



KEDBYTE

SOFTWARE · AI · AUTOMATION

HOW MONEY MOVES

Volume V — Trust, Failure and the Law

The complete plain-English guide to payments and banking,
followed by the engineer's version of the same thing.

WRITTEN BY

Shikhar Singh

Founder & Chairman

KedByte Technologies Private Limited

Five Volumes • Fifty-Four Chapters • First Edition

HOW MONEY MOVES

Volume V — Trust, Failure and the Law

Author	Shikhar Singh
Designation	Founder & Chairman, KedByte Technologies Private Limited
Published by	KedByte Technologies Private Limited
Imprint	KedByte Press
Edition	First Edition
Subject	Payments, banking systems, financial technology
Extent	Five volumes, fifty-four chapters

Copyright © KedByte Technologies Private Limited. All rights reserved.

No part of this publication may be reproduced, distributed or transmitted in any form or by any means, including photocopying, recording or other electronic or mechanical methods, without the prior written permission of the publisher, except in the case of brief quotations embodied in critical reviews and certain other non-commercial uses permitted by copyright law.

Trademarks. All product names, company names, service marks and trademarks referred to in this book are the property of their respective owners. They are used in this book in an editorial and educational capacity only, with no intention of infringement of the trademark. No such use, or the use of any trade name, is intended to convey endorsement or any other affiliation with this book or its publisher.

Accuracy and currency. This book describes technology that changes. Specifications, prices, version numbers, product names and industry figures were checked at the time of writing and are dated in the text wherever they are liable to change. The chapters on artificial intelligence in particular describe a field that moves quickly; readers are told, in the text, where to re-check a figure before relying on it. Where experts disagree, the book says so rather than choosing a side.

Disclaimer. The information in this book is provided on an “as is” basis, for education. The security chapter is written for defence: it explains how classes of attack work so that systems can be built to resist them, and it deliberately contains no working exploit code. Neither the author nor KedByte Technologies Private Limited shall have any liability to any person or entity with respect to any loss or damage caused, or alleged to have been caused, directly or indirectly by the information contained in this book.

Worked examples. Parts F and G are built around one real network fault and one real version-control session, reproduced exactly as they were observed, including the addresses, the traceroute output and the error messages. Nothing in those traces has been altered or invented.

From the Founder & Chairman

You tap a card against a machine and a shop lets you leave with the coffee. Somewhere behind that gesture, money you cannot see is supposed to move from a place you have never visited to another place you have never visited, and almost nobody involved can tell you how. Not the person who served you. Often not the person who installed the terminal. Frequently not the people who work at the bank.

That collective not-knowing is expensive. It produces businesses that cannot reconcile their own takings, engineers who store money in floating point, support teams who ask customers for the one reference number that cannot possibly help, and founders who discover the difference between a payment being authorised and a payment being settled on the day it costs them something. None of these are stupidity. They are the predictable result of an industry that documents itself for people who are already inside it.

This book is the outsider's way in, written to be useful to the insider. It begins with what money actually is, which is a record of obligation rather than a substance, and it ends with regulation, dispute and the ugly arithmetic of reconciliation. In between: the card in your pocket and the small computer embedded in it, the machine on the counter, the message formats that have carried the world's card traffic for forty years, the difference between a payment clearing and a payment settling, the rails that move money without cards at all, and what happens when any of it fails.

The method is the same one used in *How Computers Work*, and it matters more than the coverage. Every idea is explained twice. First in plain language, with an everyday comparison and a worked example, such that a twelve-year-old could follow it. Then again in the exact professional vocabulary, with the real field numbers, the real timings and the real institution names that a practitioner would use. Between the two sits a section that most writing on this subject omits entirely: an explicit statement of where the simple version stops being true. A useful simplification told without warning is how people end up confidently wrong, and in payments confidently wrong is a category of loss.

One thing about this book differs from its predecessor, and the reader should know it. How a transistor works will not change. Payments will. Scheme rulebooks are revised annually, regulatory caps move, and migrations between message standards run on live deadlines. The volumes are therefore arranged so that decay is contained: the mechanics that will outlive us sit in the early volumes, and the material that is true as of writing is gathered deliberately at the end, marked as

such. Where a figure is time-bound, this book says so rather than pretending to a permanence it does not have.

At KedByte Technologies Private Limited we build systems that move other people's money, which is a privilege that ought to come with an obligation to explain the machinery. If this book leaves one reader able to follow their own money from a card terminal to a bank ledger, and to say honestly which part of that journey they have verified and which part they have merely assumed, it has done its job.

Start at Chapter 1. It begins long before cards, and for good reason.

Shikhar Singh

Founder & Chairman

KedByte Technologies Private Limited

How to read this book

The shape of every section

Every section of every chapter is built from the same six blocks, always in the same order. Once you know the shape, you can navigate two thousand pages without ever being lost.

Block	What it is for
PLAIN in simple words	The idea for a complete beginner. No jargon at all.
PLAIN a picture in your head	One everyday comparison, and then a line telling you exactly where that comparison stops being true.
PLAIN a worked example	Concrete numbers, worked step by step.
PLAIN what is really happening inside	The mechanism, still in easy words, but with nothing hand-waved.
TECHNICAL the engineer's version	The correct professional vocabulary, exact units, real specifications, real figures, standard numbers and the commands that observe it.
WORDS remember these	Each term given twice: the plain meaning, then the technical meaning.

Two ways to read it

1. **The single pass.** Read everything, in order, from Chapter 1. This is the intended route and it works.
2. **The double pass.** Read only the blocks marked PLAIN, front to back. You will understand the whole of computing. Then return to the start and read the blocks marked TECHNICAL to be able to hold your own in an engineering conversation.

Conventions used throughout

1. Almost everything is written as short numbered lines, one idea per line, so nothing hides inside a paragraph.
2. *The honest version* marks the places where a simple explanation is not strictly true. The accurate model follows immediately.

3. The book states whether a thing is a **standard** (written down in a specification), a **convention** (everybody just does it this way), or an **implementation detail** (true of one product, not of the idea).
4. In the artificial intelligence chapters, claims are separated into established fact, active research, and marketing claim.
5. Anything that goes stale is dated. Check it again before you rely on it.
6. Every chapter ends with *Common wrong ideas* and a twenty-line summary. If you are short of time, read those two first.

Where to find things

1. **Chapter 1** is the map of the whole book.
2. **Chapter 52** is the reference desk: the numbers, the command cheat sheets and the two diagnostic decision trees.
3. **Chapter 53** is the master timeline and a glossary of more than seven hundred terms, each given in plain words and then technically.

Contents

Volume V — Trust, Failure and the Law

1

1 Where Cards Are Stolen

2

45.0 What this chapter gives you	2
The plain version	3
Where the plain version stops being true	4
The technical version	5
45.98 Common wrong ideas	14
45.99 Chapter summary in 20 lines	15

2 Card Not Present

17

46.0 What this chapter gives you	17
The plain version	17
Where the plain version stops being true	19
The technical version	20
46.98 Common wrong ideas	29
46.99 Chapter summary in 20 lines	30

3 Three-D Secure and Strong Customer Authentication

32

47.0 What this chapter gives you	32
The plain version	32
Where the plain version stops being true	34
The technical version	36
47.98 Common wrong ideas	45
47.99 Chapter summary in 20 lines	45

4 PCI DSS and Its Relatives

48

48.0 What this chapter gives you	48
The plain version	48
Where the plain version stops being true	50
The technical version	51
48.98 Common wrong ideas	61
48.99 Chapter summary in 20 lines	62

5 Chargebacks

64

49.0 What this chapter gives you	64
The plain version	64
Where the plain version stops being true	66
The technical version	67
49.98 Common wrong ideas	77
49.99 Chapter summary in 20 lines	77

6 Reconciliation

79

50.0 What this chapter gives you	79
The plain version	79
Where the plain version stops being true	82

The technical version	83
50.98 Common wrong ideas	90
50.99 Chapter summary in 20 lines	91
7 Fraud Detection	93
51.0 What this chapter gives you	93
The plain version	93
Where the plain version stops being true.	96
The technical version	97
51.98 Common wrong ideas	105
51.99 Chapter summary in 20 lines	105
8 Anti-Money-Laundering	107
52.0 What this chapter gives you	107
The plain version	107
Where the plain version stops being true.	109
The technical version	110
52.98 Common wrong ideas	118
52.99 Chapter summary in 20 lines	119
9 Regulation	121
53.0 What this chapter gives you	121
The plain version	121
Where the plain version stops being true.	123
The technical version	124
53.98 Common wrong ideas	133
53.99 Chapter summary in 20 lines	134
10 What Breaks and What It Costs	136
54.0 What this chapter gives you	136
The plain version	137
Where the plain version stops being true.	138
The technical version	140
54.98 Common wrong ideas	148
54.99 Chapter summary in 20 lines	149



VOLUME V

Trust, Failure and the Law

Fraud, disputes, reconciliation and regulation. The
perishable volume, dated deliberately and revised on its
own.

Where Cards Are Stolen

45.0 What this chapter gives you

1. You will be able to name the five places card details are actually taken — the stripe, the shim, the radio, the breach and the guess — and say which of them still costs anybody real money.
2. You will be able to explain why watching a chip talk to a terminal gets you a used ticket, while reading a magnetic stripe gets you a key that works for ever.
3. You will be able to say why the chip did not make card fraud go away, and point at £423.5 million of remote purchase losses in the United Kingdom in 2025 as the evidence of where it went instead.
4. You will be able to answer somebody worried about contactless in a queue with the actual proportion: £46.8 million of £703.4 million of unauthorised losses, under seven per cent.
5. You will be able to explain why a shim yields far less than a skimmer ever did, and name the issuer-side checks — iCVV, service code, transaction counter, fallback constraint — that decide whether it yields anything at all.
6. You will be able to distinguish a server-side breach from client-side e-skimming, and say why a merchant's own logs stay clean throughout the second one.
7. You will be able to cite PCI DSS requirements 6.4.3 and 11.6.1, mandatory since 31 March 2025, as the two controls written specifically for payment-page script theft.
8. You will be able to do the enumeration arithmetic — ten million accounts under an eight-digit issuer identification number, sixty expiry dates, a thousand three-digit codes — and explain why centralised counting rather than merchant rate limiting is the fix.
9. You will be able to read the Visa Acquirer Monitoring Program thresholds and say why the enumeration ratio deliberately counts declined authorisations on both sides of the fraction.
10. You will be able to state a United Kingdom cardholder's legal position under regulations 76 and 77 of the Payment Services Regulations 2017, including the £35 cap and the circumstances that disapply it.

Everything in the four volumes before this one described a system working. This volume describes the same system under attack, and the first thing to establish is that “card fraud” is not one crime. It is at least five, they have almost nothing in common technically, and the defences that stopped the first one did nothing at all about the fifth.

This chapter is the threat model. It is written to be useful to somebody who has to defend a payment system, not to somebody who wants to attack one, and that distinction shapes what is on the page: how each attack works and why each control exists, without the operational detail that would make any of it a recipe. Where a control is mandated by a rulebook or a regulator, the figure and the date are given, and both are marked as accurate at the time of writing, because in this volume more than any other they will move.

The plain version

Think of your card as a key to a door.

For the first forty years of the card industry, it was a very simple key: a shape stamped into a piece of metal. Anybody who could get a good look at the key could file an identical one. The magnetic stripe on the back of your card is exactly that. It is a strip of iron oxide with your card number, expiry date and a few other details written into it, in the clear, and it never changes. A machine that reads the stripe learns everything it needs to make a stripe that behaves identically. That is why, for decades, the standard card crime was to read stripes and write copies.

Then the industry replaced the key with a locksmith.

The chip in your card is a very small computer with a secret inside it that it will never tell anyone, including you. When you put the card into a machine, the machine does not ask “what is your key”. It asks a question, and the chip writes a short note in reply, sealed using the secret. Crucially, the question is different every time, because the machine throws in a fresh random number, and the chip adds a counter that ticks up on every single transaction it has ever done. So the reply is different every time too.

Copy that reply and you have a note that answers one question that will never be asked again. Watching the conversation between a chip and a machine gets you a used ticket. Watching a magnetic stripe gets you a key.

That single change is why the crime moved.

■ Where it moved to

Here are the five places card details are actually taken, in the order they became important.

The first is the stripe. Put a reader over the slot on a cash machine, or run a card through a second reader behind the counter, and you get a perfect copy. This is skimming, and where chips are used properly it is close to dead. The industry’s own numbers show why: EMVCo, the body that owns the chip specification, reports 15.6 billion cards in circulation using its specifications, and that 97 per cent of card-present transactions worldwide are chip transactions, most recently reported as of the fourth quarter of 2025. There is very little stripe left to skim.

The second is a thin device slipped into the slot of a chip reader, sitting between the card’s metal contacts and the machine’s. It is called a shim, and it listens to the conversation. It gets the card number, the expiry date, and a great deal else. What it does not get is the secret. The reply the chip wrote is a used ticket, and so a shim is worth far less than a skimmer ever was.

The third is the radio. A contactless card answers anyone who asks, from a few centimetres away, without a PIN. So the obvious idea is to put one machine near a real card and another machine near a real terminal and pipe the conversation between them, so that a card in a pocket in one place pays for something in another. That is a relay attack. It works in laboratories. It is limited in the world by two things: the radio only reaches centimetres, so somebody has to physically get close to the card, and the terminal times the conversation and gets suspicious if the answers come back late.

The fourth is the one that actually matters now, and it is not clever at all. Break into the shop’s computer, or the payment company’s computer, and take the list. In November and December 2013, attackers who had obtained a supplier’s credentials got inside the American retailer Target and took details for up to 40 million payment cards, along with contact details for up to 70 million more people. No skimmer, no shim, no radio. One break-in, forty million cards.

The modern version of the same idea does not even need to get into the shop. It gets into the shop's website. A few lines of code slipped into a checkout page will quietly copy what you type into the card fields and send it somewhere else, before it ever reaches the shop. This happened to British Airways in 2018. The Information Commissioner's Office found the attacker was in a position to access the personal data of approximately 429,612 customers and staff, including the payment card details of around 244,000 customers, and issued a penalty notice of £20 million on 16 October 2020.

The fifth is guessing. A card number is sixteen digits, but the first eight identify the bank, and the last is a check digit calculated from the others. So under one bank's range there are about ten million possible accounts. An expiry date within the next five years is one of sixty possibilities. The three digits on the back are one of a thousand. Nobody can guess a whole card. But if you already have the number from a breach, you only need the expiry and the three digits, which is sixty thousand combinations, and a machine can try them at a website far faster than a person can. This is called card testing, and the small charity donation page with no shipping address and a five-pound minimum is its favourite target.

■ What the totals look like

The Nilson Report, which counts this for the industry, published in January 2026 that worldwide payment card fraud losses were \$33.41 billion in 2024, against global card volume of \$51.920 trillion. That is 6.43 cents of fraud per \$100 spent, down from 6.58 cents the year before.

Six and a half pence in a hundred pounds. Card fraud is enormous in absolute terms and small in proportional terms, and both facts are true at once, which is why arguments about it are so bad-tempered.

In the United Kingdom, UK Finance reported on 15 June 2026 that criminals stole £1.28 billion through payment fraud in 2025. Of that, £703.4 million was unauthorised fraud, meaning transactions the customer never made at all. The breakdown is the point:

Type of unauthorised fraud, UK, 2025	Losses	Change on 2024
Remote purchase (card details used online)	£423.5m	up 3%
Lost and stolen cards	£109.8m	down 2%
Remote banking	£104.4m	down 27%
Contactless	£46.8m	up 8%
Card ID theft	£54m	down

Remote purchase fraud, on its own, is more than sixty per cent of all unauthorised card losses. That is card details typed into a website. Contactless, the thing people worry about in queues, is £46.8 million out of £703.4 million: under seven per cent.

