeff Okay, let's unpack this for a second. So it tells the business exactly how many non-fraudulent events are being flagged.

Emily Right.

Emily Meaning perfectly good credit applications from real people that are being delayed or rejected for every one actual fraudulent application the system successfully catches.

Jeff And the numbers in one of these case studies were absolutely staggering. A major fintech client was experiencing a hurt ratio of five to one five to one.

Emily That is rough.

Jeff So to catch a single fraudster, the system was punishing five completely legitimate customers. Think about that from a business operations standpoint. It is like a nightclub bouncer who kicks out five perfectly good paying patrons who are already inside the club just to catch one person at the door with a fake ID.

Emily That is a perfect analogy. You simply cannot run a profitable consumer business with that kind of collateral damage.

Jeff You'd go out of business.

Emily You would. Especially in an industry where switching costs are essentially zero. Consumers have a dozen different payment options at checkout. If a good borrower faces an unexpected delay, they just close the tab and use a different provider.

Jeff And boom, you lose that customers lifetime value instantly.

Emily Exactly. But if you try to fix the hurt ratio by simply loosening your security rules to let more people through, well, you inevitably let the fraudsters in.

Jeff Which you also can't afford.

Emily Right?

Emily And in the BNPL model, missing the fraud results in first payment defaults. The provider fronts the money to the merchant immediately. So, default means the provider absorbs the total loss. It's an operational catch twenty two.

Jeff And the adversaries on the other side of this equation are incredibly sophisticated. We aren't talking about, you know, a teenager guessing passwords in their basement.

Emily Not at all. These are highly organized operations.

Jeff The sources broke down the specific distribution of attacks that this one provider was facing. And I think it's vital we look at the exact breakdown.

Emily Yeah, the data is very revealing. According to the case studies, fifty six point nine percent of the bad applications hitting this provider were categorized as pure fraud.

Jeff Pure fraud.

Emily Yes. And in this context, pure fraud is largely driven by mass registered synthetic identities.

Jeff Okay. And just to be clear, for everyone, a synthetic identity isn't just a stolen credit card, right?

Emily Right. It is way more complex.

Jeff It is when a fraudster stitches together real data like a legitimate social security number, often belonging to a child or an elderly person with a fake name, a fake address, and a burner phone number.

Emily Exactly. They create a ghost persona that doesn't actually exist in the physical world solely to request small to medium sized loans. They will never pay back.

Jeff That is wild.

Emily It is. And the ghost persona has zero credit history, which makes them very hard to evaluate. Following that, twenty six point three percent of the attacks were categorized as abuse.

Jeff Like exploiting promotional offers.

Emily Precisely creating hundreds of accounts just to harvest a ten dollars off your first purchase coupon.

Jeff Wow. Okay.

Emily And finally, sixteen point nine percent were account takeovers where hackers used credential stuffing to break into a legitimate users existing trusted account to make fraudulent purchases.

Jeff So the defense system have to look for identity fabrication, policy abuse, and account hijacking simultaneously.

Emily Yes, all at once.

Jeff And it has to do it in roughly two hundred milliseconds.

Emily During that tiny loading screen.

Jeff But this brings up a major question for me. These BNPL companies are massive tech enterprises. They aren't running their security on a spreadsheet.

Emily Definitely not.

Jeff The sources explicitly note they already had rules engines in place. They had supervised machine learning models, so I had to ask if they already had AI. Why was it failing?

Emily It's a great.

Jeff Question. I mean, supervised machine learning is specifically designed to recognize patterns. You know, if a synthetic ID behaves in a certain way, why could their existing AI just catch it?

Emily What's fascinating here is that the failure actually comes down to the fundamental architecture of supervised machine learning.

Jeff Oh really? How so?

Emily Well, a supervised model is entirely dependent on historical labeled data. During the training phase, data scientists have to feed the AI thousands of examples of past transactions that human reviewers have explicitly labeled as fraudulent.

Jeff Wait, so the AI is basically just memorizing what yesterday's crime looked like?

Emily Exactly. It constructs its mathematical boundaries based entirely on the past. It effectively says, I will block any application that shares the precise statistical characteristics of the fraud we caught last month.

Jeff But organized fraud rings know exactly how these models work, don't they?

Emily Oh they do. They continuously probe a provider's defenses once they figure out what specific behavior triggers a block from the supervised AI. They simply alter their tactics are.

Jeff So they change their IP routing or alter their device signatures, or just use different email domains.

Emily Yep. And because those new tactics haven't been manually labeled as fraud yet, the supervised AI just waves them right through.

