moneyhub

Liar-buyer fraud: the dispute crisis draining UK banks

Matt Barr
Product Director, Moneyhub

Table of contents

The enemy within

The most expensive fraud epidemic facing UK banks today isn't driven by hackers on the dark web. It is committed by their own customers.

A new generation of 'liar-buyers' is weaponising the chargeback mechanism. What was once architected in the 1970s to protect mail-order shoppers is now being exploited and used to secure illicit refunds.

While financial institutions pour millions into AI-driven defences to fight off shadowy cyber-syndicates, they remain entirely exposed to this much more costly threat closer to home.

But losing these reversed funds is only part of the problem.

Without the data to call liar-buyers' bluff, banks cannot simply dismiss a false chargeback claim. Instead, institutions are forced to drag every single dispute through a slow, manual investigation.

Ref. 2

This has inadvertently created a lose-lose financial cycle. Banks are left underwriting an inescapable operational shadow tax, forced to pay an administrative army to process thousands of claims every day that they already know are opportunistic.

With 324 million chargebacks expected to be processed globally by 2028[1], the scale of this black hole is staggering.

Defining chargebacks

Chargebacks can be classified into three categories based on the source and intent of the claim:

Third-party criminal fraud: where a stolen card or account has been used by someone other than the cardholder.

Intentional first-party misuse or friendly fraud (liar-buyer fraud): where a bad actor purposefully initiates a false dispute against a legitimate charge in order to keep both the goods and the refund.

Unintentional first-party error: where a legitimate customer raises a chargeback against a transaction they themselves authorised, simply because they didn't recognise the merchant on their statement.

To stop this siphoning of capital, the financial sector must address this data deficit at its source.

Only then can banks strip bad actors of their anonymity and restore the chargeback to its original purpose: protecting the victim. Until then, the industry will continue to sink millions into this operational shadow tax – a drain that is as unsustainable as it is preventable.

This paper focuses exclusively on what we term 'liar-buyer fraud'. It explores how opportunists are stealing from banks and exposes the true cost of investigating this deception.

The rise of liar-buyer fraud

Liar-buyer fraud has graduated from a fringe consumer quirk to a systemic exploitation of legacy financial infrastructure. The scale of this theft now rivals traditional criminal fraud.

50% of all chargebacks

Industry and merchant reports indicate that first-party misuse now drives around half of all chargebacks – and as much as 80% of fraud-related chargebacks[3,4].

Digital shoplifting

For the modern consumer, initiating a chargeback has become a convenience – a way to bypass merchant return policies by shifting the burden of proof entirely onto the bank.

Research suggests this exploitation is powered by a concerning empathy gap: more than 72% of customers believe filing a chargeback is functionally identical to requesting a refund[5]. Consumers are oblivious to the fact that their actions impose a non-discretionary tax on the bank's bottom line (or on the small businesses they often buy from).

> **"There is a perception that these disputes are victimless. Consumers often fail to see the systemic cost; they simply see an easy path to a refund."**

- Matt Barr, Product Director, Moneyhub

The result is an entrenched culture of deception, accelerated by digital word of mouth. Across social media platforms and online forums, tactics for bypassing merchant rules and securing free refunds are actively crowdsourced. What was once universally recognised as theft has been culturally rebranded as a clever loophole or a savvy budget-stretching strategy.

The practice is now so commonplace that more than **one in five UK consumers** admits to filing fraudulent chargebacks, fully aware they are gaming the system[6]. This trajectory is steepening alongside the next generation of digital natives. Among Gen Z, a striking 40% admit to submitting a false dispute over the 2024 Christmas holiday period alone[7].

As this 'hack' becomes a mainstream habit, the rise of agentic commerce – where consumers hand purchasing power and returns to AI shopping assistants – threatens to multiply this threat exponentially. Industry analysis warns delegating retail to AI won't necessarily invent new methods of stealing; it will simply automate the loopholes financial institutions are failing to close[8].