The lesson of the plain version is this. The chip did not make card fraud go away. It made one kind of card fraud, the counterfeit card, extremely hard, and everything else moved to the internet, where there is no chip to consult.

Where the plain version stops being true

"The chip ended counterfeit fraud" is a claim about a control, and controls only work where they are switched on. The Federal Reserve Bank of Kansas City published a research briefing on 12 February 2025, using the Federal Reserve Board's biennial debit card data, which found that the United States migration to chip cards did not produce a uniform fall in card-present fraud. On dual-message

networks the counterfeit fraud rate was 9.0 basis points in 2015, fell to 6.6 in 2017, and was back to 9.1 by 2021. Over the same period the lost-or-stolen rate on those networks rose from 1.6 to 4.5 basis points, and overall card-present fraud rose from 10.1 basis points in 2017 to 14.9 in 2021. Only single-message networks showed the expected pattern, with counterfeit falling from 3.2 basis points in 2015 to 1.7 in 2021. The chip is sound. Its deployment was uneven, its fallback paths were exploitable, and fraud is a fluid: squeeze it and it goes somewhere else in the same container.

The chip protects the reply, not the number. I said the chip never reveals its secret, and that is true. What the chip does reveal, in the clear, on every transaction, is the card number, the expiry date and a block of track-equivalent data. Reading a chip conversation does not let you counterfeit a chip card. It does let you obtain the details that a website will accept. The plain version implies the chip made your card details secret. It did not. It made them non-replayable at a terminal. On a website, where there is no terminal and no cryptogram, they are as useful as they ever were, which is exactly why remote purchase fraud is the largest category in the table above.

Most stolen card data is now not stolen from cards at all, and a growing share is not stolen at all. The dominant modern source is a breached merchant, processor or checkout page, which touches no physical card and no physical terminal. Alongside it, and rising fast, is a category that involves no theft of card data in the conventional sense: the criminal already has the card number and telephones the customer to obtain the one-time passcode, then uses it to provision that card into a digital wallet on their own device. UK Finance's report of 15 June 2026 names this explicitly as a driver of 2025 losses. The result is a wallet credential that passes every check the network can make, on a device that is not yours. No skimmer, no shim, no breach. An authentication defeat dressed as a card theft.

Relay attacks are famous out of all proportion to their share of losses, and the reason is worth understanding. A relay does not break the cryptography. It borrows it. The issuer receives a genuine cryptogram from a genuine card, correctly computed, correctly verified. Nothing in the message says the card was three hundred miles away. That makes a relay attack unfalsifiable from the issuer's side and turns the dispute into an evidential argument rather than a technical one, which is what makes it interesting to researchers and to lawyers. It does not make it common. The £46.8 million of contactless fraud UK Finance reported for 2025 is dominated by ordinary use of ordinary stolen cards, not by relays.

"Card data" is not one asset. A PAN alone, a PAN with expiry, a PAN with expiry and CVV2, full track data, a PIN block, and a network token are six different things with six different resale values, six different attack paths and six different controls. Treating them as one category is the single most common analytical error in this field, and it is the reason people build defences in the wrong place.

The technical version

■ What lives where on a card

Three credentials coexist on a typical card, and each fails differently.

The magnetic stripe follows ISO/IEC 7811. Track 1, the IATA track, carries up to 79 alphanumeric characters including the primary account number, cardholder name, expiry date, a three-digit service code and discretionary data. Track 2, the ABA track, carries up to 40 numeric characters: PAN, expiry, service code and discretionary data. The discretionary data contains the card verification value written at personalisation, CVV1 on Visa or CVC1 on Mastercard. The service code matters more than its size suggests: its first digit encodes interchange rules, with values in the 2xx range signalling that the card

carries a chip and that the chip must be used where the terminal is capable, and its third digit encodes cardholder verification requirements.

Nothing on the stripe changes. It is a bearer credential in the purest sense.

The chip follows ISO/IEC 7816 for the contact interface and holds an EMV payment application, a set of records read by the terminal, a card private key or shared symmetric keys it never discloses, and an application transaction counter. Among the records is tag 57, track-2 equivalent data, which looks like Track 2 but carries a different verification value in the discretionary field: the integrated circuit card verification value, iCVV. It is deliberately different from CVV1. That difference is the entire defence against turning chip data into a stripe.

The contactless interface follows ISO/IEC 14443, operating at 13.56 MHz with a baseline data rate of 106 kbit/s in each direction, half duplex, over a designed range of a few centimetres. In full EMV contactless mode it produces a cryptogram exactly as the contact interface does. The older magnetic-stripe emulation modes, which produced a dynamic CVV or CVC3 instead, are the modes that most published contactless attacks targeted, and their retirement is a large part of why those attacks have narrowed.

The card-not-present credential is different again: PAN, expiry, and the CVV2, CVC2, CID or CAV2 printed on the card but written into neither the stripe nor the chip. It exists precisely so that stripe or chip data alone is insufficient online.

Finally, the tokenised credential. Under the EMV Payment Tokenisation Specification, a device token or DPAN is provisioned in place of the funding PAN, restricted by domain to a particular device, channel or merchant, and accompanied by a cryptogram of its own. A stolen DPAN outside its domain is inert.

■ Magnetic stripe cloning, and why the chip ended it

A stripe clone is a replay attack against a credential with no freshness. The attacker acquires the track data by any means, writes it to a stripe, and presents it. The issuer sees identical data to the genuine card and has no basis on which to distinguish them. CVV1 in the discretionary data defends against a specific and narrower threat: card details harvested from receipts, embossing or a manual imprint, which do not include it. Against a full track read it is no defence at all, because it was read too.

EMV removes replay by three mechanisms working together. The terminal supplies an unpredictable number, tag 9F37, which is fresh per transaction. The card supplies an application transaction counter, tag 9F36, which increments monotonically and never repeats. The card computes an authorisation request cryptogram over transaction data including both, using a key derived from an issuer master key that the card holds and never emits. The issuer recomputes the cryptogram and compares. A captured cryptogram authenticates one transaction, with one nonce and one counter value, and is worthless afterwards.

The migration was driven by liability rather than by technology, and the dates are worth having exactly. In the United Kingdom the liability shift took effect on 1 January 2005, moving responsibility for fraudulent magnetic swipe transactions from issuers to retailers who could not accept chip, and the national rollout completed on 14 February 2006. Counterfeit card fraud losses in the United Kingdom, reported by what was then the UK Cards Association, fell from £169.8 million in 2008 to £80.9 million in 2009. In the United States the equivalent shift for most point-of-sale acceptance took effect on 1 October 2015, roughly a decade later, with automated fuel dispensers deferred beyond that.

The stripe itself is now being withdrawn. Mastercard announced on 12 August 2021 that newly issued Mastercard credit and debit cards would not be required to carry a stripe from 2024 in most markets,

that United States banks would no longer be required to issue chip cards with a stripe from 2027, that no new Mastercard credit or debit card would be issued with a stripe by 2029, and that no Mastercard credit or debit card would have a stripe by 2033, with prepaid cards in the United States and Canada exempt at the time of the announcement. Those dates are the scheme's published plan as of writing in August 2026 and are subject to change by the scheme.

Until then, the residual controls are issuer-side and they are the ones that determine whether a shim or a chip-derived clone is worth anything. They are: validating CVV1 on any stripe-read transaction; validating iCVV separately on chip transactions and refusing to accept one where the other is expected; checking the service code against the card's actual profile so that a stripe claiming no chip is present on a card that has one is treated as suspicious; monitoring the application transaction counter for non-monotonic or duplicated values; and constraining technical fallback, where a chip read fails and the terminal reverts to the stripe, by volume, geography and merchant category.

■ **Shimming, and why it yields far less**

A shim is a passive device placed in the card slot between the card's contacts and the reader's, recording the application protocol data units exchanged under ISO/IEC 7816-3. The reason it exists is that the chip's dialogue with the terminal is not encrypted; confidentiality was never its security property. Authenticity was.

What the exchange discloses is the static record content the card returns during read records: the PAN, the expiry date, the cardholder name where present, the track-2 equivalent data of tag 57 including the iCVV, the application transaction counter, and the cryptogram computed for that one transaction. What it does not disclose is any key. The card never transmits its private key or its derived session keys, and the cryptogram is a one-way function of data including a nonce the attacker cannot control.

So a shim produces two products, and both are constrained.

The first is card-not-present abuse. PAN and expiry are enough at any merchant that does not require CVV2, and are a strong starting point at any merchant that does, because they reduce the guessing problem to the remaining fields.

The second is what Recorded Future named EMV-bypass cloning in 2020: writing chip-derived track-2 equivalent data onto a magnetic stripe and presenting it as a stripe transaction. This should fail, because the iCVV in the chip record differs from the CVV1 the issuer expects from a stripe. It succeeds only where the issuer does not validate the value, or validates the wrong one, or does not check the service code. Krebs on Security, reporting Royal Canadian Mounted Police material in January 2017, made the same point when shimmers first appeared publicly: shimmers exist because some issuers did not implement the standard correctly. The vulnerability is not in EMV. It is in a small number of issuer authorisation stacks, and it is closable by configuration.

Defences are correspondingly split between the issuer and the device. On the issuer side, the controls listed at the end of the previous subsection are the whole answer, and an issuer that applies all of them reduces a shim to a source of card-not-present data. On the device side, cash machine and unattended terminal manufacturers have added physical countermeasures to the card path, and terminals are certified against the PCI PIN Transaction Security Point of Interaction requirements, currently at version 7.0 as of writing in August 2026, which cover tamper detection, tamper responsiveness and physical security of the card reader and PIN entry path. What no device countermeasure can do is make the chip dialogue secret, because the terminal must be able to read it.

■ Relay attacks on contactless, and the timing that limits them

A relay attack is a proxy. One device presents itself to the genuine card as a terminal; a second presents itself to the genuine terminal as a card; the two are linked. Nothing is decrypted, forged or replayed. The cryptogram is genuine, the counter is correct, and the issuer's verification succeeds, because from the protocol's point of view the card really did participate.

Three things constrain it.

The first is physics. ISO/IEC 14443 is a proximity standard. The reader field falls away over centimetres, so one end of the relay must be brought very close to the victim's card. That is a physical-access problem, and it is why relay attacks appear in demonstrations and disputes far more often than in loss statistics.

The second is the protocol's own patience. ISO/IEC 14443-4 defines a frame waiting time, and EMV contactless kernels enforce an overall transaction time budget, commonly cited in the research literature at around 500 milliseconds, because the entire commercial proposition of contactless is that the customer does not wait. A relay adds latency at both radio interfaces and across whatever link joins them, and that latency competes against a budget that was set for user experience rather than for security but happens to serve both.

The third is an explicit countermeasure. EMV specifies a Relay Resistance Protocol, defined in the EMV Contactless Specifications for Payment Systems, Book C-2, the Mastercard kernel, and carried forward into EMVCo's Kernel 8. It works by adding a timed exchange. The terminal issues an EXCHANGE RELAY RESISTANCE DATA command carrying terminal entropy and starts a timer. The card replies with its own entropy and with its Device Estimated Transmission Time For Relay Resistance R-APDU, a two-byte value in units of hundreds of microseconds stating how long it expects its own response to take. The terminal records a Measured Relay Resistance Time in the same units and compares the measurement against the card's declaration, its own Terminal Expected Transmission Time For Relay Resistance R-APDU, and configured Relay Resistance Minimum and Maximum Time Difference Limits. If the measurement falls outside the window, the terminal can decline or force the transaction online.

It is important to be precise about what this achieves, because it is routinely overstated. RRP is a latency bound, not a cryptographic distance-bounding protocol in the academic sense. Its granularity is hundreds of microseconds; one hundred microseconds is thirty kilometres of light travel in free space. The protocol therefore cannot localise a card geometrically at all. What it does is impose a ceiling on the additional processing and network latency an attacker can introduce, which is a real and useful constraint against a relay built over general-purpose networks, and a much weaker one against a purpose-built low-latency link. Published analysis, including the systematisation-of-knowledge literature on EMV contactless security, notes that measured RRP timings vary with card position and orientation, which forces the acceptance window wider than a security engineer would like.

The academic record is worth citing accurately, because it is what drove the fixes. Francis, Hancke, Mayes and Markantonakis demonstrated practical relay of contactless transactions using ordinary NFC-capable mobile telephones in 2010. Radu, Chothia, Newton, Boureanu and Chen published "Practical EMV Relay Protection" at the IEEE Symposium on Security and Privacy in 2022, demonstrating a relay against Apple Pay in Express Transit mode with a Visa credential, exploiting the fact that transit modes deliberately relax cardholder verification so that a locked phone can pay at a gate; Apple and Visa subsequently addressed it. A cluster of work between 2019 and 2021, much of it by Galloway, Yunusov and colleagues, showed cardholder verification and limit bypasses achieved by

manipulating terminal transaction qualifiers and card transaction qualifiers in the contactless dialogue rather than by attacking cryptography at all.

The economic constraint deserves its own paragraph because it is the control most people actually experience. Contactless transactions are bounded by the strong customer authentication regime. In the European Economic Area, Article 11 of the PSD2 regulatory technical standards permits an exemption for contactless payments at point of sale up to €50 per transaction, with cumulative limits of €150 or five consecutive transactions since the last application of strong customer authentication. In the United Kingdom, the corresponding Article 11 of the UK SCA-RTS was set at £100 per transaction and £300 cumulative, following the Financial Conduct Authority's policy statement PS21/2, which raised those figures from £45 and £130.

That position changed recently and the change is the kind of thing this volume exists to track. On 19 December 2025 the Financial Conduct Authority announced, through Handbook Notice 136, rule changes taking effect in March 2026 which allow firms with strong fraud controls to set their own contactless limits rather than being bound by the regulatory figures. The FCA stated it did not expect immediate changes to the £100 single-transaction limit, and that firms remain obliged to reimburse consumers for unauthorised transactions and to communicate limit changes under the Consumer Duty. As of writing in August 2026, the £100 and £300 figures therefore persist as industry practice rather than as a hard regulatory cap in the United Kingdom, and any reader relying on this paragraph after 2026 must re-check it.

For a relay attacker, those limits are the binding constraint. A relayed contactless transaction below the floor limit yields a small amount, exhausts the cumulative counter quickly, and then triggers a PIN or biometric prompt that the relay cannot satisfy, because a device-based wallet requires consumer device cardholder verification on the genuine device.

■ Breaches at merchants and processors: the dominant modern source

Everything above concerns single cards. Breaches concern lists, and lists are what the economics reward.

There are two families and they need separating.

Server-side compromise means card data taken from inside a merchant's or processor's environment: from a database, from a log file, from application memory, or from network traffic. The canonical era was point-of-sale memory-scraping malware, which exploited the window in which track data existed in the clear in a till's memory between being read and being encrypted. TJX in 2007, Heartland Payment Systems disclosed in January 2009, Target in 2013 and Home Depot in 2014 are the reference cases; Target alone was up to 40 million payment card records taken between 27 November and 15 December 2013, using credentials obtained from a supplier and BlackPOS malware deployed on till systems.

The structural answer to that family is point-to-point encryption. If the reading device encrypts card data inside its own secure hardware, using keys the merchant never holds, then the merchant's environment never contains card data in a usable form and the memory-scraping window closes. This is why PCI Point-to-Point Encryption solutions are built on PCI PTS POI devices with the Secure Reading and Exchange of Data capability, and why a validated P2PE deployment collapses a merchant's PCI DSS scope so dramatically. It is also why tokenisation matters: replacing stored PANs with tokens means a stolen database of tokens is a stolen database of nothing.