Jeff It's completely blind to a novel attack.

Emily It is entirely blind. It's waiting for a label that doesn't exist yet.

Jeff Okay. But the sources also mention they tried to solve this by deploying clustering based approaches. And this is where I need some pushback.

Emily Sure, let's hear it.

Jeff Because if the whole problem with supervised learning is that it needs pre-existing labels, isn't clustering the exact solution.

Emily It seems like it would be.

Jeff Right. From a data science perspective, clustering just groups anomalies together based on similar behavior, regardless of what the label is. Why didn't that work?

Emily That is the logical next step, and it's exactly why the provider tried it. The theory behind clustering using algorithms like K-means or DBScan, is that fraudsters will naturally act similarly to one another, so they should form a distinct cluster in the data space.

Jeff Which makes sense.

Emily It does. The problem isn't the theory. The problem is the computational reality of doing it in real time.

Jeff Because of the checkout window.

Emily Right? Traditional clustering algorithms are incredibly computationally expensive. To group users accurately. The system mathematically has to calculate the distance between the new users behavior and the behavior of every other user currently in the system.

Jeff Wow. Okay. And when you are processing thousands of transactions per second, doing those mathematical comparisons across massive data sets creates latency, massive latency.

Emily It slows down the checkout. On top of that, maintaining and tuning these clustering models to avoid huge false positive spikes requires an army of data scientists constantly tweaking the parameters.

Jeff So it was adding huge operational drag precisely when the business needed to be fast and nimble.

Emily Exactly. They were stuck.

Jeff Supervised AI is too slow to adapt, and traditional clustering is too slow to compute. So if you can't rely on past data labels to catch a brand new scam, and you can't run heavy clustering algorithms at checkout, you have to completely change the underlying architecture.

Emily You really do. You have to change how the AI processes information at a structural level.

Jeff And that brings us to the core unsupervised machine learning, or UML.

Emily The shift to unsupervised machine learning is what finally breaks the deadlock. Unlike the supervised models we discussed, UML evaluates multiple dimensions of loan applications in real time to detect novel attacks as they happen.

Jeff Just evaluating the data on the fly.

Emily Yes, it identifies hidden correlations and synchronized behaviors across millions of events without needing a human to tell it what to look for beforehand.

Jeff But the way they execute this is what really caught my attention. Specifically regarding how they handle the data inputs. The sources highlight something called their feature platform.

Emily The feature platform is a game changer.

Jeff Typically, if a company wants to build a machine learning model, they have to do what's called feature engineering. A human data scientist has to sit down and figure out what specific variables or features the AI should even look at.

Emily Which is incredibly tedious.

Jeff They have to manually write code that says, hey AI, check the time difference between when the account was created and when the purchase was made, or check if this IP address matches the shipping zip code.

Emily And that manual feature engineering process takes months, it requires massive engineering resources just to define the rules of the game before the AI can even start playing right.

Jeff But automates this entire process, ingests raw data and automatically generates thousands of complex, multi-dimensional features without human intervention.

Emily The business impact of that is just massive.

Jeff It really is. The sources explicitly point out that this automation allowed the client to bypass a development cycle that typically takes comparable fintech companies up to five years to build in-house.

Emily Five years. Bypassing a five year engineering bottleneck is the difference between leading a market and going bankrupt.

Jeff Truly.

Emily Because fraud rings, do not wait five years for your data science team to finish building a model by automating feature generation, the UML engine instantly has thousands of microscopic behavioral data points to analyze.

Jeff It's just ready to go.

Emily Yep. But the true power of this system is how it organizes and cross-references those features using a knowledge graph.

Jeff Okay, here's where it gets really interesting. Let's really dive into the knowledge graph because this is where the architecture completely outclasses traditional databases.

Emily It's a fascinating structural shift.

Jeff If you think about a standard relational database, like a giant spreadsheet, looking up connections is really slow. If you want to know if two users share the same device, the system has to scan the user table, cross-referenced it with a device table, and then pull the result.

Emily Which again, takes too much computing time when you only have milliseconds to approve a loan.

Jeff Exactly. But a knowledge graph fundamentally changes how the data is stored. Instead of thinking of it like a movie detective's corkboard with red string connecting suspects. Think of it more like dropping a highly concentrated dye into a municipal water supply.

Emily Oh, I like that visual.

Jeff The graph database stores the relationships between data points natively as connections or edges, so when a new transaction hits the system, the AI drops that digital dye into the graph.