If banks cannot effectively manage human liar-buyers today, their manual back-office pipelines will buckle when AI assistants begin submitting disputes en masse.

Plausible deniability as a weapon

So, what emboldens bad actors to systematically steal from their own bank?

Whether driven by buyer's remorse or a calculated strategy to secure free goods, the mechanism is the same: opportunists have identified and weaponised a structural vulnerability in the dispute process.

When these customers-turned-liar-buyers initiate a chargeback, they exploit the 'plausible deniability' afforded by legacy payment infrastructure. They understand that their bank is operating somewhat in the dark, hamstrung by complex, often cryptic merchant billing descriptors and facing an exceptionally high regulatory bar to legally refuse a refund.

The fraudster hedges a calculated bet: if a transaction appears on a statement as a cryptic, alphanumeric string, such as A.L. RETAIL HLDG £145 or 247 MKTPL UK7H2K9, many banks lack the immediate contextual evidence or enriched data to prove what was purchased, where and by whom.

Almost universally, this bet pays off.

When the liar-buyer disputes one of these opaque transactions, back-office agents are forced to rely on light questioning and limited geographical and merchant data to get to the truth – most can't confidently challenge the claimant's story.

> **"The reality is that as long as the customer keeps saying 'no, it wasn't me,' most banks will eventually give the money back – they simply do not have the contextual data to prove otherwise. Below a certain value, the policy is often just to refund fast and move on."**
>
> - Matt Barr, Product Director, Moneyhub