Client-side compromise, usually called e-skimming or by the threat-actor label Magecart, is the modern replacement. The attacker does not enter the merchant's environment at all. They compromise a

script that the merchant's checkout page loads into the customer's browser, often a third-party script for analytics, chat or tag management, and that script copies the card fields as they are typed and sends them to an attacker-controlled endpoint. The merchant's servers never see the theft. The merchant's own logs are clean. The card data is taken between the customer's keyboard and the merchant's form handler.

British Airways in 2018 is the reference case in the United Kingdom, and the regulatory outcome is instructive. The Information Commissioner's Office issued a notice of intent on 4 July 2019 proposing a penalty of £183.39 million, and after representations issued a penalty notice on 16 October 2020 for £20 million, finding infringements of Articles 5(1)(f) and 32 of the General Data Protection Regulation as it then applied. The attacker was in a position to access the personal data of approximately 429,612 customers and staff, including the payment card details of around 244,000 customers. The reduction from the notice of intent to the final penalty is itself worth noting by anyone modelling regulatory risk: it was not a finding that the breach was less serious, but the product of representations and of the Commissioner's five-step penalty methodology.

The controls that address e-skimming are now explicit in the standard rather than implied. PCI DSS version 4.0.1 is the current version of the standard as of writing in August 2026. Its future-dated requirements ceased to be best practice and became mandatory on 31 March 2025. Two of them exist specifically for this attack:

Requirement	What it obliges	Status as of writing, August 2026
6.4.3	Maintain an inventory of all scripts executing on payment pages, authorise each one, justify its necessity, and assure its integrity	Mandatory since 31 March 2025
11.6.1	Deploy a change and tamper detection mechanism that alerts on unauthorised modification of the HTTP headers and content of payment pages as received by the consumer browser	Mandatory since 31 March 2025

The PCI Security Standards Council published an information supplement, "Payment Page Security and Preventing E-Skimming — Guidance for PCI DSS Requirements 6.4.3 and 11.6.1", on 10 March 2025, which is the primary reference for how to satisfy them.

The rest of PCI DSS matters here too and is frequently misremembered, so it is worth stating plainly. Sensitive authentication data — full track data from stripe or chip, the CVV2 family of codes, and PINs or PIN blocks — must not be retained after authorisation, by anyone, under any circumstances, even encrypted. The PAN may be stored but must be rendered unreadable wherever it is stored. Those two rules between them determine what a breach can actually yield, and they are the reason a well-run breached merchant loses PANs and expiry dates while a badly run one loses everything needed to transact.

Two reporting obligations attach in the United Kingdom, and they run on different clocks. Under UK GDPR Article 33, a personal data breach must be notified to the Information Commissioner's Office without undue delay and, where feasible, not later than 72 hours after becoming aware of it. Separately, under the Payment Services Regulations 2017, a payment service provider must report major

operational or security incidents to the Financial Conduct Authority. Card scheme rules add a third: an account data compromise event triggers scheme notification and, above thresholds set by each scheme, a PCI Forensic Investigator engagement. All three are live as of writing in August 2026.

■ Card testing and enumeration attacks

The last category is aimed at issuers, executed through merchants, and paid for by both.

The mechanism is straightforward. An attacker sends authorisation requests through a merchant's payment endpoint using candidate card credentials and reads the response. An approval, or in some configurations a decline code that distinguishes "wrong CVV2" from "no such account", tells the attacker which candidates are live. The merchant is not the target. The merchant is the oracle.

The arithmetic explains why this is worth doing and where the defences must sit. Under ISO/IEC 7812:2017 an issuer identification number may be eight digits within a PAN of up to nineteen, and the final digit is a Luhn check digit. For a sixteen-digit PAN under an eight-digit IIN, that leaves seven free digits: ten million candidate accounts, of which some fraction is issued. Expiry dates within a typical five-year validity window give around sixty combinations. CVV2 gives one thousand. Guessing a complete credential from nothing is not economic. Completing a partial credential is: given a PAN from a breach, expiry and CVV2 together are sixty thousand combinations, and an attacker who can make unlimited attempts will find them.

That "unlimited" is the vulnerability, and it was documented precisely. Ali, Arief, Emms and van Moorsel of Newcastle University published "Does the Online Card Payment Landscape Unwittingly Facilitate Fraud?" in IEEE Security and Privacy in 2017, describing what they called the distributed guessing attack. Their finding was structural rather than cryptographic: because merchants require different combinations of fields, and because no party was counting failed attempts across merchants, guesses could be spread across many sites so that no single site saw enough failures to react. They reported that Mastercard's centralised network detected the pattern within fewer than ten attempts even when distributed, and that Visa's did not. The fix is not at the merchant. It is centralised counting.

Enumeration has a variant that is particularly awkward: account verification messages, zero-value or nominal-value authorisations whose purpose is to check a card is live. They produce no clearing record, no settlement and no cardholder-visible charge, which makes them ideal for testing and invisible to the person whose card is being tested.

The scheme response, and the figures a practitioner will be measured against, is the Visa Acquirer Monitoring Program. Visa's own published material states that VAMP took effect on 1 April 2025 with an advisory period from 1 April 2025 to 30 September 2025, that it consolidates five existing fraud and dispute programmes into a single acquirer programme and thirty-eight remediation processes into one, that it incorporates enumeration criteria based on the count of enumerated authorisation transactions and an enumeration rate, and that enumeration attacks account for \$1.1 billion annually in fraud losses.

The numeric thresholds do not have to be taken on trust from acquirers and risk vendors, because Visa publishes them itself. Its fact sheet "Visa Acquirer Monitoring Program Overview", hosted on the corporate Visa site, sets them out as follows, and they stand as below at the time of writing in August 2026:

Measure	Definition	Threshold
VAMP ratio	$(\text{TC40 fraud reports} + \text{TC15 disputes}) \div \text{TC05 settled transactions, card-not-present}$	Acquirer: 50 bps Above Standard, 70 bps Excessive
VAMP ratio, merchant	As above	Excessive, reduced from 220 bps to 150 bps from 1 April 2026 in the United States, Canada, the European Union and Asia-Pacific
Minimum count	Combined fraud and dispute events per month to enter the programme	1,500, on a different basis in the CEMEA region
Enumeration ratio	$\text{Enumerated authorisations, approved and declined} \div \text{total authorisations}$	2,000 bps, that is 20 per cent
Minimum enumeration count	Enumerated authorisations per month to enter the programme	300,000

Two features of that table are worth dwelling on. The enumeration ratio counts declined authorisations in both numerator and denominator, which is correct, because a card-testing run is overwhelmingly declines and a ratio built only on approvals would be blind to it. And the VAMP ratio counts a TC40 fraud report and a TC15 dispute arising from the same transaction as two events, so one bad transaction can move the ratio twice.

Defences against enumeration are layered, and each layer catches a different thing.

At the merchant, the controls are rate limiting by internet protocol address, device and session, bot management, a minimum transaction amount, removal of free-text amount fields on donation pages, cardholder verification challenges on low-value flows, and monitoring of the approval ratio and the mix of decline codes as a leading indicator. A merchant whose approval rate collapses overnight is usually being used as an oracle, not losing customers.

At the gateway and acquirer, the controls are velocity rules by card, by BIN range and by merchant, pre-authorisation screening that never lets the request reach the network, and specific treatment of zero-value account verification traffic.

At the issuer, the controls are CVV2 attempt counters that block a card after a small number of failures, velocity monitoring across the issuer's own BIN ranges regardless of which merchant the attempts arrive from, and deliberate parsimony in decline codes. This last point connects directly to the response codes discussed in Volume III: "do not honour" is uninformative on purpose. Every extra bit of information in a decline is a bit of information for an enumeration attacker, and the industry's decision to be vague is a security decision, not laziness.

At the scheme, the controls are network-level enumeration detection of the kind Visa markets as Account Attack Intelligence, the monitoring programmes above, and financial penalties on acquirers that transmit the traffic.

■ The threat model in one table

Attack	What is obtained	Primary control	Where the control lives
Stripe skimming	Full track data, replayable	EMV chip; stripe withdrawal; CVV1 validation	Card and issuer
Shimming	PAN, expiry, track-2 equivalent with iCVV, one spent cryptogram	iCVV and service code validation; ATC monitoring; PTS POI tamper resistance	Issuer and device
Contactless relay	A genuine cryptogram, at a distance	RRP timing bound; SCA contactless limits; CDCVM on wallets	Terminal, kernel and regulator
Merchant or processor breach	Bulk PAN and expiry, sometimes more	P2PE; tokenisation; PCI DSS storage rules	Merchant and acquirer
E-skimming	Full card-not-present credential, at scale	PCI DSS 6.4.3 and 11.6.1; script control	Merchant
Enumeration and card testing	Validated credentials from partial data	Centralised velocity counting; CVV2 attempt limits; scheme monitoring	Issuer and scheme

The pattern in the right-hand column is the argument of this chapter. Every control that worked, worked because it was placed where the attacker could not choose to avoid it. EMV worked because the card is in every transaction. P2PE works because the reading device is in every transaction. Centralised velocity counting works because the network is in every transaction. Every control that failed, failed because it was placed at a merchant, and there are millions of merchants, and the attacker picks which one to use.

■ What the law does when it fails

The consumer position in the United Kingdom is governed by the Payment Services Regulations 2017, and it is more favourable than most people assume. Under regulation 76, where a payment transaction was unauthorised, the payment service provider must refund the amount and restore the account to the state it would have been in, as soon as practicable and in any event no later than the end of the business day following the day on which it becomes aware of the unauthorised transaction, unless it has reasonable grounds to suspect fraudulent behaviour by the payment service user and notifies the appropriate person under section 333A(2) of the Proceeds of Crime Act 2002.

Under regulation 77, the payer may be made liable for a maximum of £35 for losses arising from the use of a lost, stolen or misappropriated payment instrument, and even that does not apply where the loss was not detectable by the payer before the payment, where it resulted from an act of the provider's employee or agent, where the payer had already notified the provider, or where the provider did not require strong customer authentication. The £35 figure and both regulations are as they stand at the time of writing in August 2026. In practice UK Finance reports that customers are fully refunded in more than 98 per cent of unauthorised fraud cases, which is why the losses in the table at the start of this chapter fall on banks and merchants rather than on cardholders.

One structural change is pending and worth flagging for anyone reading this after 2026. On 11 March 2025 the United Kingdom government announced its intention to abolish the Payment Systems Regulator and consolidate its functions primarily within the Financial Conduct Authority. As of writing in August 2026 the PSR continues to exist and to exercise its functions, and the legislation required to complete the consolidation is awaited. The National Payments Vision has also committed the government to revoking the strong customer authentication provisions from the Payment Services Regula-

tions 2017 in favour of an outcomes-based approach, which, if delivered, will change the regulatory footing of several paragraphs in this chapter.

That is the honest state of this material. The mechanisms in this chapter — replay, interception, relay, bulk theft, enumeration — are permanent, because they follow from what a credential is. The numbers, the caps, the thresholds and the regulators are not, and this chapter will need rewriting before the rest of the book does.

45.98 Common wrong ideas

Wrong: The chip ended counterfeit card fraud. **Right:** The chip is sound, but its deployment was uneven and its fallback paths exploitable; on United States dual-message networks the counterfeit fraud rate was 9.0 basis points in 2015, fell to 6.6 in 2017, and was back to 9.1 by 2021.

Wrong: The chip keeps your card number secret. **Right:** The chip never reveals its key, but it reveals the primary account number, the expiry date and track-equivalent data in the clear on every transaction; what it removed was replay at a terminal, not disclosure.

Wrong: Most stolen card details come from stolen or copied cards. **Right:** The dominant source is a breached merchant, processor or checkout page that touches no physical card at all, and a fast-growing share involves no card-data theft whatsoever, only a telephoned customer reading out a one-time pass-code so the criminal can load the card into their own wallet.

Wrong: Relay attacks are the main contactless threat. **Right:** A relay borrows genuine cryptography rather than breaking it, which makes it interesting to researchers and lawyers; the £46.8 million of contactless fraud reported for 2025 is dominated by ordinary use of ordinary stolen cards.

Wrong: The Relay Resistance Protocol proves the card was physically at the terminal. **Right:** It is a latency bound with a granularity of hundreds of microseconds, and one hundred microseconds is thirty kilometres of light in free space, so it cannot localise a card geometrically at all.

Wrong: “Card data” is a single asset to be protected. **Right:** A PAN alone, a PAN with expiry, a PAN with expiry and CVV2, full track data, a PIN block and a network token are six different things with six different resale values, attack paths and controls, and merging them is the commonest analytical error in the field.

Wrong: A successful shim-derived clone proves EMV is broken. **Right:** It proves an issuer failed to validate iCVV, validated the wrong verification value, or ignored the service code; the defect sits in a small number of authorisation stacks and is closable by configuration.

Wrong: Card testing is a merchant problem, fixed by rate limiting at the merchant. **Right:** The merchant is the oracle rather than the target, and because attempts can be spread across many merchants the only counting that works is centralised at the network or the issuer.

Wrong: Decline codes should tell the merchant precisely what went wrong. **Right:** Every extra bit of detail in a decline is a bit of information for an enumeration attacker, so the vagueness of “do not honour” is a deliberate security decision rather than laziness.

Wrong: Cardholders end up carrying the losses in the fraud tables. **Right:** Regulation 76 requires refund by the end of the following business day, regulation 77 caps the payer at £35 and disappplies even that in most cases, and more than 98 per cent of unauthorised fraud cases are fully refunded, so the losses fall on banks and merchants.

45.99 Chapter summary in 20 lines

1. Card fraud is not one crime but at least five, and they have almost nothing in common technically.
2. The magnetic stripe is a bearer credential in the purest sense: static data that a reader learns in full and a writer can reproduce exactly.
3. The chip replaced that key with a locksmith — a secret it never discloses, a terminal nonce and a counter that ticks on every transaction — so every reply is fresh.
4. A captured cryptogram is therefore a used ticket, which is why copying a chip conversation does not produce a working counterfeit card.
5. Skimming is close to dead where chips are used properly, with EMVCo reporting 15.6 billion cards and 97 per cent of card-present transactions on chip as of the fourth quarter of 2025.
6. A shim can read the chip dialogue because that dialogue was never confidential; authenticity, not secrecy, was its security property.
7. What a shim yields is the account number, the expiry and track-2 equivalent data carrying iCVV, which is useful on a website and useless at a terminal.
8. EMV-bypass cloning succeeds only against issuers that fail to check iCVV or the service code, so the weakness lies in a few authorisation stacks rather than in EMV.
9. A relay attack borrows genuine cryptography rather than breaking it, which makes it unfalsifiable from the issuer's side and famous out of all proportion to its losses.
10. Three things constrain relays: centimetre-range radio, a transaction time budget commonly cited at around 500 milliseconds, and EMV's Relay Resistance Protocol.
11. That protocol is a latency bound and not academic distance bounding, and its granularity leaves the acceptance window wider than a security engineer would like.
12. In practice the binding constraint on contactless is the authentication regime, with the United Kingdom's £100 and £300 figures now industry practice rather than a hard cap after Handbook Notice 136.
13. The dominant modern source of card data is the breach: Target lost up to 40 million payment card records in three weeks of 2013 with no skimmer, shim or radio involved.
14. Point-to-point encryption and tokenisation answer server-side theft by making the merchant's environment hold nothing worth taking.
15. E-skimming moved the theft into the shopper's browser, where the merchant's servers never see it, and British Airways in 2018 is the United Kingdom reference case, penalised £20 million on 16 October 2020.
16. PCI DSS requirements 6.4.3 and 11.6.1, mandatory since 31 March 2025, exist for exactly that attack, while the storage rules on sensitive authentication data determine what any breach can yield.
17. Enumeration completes partial credentials rather than guessing whole ones: sixty expiry dates against a thousand codes is sixty thousand tries, which only unlimited attempts make economic.
18. The merchant is the oracle rather than the target, and the documented fix is centralised counting at the network, as the Newcastle work on distributed guessing showed in 2017.
19. The Visa Acquirer Monitoring Program supplies the discipline, with a combined fraud and dispute ratio, an enumeration ratio of 2,000 basis points that counts declines on both sides, and penalties landing on acquirers.
20. Every control that worked was placed where the attacker could not choose to avoid it — the card, the reading device, the network — and every control that failed was placed at a merchant, and the attacker picks the merchant.