Emily And it spreads instantly.

Jeff Instantly, the dye flows through the preexisting pipes. It immediately illuminates every single shared device, every shared IP subnet, every associated email domain, and every money flow.

Emily There is no searching required. The connections are inherently part of the structure.

Jeff That is so cool.

Emily That is a phenomenal way to visualize it. The graph naturally maps out the invisible network in milliseconds, and because the UML algorithm is running on top of that graph, it just looks for dense pockets of die.

Jeff Areas where things are highly connected in ways that normal consumers simply aren't.

Emily Exactly. And the sources provided some incredible specific examples of what this graph illuminated during a holiday search.

Jeff Oh yeah, the cluster sizes were unbelievable.

Emily They really were. When the BNPL provider experienced a massive spike in legitimate loan applications during promotional events, the fraud rings tried to hide in the noise.

Jeff Classic tactic.

Emily Right? But UML detected coordinated clusters operating simultaneously. And we aren't talking about small groups. The graph mapped out clusters ranging from fewer than ten to as many as ten thousand users, operating as a single synchronized ring.

Jeff Wait, really? Ten thousand synthetic users applying for loans at the exact same time.

Emily At the exact same time.

Jeff And what fascinates me is how the graph caught them, because these rings were sophisticated enough to beat the supervised AI, but they got lazy with their data generation.

Emily They usually do eventually.

Jeff The knowledge graph instantly illuminated that these ten thousand distinct users were sharing the same small pool of Social Security Numbers across multiple applications. Yep, they were routing thousands of applications through the exact same physical shipping addresses.

Emily And my personal favorite detail from the case study. The unsupervised algorithm flagged massive clusters where the exact same first name was being used across thousands of fraudulent loans, even when the last names and addresses varied.

Jeff They literally copy pasted a generic first name ten thousand times.

Emily They really did.

Jeff The old supervised AI missed it entirely because no data scientist had ever thought to create a rule that says block the transaction. If the name John appears five hundred times in ten minutes.

Emily Because why would you even think to write that rule?

Jeff Exactly. It's incredible that a highly funded global fraud ring got dismantled just because the graph database realized the data was too mathematically similar to be human.

Emily It perfectly illustrates why relying on past labels is a losing game. Fraudsters always make a mistake eventually, and unsupervised graph analysis mathematically exposes that mistake in real time.

Jeff Okay, but spotting a ten thousand person fraud ring is only step one. The next operational hurdle is execution, right?

Emily Turning insights into action.

Jeff Because if the system suddenly flags ten thousand applications during a Black Friday sale, what is the business actually do? If you just dump ten thousand flagged applications onto a human manual review team, the whole department shuts down oh instantly.

Emily It would take them weeks to review that queue.

Jeff And meanwhile, the real customers mixed in there are waiting for their espresso machines.

Emily Exactly. Identifying the anomaly is useless if your infrastructure can't act on it safely. That is why the solution detailed in the sources wasn't just deploying UML in a vacuum. DataVisor integrated the unsupervised machine learning directly with a sophisticated rules engine and a unified case management toolset.

Jeff And the mechanics of that rules engine are brilliant. The sources highlight that they utilize the capability to back test and forward test their rules.

Emily Which is crucial for not breaking things.

Jeff Let me explain what that means practically. When the UML identifies a new pattern like our ten thousand John's, the risk team wants to create a hard rule to block it.

Emily Makes sense.

Jeff But before they push that rule live and risk breaking the checkout flow, they back test it. They run the simulation, applying that new rule against the last thirty days of historical data. They can instantly see if this rule had been live last month. Would it have blocked real customers?

Emily Right? And forward testing allows them to run the rule in a shadow mode against live incoming traffic, without actually blocking the transactions. They get to observe its accuracy in real time before flipping the kill switch.

Jeff That is so smart.

Emily It ensures they don't accidentally spike their own hurt ratio. But the true operational leap comes from how the case management system integrates with the knowledge graph. If we connect this to the bigger picture, it uses a feature called One Click Investigation.

Jeff This is where the efficiency goes through the roof.

Emily It really does. Once the knowledge graph maps out that network of bad actors, you know, our pocket of concentrated dye in the water supply, it remembers those connections. Okay. So if a brand new user applies for a loan tomorrow and the graph instantly sees that their device ID or their IP subnet is intimately connected to that known ten thousand person fraud ring, the system can automatically flag them.

Jeff Wow. So the manual review team doesn't have to start an investigation from scratch.

Emily Exactly. They click once, see the entire visual network of how this new user is connected to known fraud and can ban them instantly.