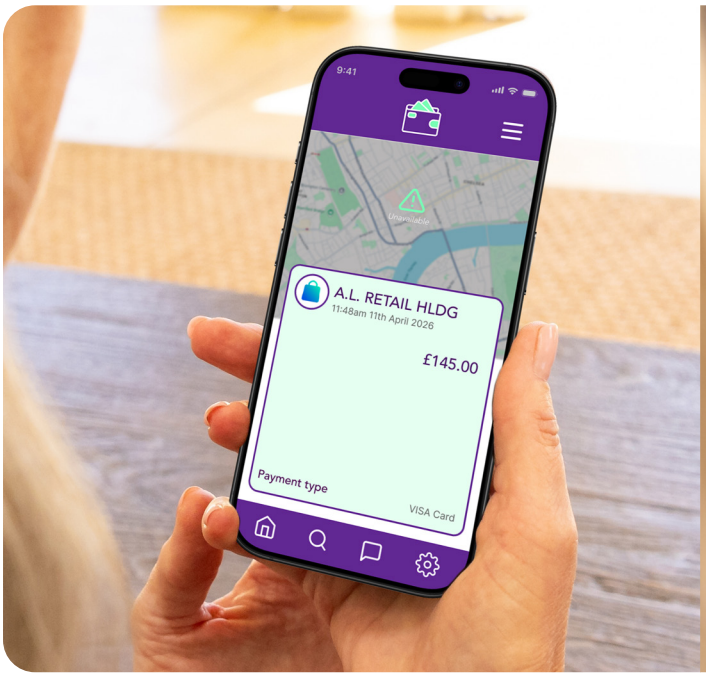

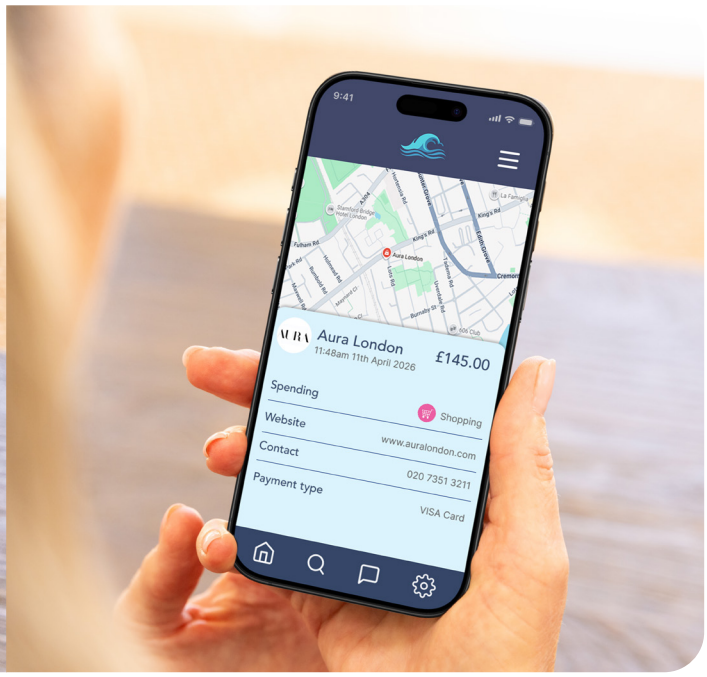

Forced to default to a refund, the bank absorbs the operational cost, processes the chargeback and effectively rewards the bad actor's deception.

Crucially, this vulnerability is not restricted to traditional high street banks. While digital-native challenger banks boast modern front-end interfaces, their in-house data enrichment can sometimes rely on the same flawed or missing merchant data as legacy institutions.

For the liar-buyer, an inaccurate map pin provides the exact same ironclad alibi as an unenriched text string.

In short, the structural vulnerabilities of an outdated payments infrastructure have transformed digital banking apps from financial tools into shortcuts for digital shoplifting.