Sources: UK Finance Annual Fraud Report 2026 and 2025 press releases; Home Office call for evidence on unauthorised fraud, July 2026; The Nilson Report, January 2026; EMVCo worldwide deployment statistics, Q4

2025; Federal Reserve Bank of Kansas City payments research briefing, February 2025; PCI Security Standards Council document library and e-skimming information supplement, March 2025; Visa corporate publications on the Acquirer Monitoring Program; FCA press release of 19 December 2025 and Handbook Notice 136; Payment Services Regulations 2017 regulations 76 and 77 as published on legislation.gov.uk; ICO penalty notice against British Airways, October 2020; Mastercard magnetic stripe announcement, August 2021; EMV Contactless Book C-2 and CPACE terminal kernel specification for the Relay Resistance Protocol; Radu et al., IEEE S&P 2022; Ali et al., IEEE Security and Privacy, 2017; Krebs on Security and RCMP material on shimmers, January 2017.

Card Not Present

46.0 What this chapter gives you

1. You will be able to explain why a chip transaction cannot be replayed while an online one is nothing but replayable facts, and why that single difference decides the rest of the chapter.
2. You will be able to say what “card not present” actually means — a combination of indicators in an authorisation message, not a fact about where the customer was standing.
3. You will be able to describe exactly what the address verification service compares, and why a criminal holding a delivery label defeats it without ever knowing the street name.
4. You will be able to read the address verification response codes and explain why G, U, S and R mean the merchant learned nothing, and what a decline-on-mismatch rule costs in conversion.
5. You will be able to explain why the shop rather than the bank carries an online fraud loss by default, and why the acquirer is the party of last resort behind it.
6. You will be able to state the effect of regulation 77(4)(d) of the Payment Services Regulations 2017 and say why a United Kingdom consumer’s exposure to unauthorised online card fraud is zero rather than £35.
7. You will be able to explain why quoting £328.4 million of 2008 card-not-present fraud without the £41.2 billion of online spending behind it is campaigning rather than reporting.
8. You will be able to map the three e-commerce indicator outcomes to their liability positions, and name the three limits that cause most disappointment about the liability shift.
9. You will be able to say why first-party misuse is invisible at the message layer, and what Compelling Evidence 3.0 does instead of arguing about it.
10. You will be able to explain why a billing descriptor that does not resemble the trading name is a fraud problem the merchant manufactured itself, and the cheapest one on the list to fix.

The plain version

Imagine a members’ club with a very good doorman.

When you arrive in person you hand him your membership card. The card has a small computer inside it. The doorman’s reader asks the card a question, the card does a sum using a secret number that only it and the club’s office know, and hands back an answer. The answer is different every single time. If somebody photographed everything about your card and turned up tomorrow with a perfect copy of the picture, they would still be stuck at the door, because the doorman would ask his question and a picture cannot do a sum.

Now imagine the same club takes bookings by post. You write a letter containing your membership number, the date your membership runs out, and the three-digit code on the back of your card. The club opens it, checks all three against its file, and books you a table.

Everything about that letter is copyable. The membership number does not change. The expiry date does not change. The three-digit code does not change. Anybody who has seen your card, or received one of your letters, or broken into a filing cabinet full of other people's letters, can write out an identical letter and send it in. The club will read it and it will agree with itself, because all the club can do is check the writing on the page against the writing in its file.

That is the whole problem of buying things on the internet. In a shop the chip does a fresh sum every time. Online there is no chip, no reader, no sum. There is only a page of facts about a card, and facts can be copied.

Nadia buys a pair of running shoes for £48.99 from a British website. The checkout page asks her for five things: the sixteen digits on the front of her card, the expiry date, the three-digit code on the back, her name, and her billing address, which is Flat 4, 27 Bramble Road, Leeds, LS6 2QD.

The shop sends all of that to its bank, which sends it to Nadia's bank, which does three checks and answers all three in the same breath. Is there £48.99 available and is the card in good standing? Yes. Does the three-digit code match the one we generated? Yes. Does the address match the one we hold on file?

That third check is much weaker than it sounds. The bank does not compare the words. It does not know or care that the street is Bramble Road or the town is Leeds. It pulls out the numbers, the house number and the digits of the postcode, compares them with the digits it holds, and reports whether both agree, one does, or neither does.

So a criminal who has stolen Nadia's card details does not need to know she lives on Bramble Road. He needs a house number and a postcode, the two least secret things about anybody. They are printed on delivery labels, typed into hundreds of websites, and usually stolen in the same breath as the card number, because the same shopping basket that took the card number also took the address to send the shoes to.

Now the part that decides everything else: who loses the money.

Suppose the transaction was not Nadia. Somebody in another country bought £48.99 of shoes with her card. She sees the line on her statement, does not recognise it, and rings her bank, which puts the £48.99 back into her account. That is not a favour; in the United Kingdom it is the law, and it has to happen quickly. Her bank then tells the shop's bank it wants the money back. The shop's bank takes the £48.99 out of the shop's account and charges the shop a fee for the trouble on top. So the shop has lost the shoes, the £48.99, the postage and the fee. The bank has lost nothing. Nadia has lost nothing.

That asymmetry is the single most important fact in this chapter. In a shop, when the chip does its sum properly, the bank carries the fraud. Online, by default, the shop carries it. Which means everything a merchant does about online fraud, every awkward extra screen, every declined order, every "please confirm this purchase in your banking app", is a merchant spending its own money to protect its own money.

For decades the easiest way to steal with a card in Britain was to copy the magnetic stripe and make a fake card, or to steal a real card and forge a signature. Then, over 2004, 2005 and 2006, the country replaced signatures with chips and PINs. It worked: between 2004 and 2008 the value of counterfeit card fraud committed in the United Kingdom fell by sixty-eight per cent. The criminals did not retire. They moved. Over the same stretch card-not-present fraud on British cards rose from £72.9 million in 2000 to £328.4 million in 2008, and by 2008 total losses on British cards hit £609.9 million, the highest since 1998, of which fifty-four per cent was card-not-present and only twenty-eight per cent counterfeit.

Squeeze a balloon at one end and the air goes to the other. Security measures rarely destroy fraud. They relocate it, towards whichever door is now cheapest to open.

Finally, there is a kind of card-not-present loss with no criminal in it, and it is the one merchants find hardest to talk about. Sometimes the person who says “I did not buy that” really did buy it. Sometimes a twelve-year-old bought it on a parent’s card. Sometimes a subscription renewed and the customer had forgotten signing up. Sometimes the statement line reads “SP*ORD9931” and the customer genuinely does not recognise their own purchase. And sometimes somebody has worked out that saying “I did not buy that” is a quick way to get free shoes.

At the level of the message that arrives at the shop, every one of those looks identical. They arrive as the same dispute, with the same code, saying the same thing. The industry calls this friendly fraud, or first-party misuse, and nothing in the payment system separates it from the real thing. That is the awkward part.

Where the plain version stops being true

“Card not present” is not a description of the shopper. It is a set of codes in a message. Whether a card is present is not a physical fact about the room. It is an assertion the acquiring side puts into the authorisation, and the issuer prices, risk-scores and disputes the transaction on the basis of that assertion rather than on reality. A customer can stand at a counter with the card in their hand and still generate a card-not-present transaction, because the shop typed the number into a virtual terminal or the chip reader was broken. Conversely, a subscription renewal at three in the morning with no human near it is card-not-present in the same category as a live checkout, yet treated differently again because it carries an indicator saying the merchant initiated it. The question is never “was the card there?” It is always “what did the message claim, and does the issuer believe it?”

Address verification does not verify an address, and often does not run at all. Two things are hidden. The first is that the check is numeric and partial, so it is defeated by knowing two facts that are barely secret. The second, and more practically important, is coverage: it depends on the issuer supporting it and on the acquiring chain passing the data through. British and American issuers generally do; a great many issuers elsewhere do not, and return a code meaning “not checked” or “not supported”. A merchant selling internationally gets a real answer on some of its traffic and a shrug on the rest, so a rule that declines on mismatch will refuse many honest foreign customers and almost no informed criminals. There is also the mundane failure mode: a cardholder who moved house and never told their bank fails the check on every purchase for years, while being entirely genuine.

The liability shift did not move money. It moved a right to argue, and it did not move it to a bank. What the shift changes is who carries the risk of one category of dispute, under scheme rules that are private contracts, not law. The party carrying the merchant’s risk in practice is the acquirer, because the acquirer is the one the scheme takes the money from. If a merchant with a large fraud problem goes out of business, the chargebacks keep arriving for months and the acquirer pays them from its own capital. That single fact explains behaviour that looks arbitrary from the merchant’s side: rolling reserves, delayed settlement for new merchants, onboarding refusals for high-risk sectors, and the acquirer’s intense interest in a dispute ratio. The acquirer is not being difficult. It is the party of last resort.

“Friendly fraud” is a commercial term, not a legal one, and a meaningful share of it is the merchant’s own fault. No regulation anywhere contains the phrase. In United Kingdom law the customer’s claim is simply that a transaction was unauthorised, and the burden of proving otherwise sits by statute on the payment service provider. The law goes further than most merchants realise: the fact that the

instrument was used, as recorded by the bank's own systems, is expressly stated not to be sufficient in itself to prove the customer authorised it. The merchant who says "the transaction went through, so obviously they authorised it" is making an argument the legislation has already anticipated and rejected. Separately, an unglamorous share of these disputes are caused by the merchant: a descriptor that does not resemble the trading name, a free trial that converts silently, a delivery date that slipped. That is not fraud in any direction. It is bad merchandising producing a fraud-shaped message.

One caution applies to every number below. Published fraud figures are not comparable with each other and often not with themselves. Some count cards or accounts rather than people; some count issuer reports of suspected fraud, which are not confirmed loss; and year-on-year percentages are routinely computed against restated prior-year figures. Every figure below therefore carries its source and its date.

The technical version

■ What the message actually says

A card-not-present transaction is not a distinct protocol. It is an ordinary authorisation request carrying a particular combination of indicators, and everything downstream, from the interchange rate to the dispute rights, keys off them.

The primary indicator is the point-of-service entry mode, carried in ISO 8583 data element 22. The values that matter here are those asserting the card details were not read from the card: manually keyed entry, and electronic commerce. Acquirer specifications commonly assign 01 to manual key entry, 05 to chip, 07 to contactless chip and 81 to electronic commerce, but this is precisely the area where, as Chapter 24 set out, the standard fixes the box number and the parties negotiate the contents. An integrator must read the specification in front of them rather than a table on the internet.

Alongside it sit several other assertions, each changing the commercial and legal treatment:

Assertion	What it claims	What it changes
Entry mode	How the card data reached the terminal	Interchange, fraud scoring, dispute rights
E-commerce indicator (ECI)	Whether the cardholder was authenticated, and to what degree	Whether fraud liability sits with issuer or merchant
MOTO indicator	The order was taken by mail or telephone	Excludes most authentication routes entirely
CIT/MIT flag	Whether the cardholder or the merchant initiated this specific transaction	Whether authentication was required at all
Recurring/instalment indicator	Part of an agreed series	Availability of an authentication exemption

A merchant's fraud position is therefore decided partly by the accuracy of its own integration. One that flags genuine cardholder-initiated e-commerce as merchant-initiated, or sends recurring transactions without the series indicator, will find a liability shift it believed it had does not exist when a dispute arrives. That is among the most common and least discussed causes of lost disputes.

■ What the merchant can check, and what each check is actually worth

The Luhn check on the primary account number, described in Chapter 42, involves no network at all: it proves only that the digits are internally consistent, and catches typing errors. The expiry date is likewise static and printed on the card. The issuer identification number is more useful, because it names the issuing institution, country and product, and the mix of card types in an order stream is stable for a real business and unstable for an attacked one.

The card verification code printed on the signature panel, called CVV2 by Visa and CVC2 by MasterCard, is the one static check with real teeth, and the reason is archival rather than cryptographic. The code is not in the magnetic stripe and not in the chip, and PCI DSS prohibits merchants and processors from retaining it after authorisation. A criminal who obtained card data from a merchant's stored records should therefore not have it; one who copied the physical card, or persuaded the cardholder to read it out, will. The code does not distinguish honest from dishonest. It distinguishes one class of data theft from another, and it is worth exactly that much.

Everything beyond address verification, treated below, is inference rather than verification: network address, device, email address, shipping address, time, basket composition, the velocity of orders sharing any of those, and how it all compares with the merchant's own history. That is the domain of Chapter 51.

■ Address verification and its limits

The address verification service is an issuer-side check invoked as part of authorisation. The acquiring side sends the billing address supplied at checkout; the issuer compares it with the address on the account and returns a single-character result code alongside the authorisation response.

The comparison is not textual. It compares the numeric portion of the street address and the numeric portion of the postcode or ZIP code. Street, town and county names play no part. That was a deliberate decision from an era of noisy address data, where "27 Bramble Rd" and "Flat 4, 27 Bramble Road" had to be treated as the same place, and it has never been revisited.

Each scheme publishes its own code set and they do not agree. The following are the codes a British merchant will meet most often, as published by the schemes and by acquirers at the time of writing in August 2026:

Code	Meaning
Y	Address and postcode both match (five-digit ZIP in the United States)
A	Address matches, postcode does not
Z	Postcode matches, address does not
N	Neither matches
F	Visa only: address and postcode both match, United Kingdom domestic
G	Non-participating issuer outside the United States; not verified
U	Issuer holds no data or the service is unavailable
S	Issuer does not support address verification
R	System unavailable, retry

Four properties of this list matter more than the list itself.

First, most of the codes are not answers. G, U, S and R all mean the merchant learned nothing. A rule expressed as “decline on anything other than a full match” therefore declines a large fraction of legitimate international orders, and merchants who write that rule discover it in their conversion figures rather than their fraud figures.

Second, the result arrives with, or after, the authorisation decision, not before it. The issuer may approve a transaction whose address check failed completely, because the two decisions are separate: approval is about funds and card status, the address code is advice. A merchant that rejects on the strength of the code must then reverse or void the authorisation, and if it does not it leaves a hold on a genuine customer’s balance for days. Failing to reverse abandoned authorisations is among the most common defects in a young integration, and it generates complaints that look like fraud disputes and are not.

Third, the check is defeated by the same breach that supplies the card number, because retail data theft almost always yields the billing address in the same record. It is not a defence against an attacker holding a full record; it is a defence against one holding a partial record, which is a real and common situation.

Fourth, it does nothing for digital goods, where there is no shipping address, and is easily neutralised for physical goods by shipping somewhere other than the billing address, which honest customers do constantly. The correct use is therefore not as a gate but as one weighted input, where a mismatch on a high-value first order from an unrecognised device is worth a great deal and a mismatch on a small repeat order from a two-year customer is worth almost nothing.

■ Who bears the loss, in law and in practice

In the United Kingdom, the statutory position for consumer card payments is set by the Payment Services Regulations 2017 (S.I. 2017/752), and it is far more favourable to the cardholder than most people working in payments assume. What follows describes the regulations as in force at the time of writing in August 2026; legislation.gov.uk recorded regulation 77 as up to date to 14 August 2026.