Jeff And the reverse is true, which is just as important.

Emily Oh, absolutely.

Jeff If a legitimate customer accidentally triggers a minor security alert, maybe they are traveling and using a weird IP address. The reviewer can click once. See that this user has absolutely zero connections to any known bad networks in the graph and instantly whitelist them.

Emily It removes the friction for the good guys.

Jeff So let's look at the actual business results here, because deploying graph databases and unsupervised learning sounds incredibly complex. Did it actually solve the problem? What's the final scorecard?

Emily It was a massive success. First, the integration of this entire technological stack happened in just one to two weeks.

Jeff One to two weeks. That is practically light speed for enterprise level financial infrastructure. Usually, ripping out an old AI model takes a quarter of a year.

Emily At least, and the financial protection scaled immediately for one specific client highlighted in the case studies, the system prevented thirty six million dollars in potential fraud losses simply by making accurate real time decisions.

Jeff Thirty six million.

Emily Yeah. For another provider, it resulted in over fifteen million dollars in annualized savings.

Jeff Incredible. They detected over three hundred and twenty distinct coordinated fraud rings. The case studies noted that some of these individual attacks had. They slipped through, the supervised AI would have caused over eighty thousand dollars in losses per single attack.

Emily Just one coordinated strike by passing the checkout screen.

Jeff But the most vital metrics bring us all the way back to the core dilemma we started with. The trade off between security and customer experience. The system achieved a seventy four percent fraud recall rate, meaning it successfully identified and caught the vast majority of the actual fraud attempts.

Emily Right. But as we know, catching fraud is easy if you just block everybody.

Jeff Exactly. And that is why the next metric is the most important one in the entire study. Alongside that high recall rate, they achieved a forty one percent reduction in false positives.

Emily A forty one percent reduction. They effectively cured the hurt ratio.

Jeff By drastically dropping the number of false positives. They ensured that the nightclub bouncer stopped kicking out the paying customers.

Emily The legitimate shoppers got their instant green check mark, while the synthetic rings gave silently blocked.

Jeff And as a direct result of that accuracy, the internal review efficiency for the human risk teams improved by a massive five x to twenty x multiplier.

Emily Which is huge for scaling a business.

Jeff Because the UML and the knowledge graph were doing the heavy lifting of connecting the dots. The human analysts weren't wasting hours manually cross-referencing IP addresses. It allowed the BNPL providers to scale their transaction volume aggressively, without having to hire hundreds of new human reviewers to keep up.

Emily It was a complete paradigm shift from reactive to proactive architecture.

Jeff It really makes you evaluate the hidden infrastructure of your own digital life differently. I want you to think about this the very next time you are shopping online. The next time you effortlessly check out, click that button to split your payment and instantly get approved without any friction.

Emily You'll never look at a loading screen the same way.

Jeff Never remember that in the literal blink of an eye. Behind that split second loading animation in unsupervised AI is rapidly navigating a massive graph database. It is mathematically mapping out thousands of behavioral connections in real time, ensuring you aren't kept waiting while simultaneously slamming the door shut on highly funded global cyber criminals.

Emily It is a quiet, incredibly complex war, and it is being won in milliseconds.

Jeff It really is a remarkable feat of engineering.

Emily And it leaves us with an intriguing broader question. I mean, this raises an important question. If unsupervised machine learning and graph architecture can uncover invisible, non-labeled connections to dismantle a ten thousand person fraud ring simply by noticing a shared first name and IP address? How might the same unsupervised approach uncover hidden positive connections in other fields?

Jeff Oh, that's an interesting point. Like what kind of fields?

Emily Imagine applying this architecture to massive data sets in medical research or urban planning fields, where we are drowning in complex data, but we don't even know what labels or what cures we should actually be looking for yet. Unsupervised learning could find patterns we didn't even know existed.

Jeff Wow. Now that is a deep dive for another day. Thanks for joining us and we will catch you next time.

Jeff You've been listening to "What the F Happened? Fraud and Financial Crime Deconstructed," a DEFEND podcast by DataVisor. If you want to keep learning between episodes, check out DEFEND Training, a set of self-paced online courses for fraud and financial crime professionals, practical and built around real world scenarios.

Emily And you can earn CPE credits through the ACFE San Francisco Bay area chapter. You can find it at datavisor.com/defend-training. The link is in the description.

Jeff This episode's audio was generated using Google's Notebook LM based on expert analysis and trusted sources. Thanks for listening. We'll see you next time.