The anatomy of digital shoplifting

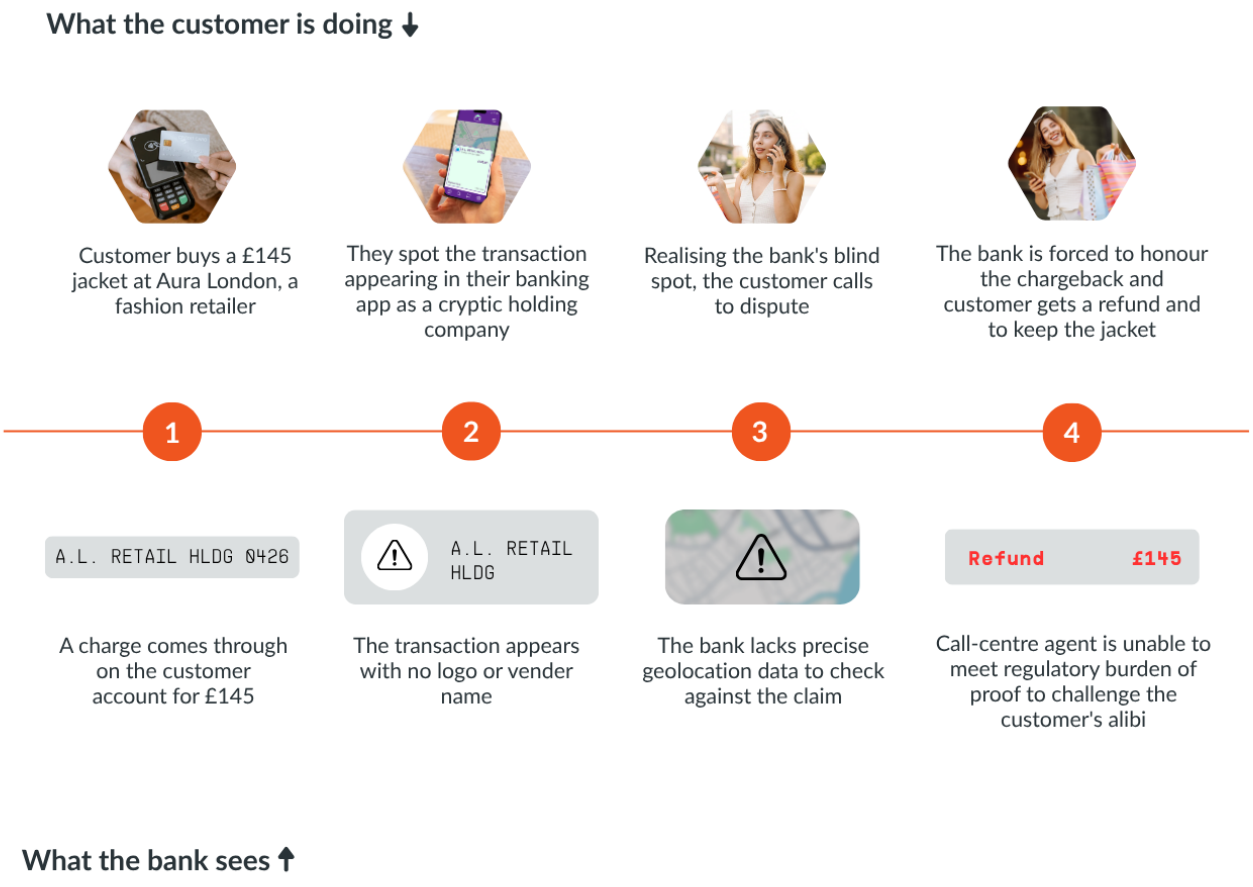

The revenue drain

For years, the banking sector has treated liar-buyer fraud primarily as a loss-prevention issue. This model is now outdated.

The true financial damage of raw or poorly enriched data extends far beyond the reversed transaction itself. Relying on blunt, highly manual processes to police an unmanageable volume of transactions creates a bottleneck that bleeds profitability.

By fighting a modern data problem with legacy infrastructure, banks are effectively levying a shadow tax on their operations, limiting revenue in three ways.

1. Ballooning payroll costs

Because the legacy payments ecosystem can't immediately distinguish between a genuine identity theft victim and an opportunist at intake, banks must treat every claim with the same gravity, triggering a mandatory, highly manual investigative pipeline.

> **"When an unrecognised transaction is reported, the raw data is missing. Who is the actual merchant? What is the real intent of the spending? Where is the physical geolocation? And did it fit into the user's wider behavioural pattern? Without that contextual backstory, back-office agents can't confidently challenge a false claim."**
>
> - Matt Barr, Product Director, Moneyhub

To resolve each dispute, back-office teams must:

1. Log the ticket
2. Retrieve any evidence available
3. Liaise with acquiring banks, and
4. Manage the file until closed.

But that's only the intake process. If the merchant contests the refund through representment – re-submitting the transaction with evidence that it was legitimate – the issuing bank's team is pulled back in to review that evidence and rule on it. Every escalation beyond that, into pre-arbitration and arbitration, adds more labour.

£800,000

For every 100,000 new disputes, banks spend nearly £800,000 in back office administrative costs, regardless of whether the claims are approved or denied.

Industry benchmarking reveals this workflow costs the issuing bank between £7.26 and £8.26 per case[2].

The investigation process is so inefficient that for every £10,400 to £11,200 worth of disputed transactions, banks are forced to add another full-time employee to the payroll[9].

This is the most expensive component of the operational shadow tax. UK banks are effectively spending millions of pounds paying skilled in-house fraud teams to process claims they already know are opportunistic.

2. The margin drain

Every time reversed funds are returned to the liar-buyer, the associated interchange fees and high-margin transaction markups vanish with them.

This is not a negligible sum. With the average UK chargeback valued at £65, banks pay regulated interchange fees of 13p (debit) to 20p (credit) on every reversed transaction.

A Tier 1 institution processing 15,000 chargebacks a week – a conservative estimate, based on confirmed cases of third party fraud alone – is not only returning hundreds of thousands in transaction value, but also losing up to £3,000 a week in interchange fees.

> **"These figures are just the tip of the iceberg. The true volume of chargebacks a bank of this size processes weekly will be many times higher, but liar-buyer fraud is rarely captured."**
>
> - Matt Barr, Product Director, Moneyhub

When compounded by lost processor markups, cross-border fees and uncapped commercial card rates, the financial exposure is significant.