Regulation 76(1) requires the payment service provider to refund the amount of an unauthorised transaction and restore the account. Regulation 76(2) requires that refund “as soon as practicable, and in any event no later than the end of the business day following the day on which it becomes aware of the unauthorised transaction”, with a single exception in 76(3) where the provider suspects fraud by its own customer and reports it in writing under the Proceeds of Crime Act 2002.

Regulation 77(1) allows the provider to hold the payer liable for up to £35 of losses arising from a lost, stolen or misappropriated instrument, and 77(3) removes all protection where the payer has acted fraudulently or has failed, with intent or gross negligence, to keep credentials safe.

Then comes regulation 77(4), where card-not-present becomes a special case. Except where the payer has acted fraudulently, the payer is not liable for any losses at all in four situations. Two concern notification. The third is where strong customer authentication was required but the payer’s own provider did not apply it. The fourth, at regulation 77(4)(d), is where “the payment instrument has been used in connection with a distance contract”, subject to narrow exceptions.

An ordinary online purchase is a distance contract. The practical effect is that for a consumer card used to buy something online in the United Kingdom, the £35 excess does not apply at all. The cardholder’s exposure to unauthorised online card fraud is, absent their own fraud, zero. Regulation 77(6) completes the circuit: where strong customer authentication was required but the payee or the payee’s provider did not accept it, that party must compensate the payer’s provider for what it paid out.

Two evidential provisions finish the picture. Under regulation 75(1) the burden of proof is on the payment service provider to show the transaction was authenticated, accurately recorded and unaffected by technical failure. Under 75(3) the use of the instrument as recorded by the provider “is not in itself necessarily sufficient to prove” either that the payer authorised it or that the payer was fraudulent or grossly negligent, and under 75(4) a provider alleging either must give the payer supporting evidence.

Read together, these provisions leave the issuer under a next-business-day refund obligation with a burden of proof it usually cannot discharge cheaply, and therefore with one economically rational response: refund the customer and push the loss down the chain by raising a chargeback. The scheme rules then determine whether the acquirer, and through it the merchant, can push it back. That is how a consumer protection statute becomes a merchant cost line, and it is why the merchant’s true cost of a single fraudulent order is never the order value alone. It is the order value, plus goods and shipping already dispatched, plus the chargeback fee, plus the labour of responding, plus that dispute’s contribution to the ratios described next, plus, over time, the controls bought to prevent it and the revenue lost to honest customers those controls turned away.

■ The price of the channel

Card-not-present acceptance is also more expensive to run before any fraud occurs, and the clearest recent United Kingdom illustration is a matter of public regulatory record. Following the United Kingdom’s withdrawal from the European Union, Mastercard and Visa increased interchange fees for card-not-present transactions between the United Kingdom and the European Economic Area on consumer cards, from 0.2 per cent to 1.15 per cent for debit and from 0.3 per cent to 1.5 per cent for credit. The Payment Systems Regulator opened a market review, and its final report concluded that the two schemes were not subject to effective competitive constraint, that the increases were made without regard to the effect on businesses, that no justification for them was identified, and that they were costing United Kingdom businesses an additional £150 million to £200 million a year.

The state of that remedy at the time of writing is worth recording precisely, because it will not stay true. As of the regulator’s page last updated in October 2025, it had decided not to proceed with the interim cap it had consulted on, citing litigation about its powers, and had instead launched a consultation on the methodology for a longer-term cap. No cap was in force. Separately, in March 2025 the government announced its intention to abolish the Payment Systems Regulator and consolidate its functions primarily into the Financial Conduct Authority, with a Treasury consultation closing in October 2025 and legislation expected when parliamentary time allows. Any reader in a later year must check both.

The principle underneath survives regulatory change: card-not-present interchange is higher than card-present interchange, in every scheme and region, because the issuer carries a higher expected fraud cost and prices accordingly. Anything that reduces that cost, above all authentication, tends to be rewarded with a lower rate.

■ Scheme monitoring: the thresholds that actually discipline a merchant

Statutory liability determines who pays for a given transaction. Scheme monitoring programmes determine whether a merchant may keep trading at all, and they shape day-to-day behaviour in a fraud team. Visa consolidated its previous Visa Dispute Monitoring Program and Visa Fraud Monitoring Program into a single Visa Acquirer Monitoring Program, VAMP, effective 1 April 2025. It combines fraud reports and disputes into one ratio rather than tracking them separately, and applies at both acquirer portfolio and individual merchant level.

The thresholds below are as reported by acquirers, processors and the Merchant Risk Council at the time of writing in August 2026, and are marked as such deliberately: the Visa Rules are a private document distributed to participants and not published, so no figure here can be cited to a primary public source, and the thresholds have already changed twice since launch.

Level	Ratio threshold	Minimum volume to enter the programme
Acquirer, above standard	0.5 per cent	1,500 combined fraud and dispute events per month
Acquirer, excessive	0.7 per cent	1,500 combined fraud and dispute events per month
Merchant, excessive (from 1 April 2026)	1.5 per cent, reduced from 2.2 per cent	1,500 combined fraud and dispute events per month

The design consequence is worth stating plainly. Because fraud reports and disputes are combined into one ratio, a merchant with almost no fraud but many “goods not received” disputes can be caught by a fraud programme, and a single large merchant can drag an entire acquirer’s portfolio over a threshold. That is why acquirers monitor individual merchants far more closely than their dispute volumes would seem to justify, and why a merchant that quietly refunds every complaint to keep its ratio down may be acting rationally even while being taken advantage of.

Alongside this, PCI DSS v4.0.1, published in June 2024 and the active version of the standard at the time of writing, brought two requirements out of their future-dated grace period on 31 March 2025. Requirement 6.4.3 requires every script executing on a payment page to be authorised, integrity-assured and inventoried. Requirement 11.6.1 requires a mechanism alerting staff to unauthorised changes to payment page scripts and HTTP headers, evaluated at least weekly. Both exist because the modern route to bulk card data theft is not the merchant’s database but the shopper’s browser, and the control is inventory and integrity rather than detection after the fact, precisely because the merchant’s own server never sees the theft happen. Chapter 48 covers the standard in full.

■ Displacement: why the internet got worse as the shop got better

The British chip and PIN migration is the cleanest natural experiment in payment security, because it happened quickly, in one country, with published statistics either side. The figures given in the plain version come from the Office for National Statistics, drawing on data from the UK Cards Association. Of the £609.9 million lost on United Kingdom-issued cards in 2008, £379.7 million was incurred domestically and £230.1 million abroad; the composition was fifty-four per cent card-not-present, twenty-eight per cent counterfeit, nine per cent lost and stolen, eight per cent card identity theft and two per cent mail non-receipt.

Two refinements stop this being a simple morality tale.

The first is that displacement also happened geographically. Counterfeit fraud accounted for fifty-eight per cent of losses incurred abroad on British cards in 2008, because criminals took cloned British magnetic stripes to countries that had not yet migrated to chip. The stripe stayed on the card for years after the chip arrived, precisely so British cards would work abroad, and that compatibility was the vulnerability. Displacement follows the gradient of the weakest reader anywhere in the world that will accept your card.

The second is arithmetic honesty. Card-not-present fraud rose about three hundred and fifty per cent between 2000 and 2008, while the value of online shopping in the United Kingdom rose from £3.5

billion to £41.2 billion, more than tenfold. The fraud rate per pound spent therefore fell substantially while the absolute loss more than tripled. Both statements are true and only one is usually quoted. A book that reports the £328.4 million without the £41.2 billion is not reporting; it is campaigning.

The same pattern is visible today at European scale. The joint European Banking Authority and European Central Bank report on payment fraud published on 15 December 2025, covering data from 2022 to 2024, found card payment fraud seventeen times higher when the payee was outside the European Economic Area, where strong customer authentication is not legally required and frequently is not applied. Total reported payment fraud across the Area rose from €3.4 billion in 2022 to €4.2 billion in 2024, of which fraud on cards issued in the European Union and Economic Area was €1.329 billion, a twenty-nine per cent increase. The report concluded that strong customer authentication remains effective against the fraud types it was designed to defeat, while new types, chiefly the manipulation of legitimate payers into authenticating fraudulent transactions themselves, are rising.

That is confirmed on the British side. UK Finance's Annual Fraud Report 2026, published on 15 June 2026 and covering 2025, reported remote purchase card fraud losses of £423.5 million on 3.2 million cases, up three per cent by value and thirteen per cent by case count, and explicitly identified criminals compromising one-time passcodes in order to authenticate fraudulent transactions or to enrol stolen card credentials into digital wallets. The preceding edition, published on 28 May 2025, reported 2024 losses up eleven per cent to just under £400 million across nearly 2.6 million cases, and made the same observation. The two do not reconcile arithmetically, because figures are restated between editions, which is routine and legitimate and is exactly why this chapter dates every number. The 2025 edition is consistently reported, in the professional summaries written of it by firms such as TLT and circulated through Lexology, as having found that seventy-five per cent of United Kingdom e-commerce card-not-present fraud occurred at merchants acquired outside the United Kingdom, alongside £154.2 million of card fraud incurred abroad. UK Finance places the full report behind registration rather than publishing it openly, so that particular figure cannot be checked against the primary text, and it is given here on the authority of those summaries rather than of the report itself.

The mechanism to carry away does not date. Authentication closed the route that consisted of re-playing stolen static credentials. It did not close the route that consists of persuading the legitimate human to complete the authentication. The defence against the second is not a stronger credential; it is behavioural detection at the issuer, out-of-band confirmation for high-risk actions such as wallet enrolment, and consumer messaging that treats a one-time passcode as something never to be repeated to another person under any circumstances, including to somebody claiming to be the bank.

■ Liability shift, and what it did to merchant behaviour

There are two distinct liability shifts in cards and they are routinely confused. The first is the chip liability shift, in the card-present world: where a counterfeit card is used at a terminal that could not read a chip, the party that failed to support the technology bears the fraud. That is the shift that drove terminal replacement.

The second is the authentication liability shift, which applies in the card-not-present world and is the one that matters here. Where a merchant submits a transaction successfully authenticated through 3-D Secure, fraud liability moves from the merchant to the issuer, and the issuer loses the right to raise a chargeback under the fraud reason codes. The mechanism is the e-commerce indicator carried in the authorisation, together with the cryptographic authentication value the authentication produced. The mapping, as documented by acquirers including Worldpay at the time of writing in August 2026, is:

Visa ECI	Mastercard ECI	Meaning	Liability
05	02	Cardholder fully authenticated	Shifts to issuer
06	01	Authentication attempted; card not enrolled or issuer unavailable	Generally shifts, subject to scheme rules
07	00	Authentication not performed	Remains with merchant

Three limits on this shift are load-bearing and are the source of most disappointment.

It covers fraud disputes only. It does not cover “goods not received”, “goods not as described”, “credit not processed” or any other non-fraud condition. A merchant that authenticates every transaction still receives chargebacks, and for many merchants those are the majority.

It does not apply to data-only authentication, where risk data is passed to the issuer to improve approval rates without a full authentication taking place. That is a conversion optimisation, not a liability instrument, and is regularly mis-sold as both.

It does not apply where the merchant, or its acquirer, requested an exemption. If you asked not to authenticate, you cannot claim the protection authenticating would have given you. That is the central trade-off of the modern checkout, and it is a commercial decision made per transaction, at speed, by software.

The behavioural effect over two decades was what an economist would predict. Merchants who could adopt authentication cheaply did so and stopped caring about fraud disputes; merchants who could not, because they took orders by telephone or sold digital goods into markets with poor coverage, stayed exposed and built fraud teams instead. The shift did not reduce fraud by itself. It reallocated the cost of preventing it onto whichever party could prevent it most cheaply, which is what a well-designed liability rule is supposed to do.

Strong customer authentication changed the calculus again by making authentication a legal requirement rather than a competitive option. The rules, set in the Payment Services Regulations 2017 and the associated technical standards, applied from 14 September 2019, but the Financial Conduct Authority granted successive forbearance for card-not-present e-commerce: an industry plan to deliver by 14 March 2021, a six-month extension for the coronavirus crisis, and a further six months announced in 2021 taking the final date to 14 March 2022, which the Authority described as the latest at which it expected full compliance.

Once authentication became compulsory, competition moved to the exemptions, which Chapter 47 treats in full. One belongs here because it prices the whole channel. Under Article 18 of the technical standards, a provider may skip authentication on a remote transaction it identifies as low risk, provided the amount is below an exemption threshold value and its own fraud rate for that transaction type sits at or below a reference rate tied to that threshold. The Annex to Commission Delegated Regulation (EU) 2018/389 sets those out as follows:

Exemption threshold value	Reference fraud rate, remote card payments	Reference fraud rate, remote credit transfers
€500	0.01 per cent	0.005 per cent
€250	0.06 per cent	0.01 per cent

Exemption threshold value	Reference fraud rate, remote card payments	Reference fraud rate, remote credit transfers
€100	0.13 per cent	0.015 per cent

The United Kingdom version restates the same three-tier structure in sterling. Article 18 as it stands in the Financial Conduct Authority's Handbook points to a table set out in an Appendix, and that table, reproduced in the Authority's own Policy Statement PS19/26 of October 2019, gives exemption threshold values of £440, £220 and £85 against reference fraud rates of 0.01, 0.06 and 0.13 per cent for remote card payments and 0.005, 0.01 and 0.015 per cent for remote credit transfers. The sterling figures look untidy because they are the euro ones converted at the Bank of England's spot rate on exit day and rounded down to the nearest five pounds, the same method that turned the €30 low-value exemption into £25.

Article 19 computes that fraud rate as the total value of unauthorised or fraudulent remote transactions over the total value of all remote transactions of that type, on a rolling ninety-day basis, covering authenticated and exempted transactions alike. Article 20 requires the exemption to cease if the rate exceeds the reference rate for two consecutive quarters. That is the discipline: a provider using transaction risk analysis aggressively to remove friction is buying conversion with a fraud budget it must account for quarterly.

One dating caution is unavoidable. The retained version of Commission Delegated Regulation (EU) 2018/389 carries, on the United Kingdom statute book at the time of writing, a recorded future effect: revocation by Schedule 1, Part 3 of the Financial Services and Markets Act 2023, as part of the programme replacing assimilated European Union financial services law with rules made by the regulators. The substance is expected to be restated in the Financial Conduct Authority's own rulebook. Any reader after 2026 must confirm the current instrument before quoting an article number.

■ First-party misuse: the category the system cannot see

The industry's terms are friendly fraud, first-party misuse, first-party fraud and chargeback abuse. Visa's own documentation uses "friendly fraud or first party misuse" interchangeably; Mastercard's programme is called First-Party Trust. There is no agreed definition, and no regulatory definition at all.

The reason it is structurally hard rather than merely annoying is that it is invisible at the message layer. A dispute raised because a stranger stole a card and a dispute raised because a customer changed their mind arrive under the same reason code: Visa dispute condition 10.4, "Fraud — Card-Absent Environment", and Mastercard message reason code 4837, "No Cardholder Authorization". Nothing in the message distinguishes them, and nothing can, because the distinguishing fact is a state of mind.

The economics finish the job. The issuer is under a next-business-day refund obligation, carries the burden of proof, and is told by regulation 75(3) that its own record of the instrument's use does not by itself discharge it. Investigating a £48.99 dispute properly costs more than £48.99. So it refunds and charges back, and the merchant, the only party with evidence about what actually happened, is the party furthest from the decision.

Both major schemes have responded with the same architectural idea: replace argument with a deterministic data match.

Visa's Compelling Evidence 3.0, which took effect on 15 April 2023 as an update to dispute condition 10.4, lets a merchant establish a historical footprint by supplying two prior transactions on the same

credential. Per Visa's merchant readiness documentation of March 2023, those transactions must be at least 120 calendar days old and no more than 365 days old measured from the dispute date, must carry no active fraud report and no active fraud dispute, and must be from the same merchant. At least two of four core data elements must match across all three: user identifier, network address, shipping address, and device identifier or fingerprint; at least one of the two matches must be the network address or the device identifier. Where the criteria are met, liability shifts back away from the merchant. Merchants may respond before a dispute exists, through Verifi's Order Insight service, or afterwards through their acquirer as a pre-arbitration submission in Visa Resolve Online; the submission may be attempted only once, and the response window is thirty days. Mastercard's First-Party Trust programme applies the same principle, requiring a match of one element from each of three categories.

Two consequences deserve naming. The first is that it works, because it removes human judgement: a merchant meeting the criteria wins automatically, while one that does not falls back to ordinary representment where the issuer weighs the evidence and the outcome is uncertain. That difference in expected value has turned qualification into a product requirement rather than a dispute-team concern.

The second is that it quietly converts device fingerprinting and client-side network address capture from optional analytics into commercial necessities, because those are the two elements a merchant cannot reconstruct after the event and one must be present to qualify. A server-side log of the address that received the checkout request often records a load balancer or content delivery network edge rather than the customer. The privacy consequence is real and rarely discussed: a scheme dispute rule has, in effect, mandated persistent device identification across the retail internet.

The scale is not precisely measurable, for the definitional reasons given, but the direction is not in doubt. The Merchant Risk Council's 2026 Global eCommerce Payments and Fraud Report, released on 18 March 2026 and drawn from 1,278 merchant professionals in 37 countries, found sixty-four per cent of merchants reporting increasing rates of first-party misuse, a quarter of those reporting increases of twenty-five per cent or more, while the average number of distinct fraud attack types fell from 4.2 to 3.7. More attacks of one kind, fewer kinds of attack: the profile of a maturing threat.

And it must be said, because merchants rarely say it, that a substantial share of first-party misuse is manufactured by the merchant. A billing descriptor that does not match the trading name produces disputes from customers who genuinely do not recognise the charge, and descriptor drift between the initial charge and the renewal also breaks the scheme's own merchant-matching logic and disqualifies the merchant from the remedy. A free trial that converts without a reminder produces disputes. A dispatch date that slips produces disputes filed under fraud codes, because that is what the banking app offers first. None of that is fraud. It is a merchandising problem wearing a fraud problem's clothes, and it is the cheapest thing on this list to fix.

■ What a competent card-not-present programme actually looks like

The working model is a constrained optimisation rather than a defence. The merchant minimises fraud losses plus dispute handling cost plus revenue lost to declined honest customers, subject to keeping its combined fraud and dispute ratio below the scheme threshold and its fraud rate low enough that its acquirer can keep claiming risk-based exemptions on its behalf. Every control buys one of those terms at the price of another.

Velocity limits and per-issuer throttling belong in that stack for a reason worth naming: a decline is not a neutral outcome. It is an answer, and an endpoint that answers quickly and often is itself a source of

information to somebody testing stolen numbers in bulk, which is why scheme monitoring now treats that testing traffic as a category of its own.

None of these controls is optional in a serious operation and none is sufficient. The internet remains the hard case not because any single defence is weak, but because the underlying credential, sixteen digits and an expiry date, is a static secret that must be shared with every party the cardholder transacts with, and a secret shared with thousands of parties is not a secret. Everything in this chapter is a compensating control for that one architectural fact. The two developments that attack the fact itself are network tokenisation, which replaces the shared number with a merchant-specific or device-specific one, and authentication, which adds a dynamic element to a static exchange. Chapter 45 covered the first. Chapter 47 covers the second.

46.98 Common wrong ideas

Wrong: “Card not present” describes a shopper who was not in the shop. **Right:** It is an assertion the acquiring side puts into the authorisation message, so a customer standing at the counter can generate one when the chip reader is broken and the number is typed into a virtual terminal.

Wrong: Address verification verifies the address. **Right:** It compares only the numeric portion of the street address and the numeric portion of the postcode, so street, town and county names play no part and two barely secret facts defeat it.

Wrong: A full address match is a meaningful gate, so decline everything else. **Right:** Most of the code set means the check did not run, and a rule expressed that way declines a large fraction of legitimate international orders while stopping almost no informed criminal.

Wrong: The address result arrives in time to prevent the authorisation. **Right:** It arrives with or after the authorisation decision, so a merchant rejecting on it must reverse or void the hold, and failing to do so generates complaints that look like fraud disputes and are not.

Wrong: The liability shift moves the loss to a bank. **Right:** It moves a right to argue under private scheme rules, and the party actually carrying the merchant’s risk is the acquirer, which is why rolling reserves, delayed settlement and onboarding refusals exist.

Wrong: “Friendly fraud” is a recognised legal category. **Right:** No regulation anywhere contains the phrase; in United Kingdom law the claim is simply that a transaction was unauthorised, and the burden of proving otherwise sits by statute on the payment service provider.

Wrong: The transaction went through, so obviously the customer authorised it. **Right:** Regulation 75(3) expressly states that use of the instrument as recorded by the provider is not in itself necessarily sufficient to prove authorisation, so that argument has already been anticipated and rejected.

Wrong: Authenticating every transaction ends chargebacks. **Right:** The authentication liability shift covers fraud disputes only, and for many merchants “goods not received”, “not as described” and “credit not processed” are the majority.

Wrong: Data-only authentication buys the liability shift more cheaply. **Right:** Passing risk data to the issuer to improve approval rates is a conversion optimisation, not a liability instrument, and it is regularly mis-sold as both.

Wrong: Published fraud figures can be compared year on year straight out of the report. **Right:** Some count cards rather than people, some count suspected rather than confirmed fraud, and percentages are routinely computed against restated prior-year figures, which is why every number in this chapter carries its date.

46.99 Chapter summary in 20 lines

1. In a shop the chip does a fresh sum every time; online there is no chip, no reader and no sum, only a page of facts about a card, and facts can be copied.
2. Every element the checkout asks for — number, expiry, code, name, billing address — is static, and a criminal who has any complete record has all of them.
3. “Card not present” is not a physical condition but a set of indicators in the authorisation, chiefly the point-of-service entry mode in ISO 8583 data element 22.
4. Entry mode, the e-commerce indicator, the mail-order flag, the cardholder-or-merchant-initiated flag and the recurring indicator each alter interchange, fraud scoring and dispute rights.
5. A merchant that mis-flags its own traffic discovers a liability shift it believed it had does not exist when the dispute arrives, which is among the commonest causes of lost disputes.
6. The address verification service compares numbers only, so a house number and a postcode — the two least secret things about anybody — are enough to pass it.
7. Most of its response codes are not answers at all, and the ones that are get defeated by the same breach that supplied the card number in the first place.
8. The card verification code has real teeth for an archival reason: PCI DSS forbids its retention, so a thief working from stored records should not have it while one working from the physical card will.
9. In the United Kingdom regulation 76 requires the issuer to refund an unauthorised transaction by the end of the business day following the day it becomes aware of it.
10. Regulation 77(4)(d) removes even the £35 excess where the instrument was used in connection with a distance contract, so consumer exposure to unauthorised online card fraud is zero.
11. Regulation 75 puts the burden of proof on the provider and denies that its own record of use discharges it, which leaves the issuer one economically rational move: refund and charge back.
12. The loss therefore lands on the merchant through the acquirer, and the acquirer is the party of last resort when a failed merchant’s chargebacks keep arriving.
13. The merchant’s true cost of one fraudulent order is the order, the goods, the shipping, the fee, the labour, the ratio, the controls bought afterwards and the honest customers those controls turn away.
14. The channel is dearer before any fraud occurs, and the United Kingdom to European interchange increases were found to cost businesses an additional £150 million to £200 million a year.
15. Scheme monitoring, consolidated into the Visa Acquirer Monitoring Program from 1 April 2025, combines fraud reports and disputes into one ratio, so a merchant with no fraud can be caught by a fraud programme.
16. The British chip and PIN migration is the cleanest natural experiment available: counterfeit fraud fell sixty-eight per cent between 2004 and 2008 while card-not-present fraud rose to £328.4 million.
17. Displacement was geographic as well as channel-wise, following the gradient of the weakest reader anywhere in the world that will accept the card.
18. Arithmetic honesty requires the £41.2 billion of online spending alongside the £328.4 million, because the rate per pound fell while the absolute loss more than tripled.
19. Authentication closed the route that replays stolen static credentials and not the route that persuades the legitimate human to authenticate, which is why passcode compromise and wallet enrolment now drive the numbers.
20. Every control in the chapter is a compensating measure for one architectural fact — a static secret shared with thousands of parties is not a secret — and only tokenisation and authentication attack that fact itself.

Chapter sources: Payment Services Regulations 2017 (S.I. 2017/752), regulations 75 to 77 on legisla-

tion.gov.uk, current to August 2026; Commission Delegated Regulation (EU) 2018/389, Articles 16 and 18 to 20 and the Annex, with the FCA Handbook Technical Standards on Strong Customer Authentication, Chapter 3, last updated 19 March 2026; the FCA's Strong Customer Authentication page, its Dear CEO letter of 20 August 2019, and its statement extending the e-commerce deadline to 14 March 2022; UK Finance Annual Fraud Report press releases of 15 June 2026 and 28 May 2025; the joint EBA and ECB payment fraud reports of 15 December 2025 and 1 August 2024; ONS Social Trends, sourced to the UK Cards Association, for the 2000 to 2008 card fraud series; the Payment Systems Regulator's cross-border interchange market review, page last updated October 2025; HM Treasury's March 2025 announcement on consolidating the PSR into the FCA; Visa's Compelling Evidence 3.0 Merchant Readiness document of March 2023; Mastercard Developers documentation for First-Party Trust; PCI Security Standards Council material on PCI DSS v4.0.1 and Requirements 6.4.3 and 11.6.1; the Merchant Risk Council press release of 18 March 2026; and acquirer documentation from Adyen, Worldpay and Stripe for address verification codes and e-commerce indicator mappings.

Three-D Secure and Strong Customer Authentication

47.0 What this chapter gives you

1. You will be able to name the three domains, the component that sits in each — 3DS Server, Directory Server, Access Control Server — and say which one actually authenticates the cardholder.
2. You will be able to explain why the protocol is called three-domain rather than three-dimensional, and why the second wire exists at all.
3. You will be able to distinguish authentication from authorisation, and explain how a cardholder can pass one and be declined by the other ninety milliseconds later inside the same bank.
4. You will be able to read a `transStatus` value and say whether the exchange ended frictionless, went to challenge, or produced only evidence of an attempt.
5. You will be able to name the three artefacts an authentication produces — the authentication value, the e-commerce indicator and the directory transaction identifier — and say what happens when they do not reach the authorisation intact.
6. You will be able to state the counting rule of strong customer authentication and explain why a registered phone plus a fingerprint on it satisfies it while a passcode plus another passcode does not.
7. You will be able to explain dynamic linking and why changing the amount by a penny makes the approval worthless.
8. You will be able to list the exemptions with their figures as they stood in August 2026, and say which of them the United Kingdom has already diverged on or deleted.
9. You will be able to explain why an acquirer-claimed exemption and an issuer-granted one look identical to the shopper and carry opposite consequences when the money turns out to be stolen.
10. You will be able to run the merchant's arithmetic — roughly eight abandoned sales per fraudulent transaction prevented on a £420 bicycle at twelve per cent margin — and say why it inverts completely on a £6 digital item.

The plain version

Two strangers are trying to agree about a third person who is not in the room.

The shop does not know you. It knows a card number you typed in. Your bank knows you well, but has never heard of the shop and cannot reach you through its checkout page. There is no wire between them carrying a question like “is this really her?” The only wire carries money, and money messages are not conversations. They are one-line demands with one-word answers. So the industry built a second wire. Not for money. For a question.

Picture a market stall and a bank on opposite sides of a town, with a switchboard operator between them. The stallholder cannot ring the bank directly: he does not know which of the town's forty banks issued your cheque book, and none of them would take his call anyway. What he can do is hand a sealed note to the switchboard, which reads the first few digits printed on your cheque, identifies the bank, and passes the note across. Your bank opens it, decides, and hands a stamped answer back the same way.

That is the shape of Three-D Secure. Three domains: the shop's side, the bank's side, and the switchboard in the middle belonging to neither. Hence three-domain, hence 3-D. It has nothing to do with three dimensions, and the number of people who assume otherwise is one of the small ongoing embarrassments of the payments industry.

Now, what is in the note?

Priya is buying a bicycle for £420 from a British website at ten past eight on a Thursday evening. When she clicks pay, the shop's software quietly assembles a description of the moment. Not just the card number and the amount. The make and model of her phone, its screen size, its language setting, its time zone. Whether she is in the shop's own app or in a browser, and which browser. The address she wants the bicycle delivered to, and whether it matches the address her card is registered to. Her email address. How many times she has bought from this shop before. Whether this is a one-off purchase or the first payment of a subscription.

All of that goes into the sealed note, and Priya's bank now has something it never had in the old world: context. This is the same phone she used to buy petrol on Tuesday. This is her home broadband. This is her actual delivery address. The amount is large but not absurd for her. Nothing here looks wrong.

So the bank says yes. It stamps the note "authenticated", hands it back, and Priya never sees a thing. The bicycle checkout carries on as though nothing happened. From her side the button just worked. This is the frictionless flow, and on a well-implemented system it is what happens to the large majority of purchases. The check happened. She simply was not asked to participate in it.

Now change one detail. Suppose the phone is one the bank has never seen, the delivery address is in a city Priya has never shopped in, and it is four in the morning. The bank does not feel comfortable. It cannot telephone her, but it can do something better: reach her through the banking app on the phone it does trust. Her checkout pauses, a panel appears, her banking app buzzes. It shows her the merchant's name and the exact amount, £420, and asks her to approve with her fingerprint. She does. The app tells the bank, the bank tells the switchboard, the switchboard tells the shop, and the checkout resumes. That is the challenge flow. It cost her about eight seconds and it is the reason the whole system exists.

Here is the second half of the chapter, and it is a law rather than a piece of software.

In Europe and in the United Kingdom it is not optional for your bank to check that it is really you. There is a rule called strong customer authentication, and it is a counting rule. There are three kinds of evidence a bank can use: something you know, like a passcode; something you have, like a specific phone the bank has tied to your account; something you are, like a fingerprint or a face.

The bank must use at least two of the three, from different families, and they must be independent, meaning a thief who defeats one has not thereby defeated the other. A passcode plus a second passcode is not two things. A password plus a fingerprint is. A registered phone plus a fingerprint on that phone is, which is why almost every bank in Britain settled on exactly that: the phone proves possession, the

fingerprint proves inherence, and the app shows you the amount and the payee so you know what you are approving.

That last part is not decoration. The rule requires the code your phone produces to be welded to that specific amount and that specific merchant. Change the amount by a penny and the approval becomes worthless. This is called dynamic linking, and it exists so nobody can persuade you to approve one thing while a different thing is actually happening.

Now the awkward part, and the reason this chapter has a second half at all.

Checking is expensive. Not in money, in customers. Every time you interrupt a person in the middle of buying something, some of them stop buying. They lose the thread, the app does not open, the phone is upstairs, the passcode does not work, the panel looks like a scam. A rule that challenged everybody for everything would make online shopping meaningfully worse, and regulators knew it when they wrote it.

So the rule comes with holes deliberately cut in it. They are called exemptions, and each names a situation where the risk is low enough that the interruption is not worth it.