3. The primacy penalty

Beyond immediate financial losses and payroll bleed, this data deficit creates collateral damage that jeopardises the bank's most valuable asset: customer retention.

While opportunists weaponise the lack of clarity to commit liar-buyer fraud, the honest majority of the bank's customer base suffers in the background. If they can't decode their transaction history, a confusion gap is created which erodes trust in the digital experience.

The battle for account primacy has never been more fiercely contested. In early 2026, current account switching rose 43%, with customers naming the online or app experience – not the interest rate – as the top reason they made the move[10].

Consumer loyalty across the sector is demonstrably fragile.

With the annual revenue yield of a primary customer calculated at around £270[11], no financial institution can afford to let opaque payment data dictate a frustrating user experience that pushes their most profitable customers straight into the arms of the competition.

Using technology to dismantle the liar-buyer

You cannot arrest your way out of liar-buyer fraud.

Attempting to fight a 2026 digital shoplifting epidemic using a 1970s criminal investigation process, as many in the banking sector are now, also won't work. Most fraud prevention techniques still focus on identity, building multi-million-pound defences to keep criminals out of accounts.

But this type of fraud is an inside job. When a verified customer uses their recognised smartphone to initiate a chargeback, an identity check simply proves the bank knows exactly who is exploiting the system.

The challenge for dispute teams is no longer identifying the customer; it is proving beyond a doubt that the customer made the purchase.

Market leaders have already abandoned broad back-office investigations in favour of a proactive, data-driven framework. The goal isn't to investigate all disputes but to intercept and expose them.

This requires a fundamental shift in how institutions treat their underlying data.

By rebuilding the dispute framework around high-fidelity, enriched transaction data, banks can systematically strip bad actors of their plausible deniability. Establishing this data foundation unlocks a multi-layered defence, allowing banks to reduce the volume and success of opportunistic claims.

The 4D Intelligence Framework

To strip away the cover of an opaque billing string and prove a customer made a purchase, banks must stop relying on flat, one-dimensional data. Instead, they should look to translate raw payment data using a process we call **The 4D Intelligence Framework.**

This approach helps equip dispute teams with the evidence they need to challenge a false claim, across four key pillars:

Identity (the who): translating alphanumeric strings into recognisable merchant identities, with brand logos and actionable contact information.

Intent (the what): going beyond the merchant name to determine the true nature of the spend.

Geolocation (the where): pinpointing the precise physical location of the point-of-sale terminal, overriding default mapping tools that frequently misattribute purchases to a corporate holding address.

Behaviour (the when): moving beyond simple timestamps to detect regular series spending. Applied AI technology links card purchases into behavioural routines, such as a consumer's regular Friday night takeaway, establishing a pattern of behaviour without a formal standing order.

> **“We’re building up a contextual image of the bank’s customer over time. When agents have easy access to that person’s complete backstory, they have the leverage to ask more questions and push back on false claims.”**

- Matt Barr, Product Director, Moneyhub

While many in the industry recognise the need to arm agents with more information, early solutions have limits.

Some try to build systems in-house, relying heavily on basic Merchant Category Codes (MCC). The problem is that MCCs are blunt instruments; they often stall at 70-80% accuracy and lack the granularity to truly understand consumer intent or accurately map a customer’s daily routine[12]. Reaching the level of detail required to fight liar-buyer fraud is a serious investment, typically requiring years of development and machine-learning models trained over billions of data points.

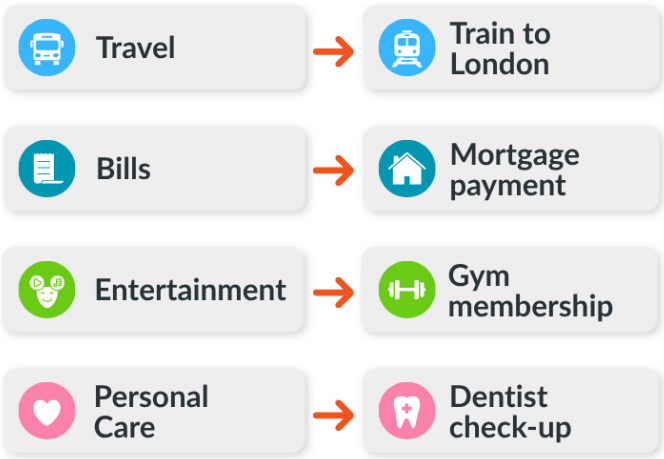

Others rely on basic API mapping models. While these consumer-facing tools can successfully provide the account holder with a merchant logo and location pin in their app, the licensing terms prohibit data caching.

This means the bank cannot legally store the enriched data. Building a behavioural history over time becomes difficult, so when a liar-buyer calls to dispute a charge, the customer agent is forced to view the transaction in isolation. They don’t have the historical context needed to confidently challenge the claim.

By deploying an AI-first categorisation and enrichment solution, banks can move from rented visual overlays to owned operational intelligence. This machine-learning infrastructure is designed to take a garbled text string with raw data and autonomously translate it into the four layers of 4D Intelligence: identity, intent, geolocation and behaviour.

Deploying active defence

With this infrastructure in place, banks can access a permanent historical record of the customer’s true spending behaviour. Instead of flying blind, back-office teams are empowered to dismantle the liar-buyer’s advantage across two expensive stages of the chargeback lifecycle: the dispute source and representment.

The moment a dispute is initiated, a categorisation and enrichment solution could leverage AI and behavioural analytics to cross-reference the transaction against the customer’s mapped spending history. Rather than pushing every single dispute into an expensive back-office queue that leaves call centres with 67-minute average wait times[12] and forces banks to hire hundreds of staff, the technology can autonomously segment the threat into likely criminal fraud, billing confusion and liar-buyer claims.

Confirmed victims of third-party criminal fraud are fast-tracked for immediate refunds, while high-probability liar-buyers are flagged for robust evidence gathering.

> **"If you look at one transaction in isolation, the liar-buyer has the advantage. But if an agent has a mapped timeline, that contextual history provides the exact friction they need to push back."**
>
> - Matt Barr, Product Director, Moneyhub

While intelligent triage intercepts casual opportunists during the initial phone call, some claims inevitably escalate into formal chargebacks.

Historically, fighting every single chargeback has required too much manual labour to be cost-effective, but by automating the representment process, whereby false disputes are challenged to recover the reversed funds with the important addition of enriched data, institutions can secure a powerful return on investment.

For example, banks can use Moneyhub's applied AI tool, the Categorisation and Enrichment Engine, to run a pre-emptive assessment before an analyst even touches a case file.

Trained specifically to strip manual labour out of the dispute process, it can read the customer's purchasing history in an instant. If it detects a deep, multi-year spending pattern with the disputed merchant's parent company, the bank can internally flag the chargeback will likely face successful merchant representment. Knowing the merchant has the evidence to fight back and win, the bank can decline to progress the dispute – bypassing the cost of processing official paperwork.

And if a merchant submits a representment pack to challenge a claim, the bank's operational burden vanishes. The technology programmatically validates the merchant's evidence against the customer's enriched footprint, meaning little to no cross-referencing of cryptic store IDs or trading names.