If the purchase is small, you can skip the check. In the United Kingdom, as the rules stand at the time of writing in August 2026, small means £25 or less, and only a few times in a row before the bank has to check you properly again. If you are paying the same amount to the same company every month, the bank checks you once at set-up and then leaves you alone. If you have told your bank a particular shop is one you trust, it can stop asking. And the largest hole of all: if the bank runs its own fraud analysis, decides this payment is genuinely low risk, and can prove its overall fraud losses are below a very low threshold, it may skip the check entirely, up to a value ceiling that rises the cleaner its record is.

That last one is worth pausing on, because it is a beautiful piece of regulatory design. It does not tell banks to be safe. It gives them a commercial reward — less friction, more completed sales, happier customers — that they only get to keep while their fraud numbers stay low. Let the fraud rate rise and the reward is taken away automatically.

Finally, the money. Why would a shop volunteer for any of this? Because of who pays when it goes wrong. As the previous chapter showed, when somebody uses a stolen card number online and the real cardholder complains, the default is that the shop loses the goods and the money. But if the shop asked for authentication, and the bank said yes, and the bank was wrong, then in most cases the loss belongs to the bank. The industry calls this the liability shift, and it is the entire commercial engine of 3-D Secure. The shop is not buying security. It is buying a receipt that says: your customer, your call, your loss.

So the shop is caught between two costs. Ask everybody and lose sales to friction. Ask nobody and eat the fraud. The whole of modern e-commerce payments is an argument about where between those two the answer sits, conducted transaction by transaction, in about four hundred milliseconds.

Where the plain version stops being true

Authentication and authorisation are two different messages, on two different networks, with two different answers, and one does not imply the other. The plain version makes it sound like one continuous conversation. It is not. The 3-D Secure exchange happens first, over the internet, between the shop's authentication provider and the bank's authentication server, and it produces a verdict about identity. Only then does the shop send a completely separate authorisation request over the card

network, which produces a verdict about money. A cardholder can pass authentication perfectly and be declined ninety milliseconds later for insufficient funds, for exceeding a velocity limit, or because the fraud engine on the authorisation side reached a different conclusion from the one on the authentication side. Those two systems inside the same issuer frequently disagree, because they are usually different products, sometimes from different vendors, sometimes bought a decade apart. Merchants who have never seen the inside of an issuer assume the bank is one mind. It is at least two, and they are not always on speaking terms.

3-D Secure is not strong customer authentication, and strong customer authentication does not require 3-D Secure. These are constantly conflated and they are different categories of thing. Strong customer authentication is a legal obligation on payment service providers, defined in regulation, that says how confident a firm must be about who it is dealing with. Three-D Secure is a messaging protocol published by EMVCo, a body owned by the card schemes, that carries data between merchants and issuers. The protocol is a delivery van; the regulation is the law about what may be in the van. The practical consequence is that 3-D Secure by itself authenticates nobody. The European Banking Authority made this explicit in its opinion of 21 June 2019: information transmitted using a communication protocol such as EMV 3-D Secure does not, on its own, constitute an inherence element, because the protocol carries no biometric measurement. What authenticates the cardholder is whatever the issuer's access control server does behind the protocol — the push to the banking app, the biometric, the one-time passcode. The van is not the cargo. And in jurisdictions with no authentication law at all, most obviously the United States, 3-D Secure is used enthusiastically anyway, purely for the liability shift.

The liability shift is a private contract, it is narrower than merchants believe, and it sits alongside a statutory allocation that works differently. The shift is not in any statute. It lives in the card schemes' operating rules, which are commercial documents between the schemes and their licensed members, revised at least annually, and which a merchant is bound by only indirectly through its acquirer's contract. It covers fraud-type disputes. It does not cover a customer who says the bicycle never arrived, or arrived buckled, or that they cancelled the order, or that the subscription was misdescribed. Those are processing and consumer disputes and the merchant carries them regardless of how beautifully the transaction was authenticated. Running underneath all of this, and quite separate from it, is the statutory position. In the United Kingdom the Payment Services Regulations 2017 provide that where strong customer authentication is required but the payee or the payee's payment service provider does not accept it, that party must compensate the payer's payment service provider for the resulting loss. That is law, it runs between institutions rather than between a scheme and a merchant, and it does not care what the scheme rulebook says.

Exemptions are requested by one party and granted by another, and the merchant is rarely the one who decides. The plain version implies a merchant can choose to skip the check. It cannot. A merchant, through its acquirer, can flag a transaction as one it believes qualifies for an exemption, and it can express a preference about whether it wants a challenge. The issuer decides. An issuer that disagrees will challenge anyway, or will decline the authorisation with a response code that means, in effect, go back and authenticate properly. Worse, the two sides can be applying different exemptions with different consequences: an exemption claimed by the acquirer generally leaves the fraud liability with the merchant, whereas the same transaction let through unchallenged on the issuer's own risk analysis generally does not. Two transactions that look identical to the shopper, both frictionless, can therefore have opposite consequences if the money later turns out to be stolen. Merchants who optimise for "frictionless rate" without knowing which side granted the exemption are optimising a number that does not mean what they think it means.

The technical version

■ The three domains and the components in them

The three-domain model divides the world into the acquirer domain, the issuer domain, and the interoperability domain that connects them.

In the acquirer domain sits the 3DS Requestor, which is the merchant or whoever acts for it, and the 3DS Server, the component that actually speaks the protocol. Historically this was the Merchant Plug-In, and the abbreviation MPI survives in most gateway APIs. In app-based flows there is also a 3DS SDK embedded in the merchant's mobile application, collecting device information natively rather than through a browser.

In the issuer domain sits the Access Control Server, universally the ACS. It holds the issuer's authentication logic: its risk model, its decision about frictionless versus challenge, its challenge user interface, and its cryptographic signing. Most issuers do not build one. They buy or rent one from a specialist vendor, which is why the challenge screens of unrelated banks in the same country often look suspiciously alike.

In the interoperability domain sits the Directory Server, the DS, operated by the card scheme. It is the switchboard: it maintains the mapping from card ranges to participating issuers' access control servers, enforces protocol version compatibility, applies the scheme's own rules, and routes every message. Each scheme runs its own and brands the consumer-facing product separately — Visa Secure, Mastercard Identity Check, American Express SafeKey, Discover ProtectBuy, JCB J/Secure. The protocol underneath is EMVCo's and common to all of them; the rules layered on top are each scheme's own and are not. EMVCo operates none of this. It publishes the specification and runs the approval processes attesting that a given ACS, DS, 3DS Server or SDK conforms to it.

■ From 3DS1 to 3DS2, and the death of the redirect

The first generation, formally 3-D Secure 1.0.2, was designed for a web that no longer exists. Its flow was a redirect. The merchant's plug-in asked the directory whether the card was enrolled, using a verify enrolment request and response pair. If the answer was yes, the shopper's browser was thrown out of the checkout onto a page hosted by the issuer, where they were asked for a static password. That page was frequently unstyled, frequently unbranded, and frequently arrived without warning. A meaningful proportion of issuers used activation during shopping, which meant a shopper who had never enrolled was asked to invent a password mid-purchase, on a page they had never seen, that looked exactly like the phishing attack it was training them to fall for. Abandonment was severe, mobile support was poor to non-existent, in-app purchases had no sensible flow at all, and the static password was a knowledge factor of low quality that could be captured and reused.

EMVCo's replacement, EMV 3-D Secure, was published in its first form in 2016, with version 2.1.0 following in 2017 and version 2.2.0 in December 2018. Version 2.3 was published by EMVCo in October 2021, followed by 2.3.1.0 in 2022 and 2.3.1.1 in May 2023. At the time of writing in August 2026, a draft of version 2.4.0.0 had been published by EMVCo on 3 June 2026 with a public comment period closing on 1 July 2026, which means the next revision of the protocol is in flight as this book goes to press and any version-specific detail below should be checked against the current specification.

The first generation was retired on a scheme-by-scheme timetable rather than by EMVCo. Visa's published notice, article AI10745 dated 4 February 2021, stated that effective 16 October 2021 it would discontinue support of the 3DS 1.0.2 Attempts Server for non-participating issuers, and that effective 15 October 2022 it would discontinue support of 3DS 1.0.2 and all related technology. Mastercard's timetable ended in the same month. Industry documentation and gateway migration guidance con-

sistently give 14 October 2022 as the date its support for 3DS 1.0.2 ceased, but Mastercard circulates its announcement bulletins to its customers rather than publishing them, so there is no public primary document to cite for it. For practical purposes, from late 2022 onwards, 3DS1 is gone.

■ The message flow in detail

Every message in EMV 3DS is a JSON structure with a defined element set and a version number. The sequence, in the order a practitioner sees it in logs, runs as follows.

The 3DS Server may first send a Preparation Request, PReq, to the Directory Server, and receive a PRes. This is not part of any individual purchase; it retrieves the protocol versions supported for a card range and, importantly, the 3DS Method URL if the issuer's ACS publishes one.

If a 3DS Method URL exists in a browser flow, the 3DS Server loads it in a hidden iframe before authentication begins, letting the ACS run its own device profiling directly against the browser and associate the result with a transaction identifier. The specification bounds how long the 3DS Server waits — ten seconds in the current versions — after which it proceeds regardless. In app flows the equivalent work is done by the 3DS SDK, which gathers operating-system-level device information and encrypts it so only the intended recipient can read it.

Then the substantive message. The 3DS Server sends an Authentication Request, AReq, to the Directory Server, which routes it to the ACS. This carries the context described in the plain version, organised into groups: cardholder data including billing and shipping address, email and telephone; account information such as how long the account has existed with the merchant and whether the card was added recently; purchase data including amount, currency, and whether it is a one-off, a recurring series or an instalment plan; merchant data including the merchant category code and country; browser or SDK device data; and indicators covering the merchant's challenge preference, any exemption claimed, and whether the request is merchant-initiated.

The ACS replies with an Authentication Response, ARes, and the single most important field in it is transStatus.

transStatus	Meaning
Y	Authentication or account verification successful
N	Not authenticated or verified; transaction denied
U	Authentication or verification could not be performed; technical or other problem
A	Attempts processing performed; not authenticated, but proof of an attempt is provided
C	Challenge required; further authentication is needed
D	Challenge required; decoupled authentication confirmed (from version 2.2)
R	Authentication or verification rejected; the issuer is rejecting the request
I	Informational only; the requestor's challenge preference has been acknowledged (from version 2.2)

A transStatus of Y or A ends the exchange there. That is the frictionless flow. Anything requiring cardholder interaction returns C, and the second phase begins.

In the challenge phase the 3DS Server or SDK exchanges Challenge Request and Challenge Response messages, CReq and CRes, with the ACS. In a browser this renders in an iframe using one of a fixed set of standard challenge window sizes, so the ACS knows in advance how much space it has. In an app the SDK renders the challenge natively using a small vocabulary of interface types — single select, multi select, text entry, out-of-band, and in later versions HTML — so it looks like part of the merchant's own app rather than a web page bolted into it.

Out-of-band is the important one, and it is the pattern almost all European issuers use. The challenge screen does not ask for anything. It says: approve this in your banking app. The actual authentication happens on a separate channel, on a device the issuer already trusts, and version 2.3 added automated transitions between the merchant app and the banking app so the shopper need not navigate manually and come back. Decoupled authentication, added in version 2.2, goes further: authentication happens entirely outside the payment session, with no challenge screen at all.

When the challenge concludes, the ACS sends a Results Request, RReq, to the Directory Server, which forwards it to the 3DS Server, which acknowledges with an RRes. This is the authoritative result. A merchant that reads the outcome from the browser rather than from the RReq is trusting the client, which is a category of mistake that needs no further explanation.

Version 2.2 also introduced 3DS Requestor Initiated messages, 3RI, which allow an authentication with no cardholder session at all — for a recurring payment, a split shipment or an account verification.

■ What the authentication produces, and how it reaches the money

Authentication produces three artefacts that must be carried into the authorisation message, and the carriage is where implementations go wrong.

The first is the authentication value: the Cardholder Authentication Verification Value, CAVV, in Visa's vocabulary, or the Accountholder Authentication Value, AAV, in Mastercard's. It is a cryptogram generated by the ACS which the issuer's authorisation system can verify. It is the proof.

The second is the Electronic Commerce Indicator, ECI, a two-digit field that states what kind of authentication occurred. The schemes number them differently, which is a permanent source of confusion:

Visa ECI	Mastercard ECI	Meaning
05	02	Fully authenticated
06	01	Attempted; authentication not completed but an attempt is evidenced
07	00	Not authenticated; no 3-D Secure performed or it failed

The third is the transaction identifier assigned by the Directory Server, dsTransID, which ties the authentication and the authorisation together for dispute purposes.

If these are absent, malformed, or attached to an authorisation whose amount does not match the amount authenticated, the merchant has paid the friction cost and received nothing for it. A substantial share of merchant complaints that the liability shift “did not work” resolve to exactly this.

■ Strong customer authentication as a legal instrument

The obligation originates in Directive (EU) 2015/2366, the second Payment Services Directive, Article 97. The operative detail is in Commission Delegated Regulation (EU) 2018/389, the regulatory technical standards on strong customer authentication and common and secure open standards of communication, adopted on 27 November 2017. Article 38 of that regulation provided that it applied from 14 September 2019, save for certain interface documentation provisions which applied from 14 March 2019. That is the date every practitioner means when they say “SCA came in”.

In the United Kingdom the obligation is in regulation 100 of the Payment Services Regulations 2017, SI 2017/752, which requires strong customer authentication when a payer accesses a payment account online, initiates an electronic payment transaction, or carries out any action through a remote channel which may imply a risk of payment fraud or other abuse, and which additionally requires dynamic linking to a specific amount and payee for remote electronic payments. The detailed standards are the United Kingdom’s own version of the technical standards, made by the FCA and reproduced in the FCA Handbook; all UK figures in this chapter are taken from that instrument as it stood in August 2026.

The United Kingdom did not enforce the e-commerce element on the European timetable. The FCA agreed a phased implementation, extended it twice, and stated that the deadline for full compliance for e-commerce card transactions was 14 March 2022. That date, and not September 2019, is when British online checkouts changed.

The mechanics inside the regulation constrain what any authentication product may do. Article 4 requires that authentication produce a single-use authentication code, that no information about the underlying elements can be derived from it, that a failed attempt must not reveal which element was wrong, that consecutive failed attempts be limited to no more than five before the action is blocked, and that a session of online account access must not survive more than five minutes of inactivity. Article 5 sets out dynamic linking: the payer must be made aware of the amount and the payee, the code must be specific to both, and any change to either must invalidate it. Articles 6 to 8 govern the three categories of element and Article 9 requires their independence, including specific mitigations where a single multi-purpose device — that is, a phone — carries more than one element.

What counts as each element was clarified by the European Banking Authority in Opinion EBA-Op-2019-06 of 21 June 2019, which remains the reference document even where it is no longer formally binding. Its conclusions were unwelcome to parts of the industry. Card details printed on the card are neither a knowledge element nor evidence of possession. A one-time passcode is not a knowledge element, because it is not something the user knows in advance. An application installed on a device is not evidence of possession unless it has been bound to that device. A dynamic card security code does evidence possession, and an app or browser bound to a device through a security chip or a private key does too. On the inherence side, fingerprint, iris, face and hand geometry, voice, vein recognition, keystroke dynamics and the angle at which a device is held all qualify; information transmitted using EMV 3-D Secure does not.

■ The exemptions, with figures, dated

Every figure below is stated as it stood at the time of writing in August 2026, and every one is capable of changing without the underlying mechanism changing at all.