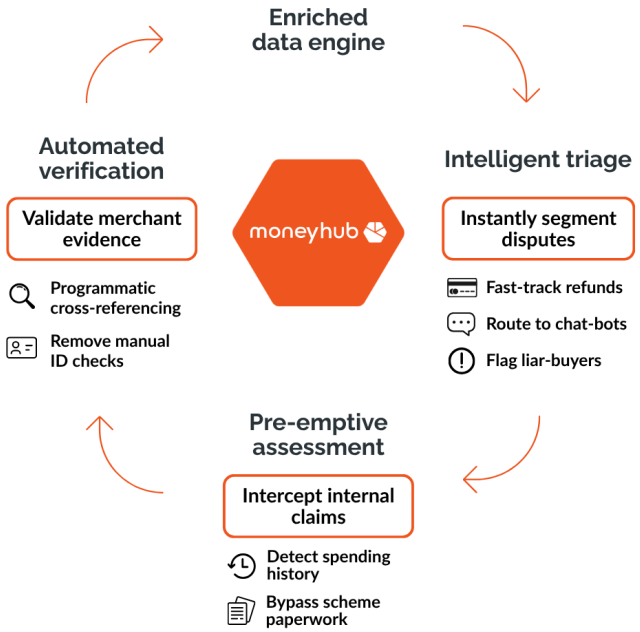

By leveraging this intelligence to build an irrefutable, automated business case, financial institutions are provided with the tools to recover direct and indirect lost revenue, penalise serial abusers and send a deterrent signal to the wider ecosystem.

The future of transaction intelligence

The multi-million-pound chargeback crisis is no longer just a byproduct of external cybercrime. Driven by the structural vulnerabilities of a legacy payments ecosystem, liar-buyer fraud has escalated from a fringe consumer behaviour into a seismic financial drain.

Trapped in a highly manual investigation pipeline, banks have been left bankrolling a phantom workforce: a costly shadow tax, born not from fraud itself, but from the sheer human capital required to investigate it. As agentic commerce threatens to automate these loopholes even further, throwing people at a data problem is no longer sustainable.

This is the consequence of an industry-wide intelligence gap – one that should no longer be considered an inevitable cost of doing business. To stop liar-buyers from siphoning capital out of the system, the sector must abandon its reliance on reactive, post-dispute investigations.

By investing in advanced categorisation and enrichment technology, banks can empower fraud teams to automatically triage claims at the source, instantly invalidate opportunistic alibis and orchestrate a potentially lucrative representment strategy.

Ready to stop underwriting digital shoplifting?

Discover how Moneyhub's applied AI Categorisation and Enrichment Engine can help you recover lost revenue today.

References

1. https://www.mastercard.com/global/en/news-and-trends/Insights/2025/2025-global-chargebacks-outlook.html

2. Moneyhub estimate. The £7.26 to £8.26 operational cost per case was calculated by applying a £0.80/$1.00 exchange rate against the $9.08 to $10.32 baseline processing cost sourced from Mastercard, What's the true cost of a chargeback in 2025? (2025).

3. https://merchantriskcouncil.org/advocacy/first-party-misuse-aka-friendly-fraud

4. https://merchantriskcouncil.org/learning/resource-center/member-news/news/2024/2024-chargeback-field-report-merchant-survey-reveals-latest-trends-impacts-of-friendly-fraud

5. https://chargebacks911.com/the-2023-chargeback-field-report-is-here/

6. https://paystrax.com/wp-content/uploads/2026/03/Friendly-Fraud-PAYSTRAX.pdf

7. https://www.socure.com/news-and-press/first-party-fraud-over-holiday-shopping-season

8. https://brand.adyen.com/share/Zg2pFLyWGdHB9ux69qH3/assets/126754

9. Moneyhub estimate. The £10,400 to £11,200 FTE cost was calculated by applying a £0.80/$1.00 exchange rate against the $13,000 to $14,000 baseline value sourced from Mastercard, What's the true cost of a chargeback in 2025? (2025).

10. https://www.currentaccountswitch.co.uk/news-insights/switching-data/current-account-switch-service-reports-319-529-switches-in-q1-2026

11. https://www.starlingbank.com/news/starling-delivers-fifth-year-of-profitability-and-accelerates-global-growth-strategy/

12. https://moneyhub.com/use-cases/identify-fraud-real-time-banking-app/

moneyhub

Build for real life.

www.moneyhub.com