Article	Exemption	EU position (Regulation (EU) 2018/389, as at August 2026)	UK position (FCA technical standards, as at August 2026)
10	Account information accessed directly	Balance and last 90 days of transactions; SCA on first access and at least every 90 days	Same, with Article 10 applying where no account information service provider is involved
10A	Account information accessed via an AISP	No equivalent article	SCA not required provided it has been applied on at least one previous occasion when the AISP accessed the data
11	Contactless at point of sale	EUR 50 single, EUR 150 cumulative, or five consecutive transactions since the last SCA	No monetary limits; exemption available where the transaction is identified by the payment service provider as posing a low level of risk
12	Unattended terminals	Transport fares and parking fees only	Same
13	Trusted beneficiaries	SCA required to create or amend the list; not required for payments to those on it	Same
14	Recurring transactions	SCA on creation, amendment or first initiation of a series of the same amount to the same payee; not on subsequent ones	Same
15	Transfers between a person's own accounts	Where both accounts are at the same account servicing provider	Same
16	Low-value remote payments	EUR 30 single, EUR 100 cumulative, or five consecutive transactions since the last SCA	£25 single, £85 cumulative, or five consecutive transactions since the last SCA
17	Secure corporate processes	Non-consumer payers only, where the competent authority is satisfied as to equivalent security	Same, with the FCA as the satisfied authority
18	Transaction risk analysis	Real-time risk analysis plus a fraud rate at or below the reference rate for the value band	Same mechanism, sterling bands

Three of these deserve expansion.

The contactless exemption is the clearest illustration of why this volume of the book is the perishable one. The European figures of EUR 50 and EUR 150 are as they appear in Regulation (EU) 2018/389

at the time of writing. The United Kingdom diverged progressively, ending with limits of £100 for a single transaction and £300 cumulative, and then abolished the numbers entirely. By the Technical Standards (Strong Customer Authentication and Common and Secure Methods of Communication) (Amendment) Instrument 2025, FCA 2025/62, made under Handbook Notice 136 and announced by the FCA on 19 December 2025, Article 11 was rewritten with effect from 19 March 2026 to remove the monetary caps and replace them with a purely risk-based test: the exemption is available where the payer initiates a contactless transaction identified by the payment service provider as posing a low level of risk. The FCA's announcement noted that most firms were expected to keep their existing limits for the foreseeable future, so the number a British shopper meets in a shop may not have moved even though the rule behind it has been deleted. Any reader consulting this chapter after 2026 should assume the position has moved again and check the FCA's technical standards directly.

The recurring transactions exemption is narrower than the phrase suggests. Article 14 covers a series of transactions of the same amount to the same payee. A subscription that varies in amount — metered usage, a variable utility bill, a plan that changes tier — is not within it. Such payments are usually handled instead as merchant-initiated transactions, which are outside the scope of the obligation altogether rather than exempt from it, provided the mandate itself was established with strong customer authentication. The distinction between exempt and out of scope is not pedantry: exemptions are counted, monitored and can be withdrawn, whereas out-of-scope transactions are simply not the subject of the obligation.

The transaction risk analysis exemption is the commercially important one, and it is a bargain rather than a permission. Article 18 allows a provider to skip authentication on a remote electronic payment it has identified as low risk, but only where the value is at or below an exemption threshold value, and only where that provider's own fraud rate for the relevant transaction type is at or below the reference rate for that band. The rate is calculated under Article 19 as the total value of unauthorised or fraudulent remote transactions divided by the total value of all remote transactions of that type, on a rolling quarterly basis of ninety days, measured across all of that provider's traffic including transactions it authenticated properly.

EU annex (as at August 2026)	Reference fraud rate, remote card	Reference fraud rate, remote credit transfer
ETV EUR 500	0.01%	0.005%
ETV EUR 250	0.06%	0.01%
ETV EUR 100	0.13%	0.015%
UK appendix (as at August 2026)	Reference fraud rate, remote card	Reference fraud rate, remote credit transfer
ETV £440	0.01%	0.005%
ETV £220	0.06%	0.01%
ETV £85	0.13%	0.015%

The sterling figures come from the FCA's instrument FCA 2020/70, which converted the euro amounts at a stated rate of £0.8876 to the euro, and they are corroborated by the FCA's own notification form NOT004 in SUP 15 of the Handbook, which is laid out in the three bands of £440, £220 and £85.

The enforcement mechanism is automatic and is set out in Article 20. A provider whose monitored fraud rate exceeds the applicable reference rate must report that immediately to the competent authority — in the United Kingdom, using form NOT004 — with a description of how it intends to restore

compliance. If the rate exceeds the reference rate for two consecutive quarters, the provider must immediately stop using the exemption in that threshold band, and may not resume until its calculated rate has been at or below the reference rate for a quarter, having notified the authority and evidenced the restoration first. Article 21 separately requires quarterly recording of transaction values, fraud values, resulting fraud rates, average values and the count and percentage of transactions under each exemption, broken down between remote and non-remote.

That is the whole design in one sentence: the reward for good fraud control is less friction, and it is withdrawn on a timetable the firm does not control.

■ Who claims an exemption and who grants it

There are two independent routes and the distinction determines who carries the loss.

The issuer route is the ordinary one: the merchant sends an AReq, the ACS applies its own risk analysis, decides the transaction is low risk, and returns transStatus Y without a challenge. The issuer has exercised its own exemption, and for liability purposes the transaction is generally treated as authenticated.

The acquirer route runs the other way. The merchant, through its acquirer, marks the authorisation as exempt — commonly under transaction risk analysis, low value, or trusted beneficiary — and performs no authentication at all. This uses the acquirer's fraud rate rather than the issuer's and generally leaves fraud liability with the merchant. It is nonetheless widely used, because a merchant with excellent fraud performance may rationally prefer to keep both the liability and the conversion.

Either way the issuer has the final word, and its instrument is the soft decline: an authorisation response that does not mean “no” but “not like this”. Visa's response code 1A, additional customer authentication required, is the one most often seen in Europe, and code 65, whose older meaning was exceeds withdrawal frequency limit, carries the equivalent sense on Mastercard. Both are recorded in that sense throughout the published response-code documentation of acquirers and gateways, Worldpay's and Adyen's among them; the schemes' own authorisation specifications are not public documents, so the mapping cannot be cited to a primary source here. The correct merchant behaviour is to run a 3-D Secure authentication and retry, not to show the customer a failure. Merchants who did not implement that retry loop discovered in the spring of 2022 that they had built a checkout which declined a substantial share of perfectly good British customers and told them their cards had been refused.

Version 2.2 added the fields that make this expressible: an exemption indicator, a challenge indicator, and support for whitelisting so a cardholder can add a merchant to their trusted beneficiary list during the challenge itself. Delegated authentication, where an issuer contracts out the authentication act to a merchant, wallet or gateway meeting its standards, sits alongside rather than inside the protocol.

■ Liability, precisely

The card schemes' rules provide, in substance, that a merchant obtaining a fully authenticated result is protected against the fraud-type dispute categories: Visa dispute condition 10.4, other fraud, card-absent environment, and Mastercard reason code 4837, no cardholder authorisation, are the two a European merchant meets most often. Attempted authentication, ECI 06 or 01, may also carry protection, but the circumstances are scheme-specific and have narrowed over time. The authoritative statement in every case is the current edition of the relevant scheme's rules, revised at least annually and reaching the merchant only through its acquirer's contract. No book is a substitute for them, and this one is not.

What the shift does not do bears repeating. It does not apply to non-fraud dispute categories, nor where the authentication values were not correctly presented in authorisation, nor where the authenticated and authorised amounts differ. It does not protect against a first-party misuse claim dressed as a fraud claim, because at the level of the dispute message those look identical, as the previous chapter set out.

Underneath the scheme rules, the statutory allocation continues to operate. In the United Kingdom, regulation 75 of the Payment Services Regulations 2017 puts the burden of proving authentication and correct execution on the payment service provider and expressly provides that use of the instrument as recorded by the provider is not in itself necessarily sufficient. Regulation 76 requires refund of an unauthorised transaction by the end of the business day following the day the provider becomes aware of it. Regulation 77 sets the payer's own liability at a maximum of £35 for losses arising from a lost, stolen or misappropriated instrument, disapplies even that where the provider failed to apply strong customer authentication as required, and provides that where the payee or the payee's payment service provider does not accept strong customer authentication, that party must compensate the payer's payment service provider for the resulting loss. All of these provisions are as they stood at the time of writing in August 2026.

■ The cost of friction, in numbers

Vendor abandonment statistics vary wildly and are rarely reproducible. The supervisory data is better.

The European Banking Authority and the European Central Bank published the 2025 edition of their joint report on payment fraud on 15 December 2025, covering data to the end of 2024. It found that of electronically initiated card payments made with cards issued in the European Economic Area, around 64% of the total value was authenticated with strong customer authentication, while only about 40% of transactions by volume were, the rest travelling under exemptions. Contactless accounted for around 79% of the volume of non-SCA card payments. Within the EEA, the fraud rate on SCA-authenticated card payments was about 0.016% by value, roughly half the rate on non-authenticated ones. Card payment fraud was around seventeen times higher where the counterpart was outside the EEA, where strong customer authentication is not legally required and frequently not applied. Remote card payments carried a fraud rate about thirteen times that of non-remote payments by value, 0.091% against 0.007%. Card fraud on EEA-issued cards totalled EUR 1.329 billion in 2024, within total fraud of about EUR 4.2 billion across all instruments. All of these figures are as published in that report; the next edition will supersede them.

The British picture from UK Finance's Annual Fraud Report 2026, published on 15 June 2026 and covering 2025, is that remote purchase card fraud losses rose 3% to £423.5 million across 3.2 million cases, an increase of 13% in case numbers, within total payment fraud of £1.28 billion. Unauthorised fraud losses overall fell 5% to £703.4 million. UK Finance specifically noted criminals compromising one-time passcodes in order to register digital wallets or make fraudulent transactions.

That last point matters for anyone designing controls, and it can be stated without giving anybody a manual. The weakest widely deployed possession factor is a one-time passcode delivered over a channel the customer can be talked into repeating aloud. The regulation's own logic — dynamic linking, so the customer sees the amount and payee they are approving; independence, so compromising one element does not compromise another; device binding, so possession means a specific device rather than knowledge of a code — exists precisely because social engineering attacks the human rather than the cryptography. The defensive conclusions follow: prefer app-based approval bound to a registered device over a passcode sent as text; display amount and payee in the approval itself; treat enrolment of a new device or wallet as a high-risk event deserving its own authentication rather than an administra-

tive step; and monitor for the pattern where authentication succeeds but the surrounding behaviour does not match the customer.

The commercial arithmetic merchants actually run compares two expected costs. Challenge a cohort and you lose some proportion of it to abandonment, immediately and certainly, at the full margin. Do not challenge and you keep those sales but carry the fraud on whatever fraction turns out to be stolen, plus dispute fees, plus the effect on the dispute ratio, which at sufficient volume attracts scheme monitoring programmes and their own charges. On a £420 bicycle at a 12% margin the merchant makes £50.40, while a single fraudulent transaction costs the £420, the goods and a dispute fee: the merchant can therefore afford to challenge and lose roughly eight sales for every fraudulent transaction prevented. On a £6 digital item at a 90% margin the arithmetic inverts completely, which is why the low-value exemption exists and why nobody is asked to fingerprint a newspaper subscription.

■ What is going to change, and when

This is the chapter most likely to be wrong first, so it should say plainly where the ground is moving.

In the European Union, the second Payment Services Directive is being replaced by a package comprising a third directive, PSD3, and a directly applicable Payment Services Regulation. Parliament and Council reached provisional political agreement on 27 November 2025, COREPER endorsed the trilogue texts on 22 April 2026, and final compromise texts were published by the Council on 23 April 2026. At the time of writing in August 2026 the author has not verified formal adoption by both institutions or publication in the Official Journal, so the application date should be treated as unsettled. The authentication changes in the agreed texts include an explicit right of access to strong customer authentication for customers without smartphones, with disabilities or with low digital skills; clarified treatment of merchant-initiated transactions, requiring authentication at mandate set-up rather than on each subsequent payment; explicit classification of delegated authentication as outsourcing, with the delegating provider retaining liability; and new technical standards from the European Banking Authority governing exemptions and transaction risk analysis. Reports of the final text also indicate a redefinition permitting two elements from the same category in the specific case of inherence; verify that against the adopted text rather than against this paragraph.

In the United Kingdom, the technical standards themselves are on borrowed time. Legislation.gov.uk records the United Kingdom's version of Regulation (EU) 2018/389 as subject to revocation by Schedule 1 Part 3 of the Financial Services and Markets Act 2023, listed as a change not yet applied as at 17 August 2026, with firm-facing requirements in assimilated law expected to move into FCA rules. When that happens the article numbers used throughout this chapter will cease to exist even where the substance survives. The contactless amendment of March 2026 previews the pattern: a numerical rule replaced by a risk-based standard, the supervisory burden shifting from compliance with a threshold to demonstrating the quality of a firm's own monitoring.

And in the protocol, EMVCo's draft version 2.4.0.0 was out for public comment until 1 July 2026. Version 2.3's direction of travel — FIDO and WebAuthn data carried inside the authentication message, Secure Payment Confirmation, device binding, automated transitions to and from the banking app — points at where this ends: authentication that is cryptographic, bound to a device, invisible in the common case, and no longer recognisable to the customer as a separate step at all. The regulation will still be counting to two.

47.98 Common wrong ideas

Wrong: A successful authentication means the payment will go through. **Right:** Authentication and authorisation are separate messages on separate networks with separate answers, and a cardholder can authenticate perfectly and be declined moments later for funds, velocity or a fraud engine that reached a different conclusion.

Wrong: Three-D Secure is strong customer authentication. **Right:** One is a messaging protocol published by EMVCo and the other is a legal obligation on payment service providers; the European Banking Authority stated in June 2019 that data carried over EMV 3-D Secure is not itself an inherence element, because the protocol carries no biometric measurement.

Wrong: The liability shift is law. **Right:** It lives in the card schemes' operating rules, which are commercial documents revised at least annually and reaching a merchant only through its acquirer's contract, while the statutory allocation runs underneath it and does not care what the rulebook says.

Wrong: The merchant can decide to skip the check. **Right:** A merchant can flag a transaction as one it believes qualifies and express a preference, but the issuer decides, and its instrument for disagreeing is the soft decline.

Wrong: A high frictionless rate is a good number on its own. **Right:** Two identical-looking frictionless transactions can carry opposite liability depending on whether the acquirer claimed the exemption or the issuer granted it, so the rate means nothing without that split.

Wrong: A soft decline means the card was refused and the customer should be told so. **Right:** Codes such as Visa's 1A mean "not like this" rather than "no", and the correct behaviour is to run an authentication and retry, which is precisely what merchants who lost good British customers in spring 2022 had failed to build.

Wrong: A one-time passcode is something the customer knows, so it is a knowledge element. **Right:** The European Banking Authority's opinion is that it is not, because it is not something the user knows in advance, and an app is not evidence of possession unless it has been bound to the device.

Wrong: The recurring exemption covers subscriptions. **Right:** Article 14 covers a series of the same amount to the same payee, so a metered or tier-changing subscription is handled as a merchant-initiated transaction, which is out of scope rather than exempt, and the difference matters because exemptions are counted, monitored and withdrawable.

Wrong: The authentication result can be read from the browser when the challenge finishes. **Right:** The authoritative result is the Results Request sent by the access control server through the directory to the 3DS Server; reading it from the client is trusting the client.

Wrong: Strong customer authentication arrived in British e-commerce in September 2019. **Right:** The Financial Conduct Authority granted successive forbearance and set full compliance for card-not-present e-commerce at 14 March 2022, and that is when British checkouts actually changed.

47.99 Chapter summary in 20 lines

1. Three-D Secure exists because the money wire carries one-line demands and one-word answers, not questions about who is buying, so the industry built a second wire for the question.
2. The three domains are the acquirer's, the issuer's and the interoperability domain between them, and the name has nothing whatever to do with three dimensions.

3. The components are the 3DS Server acting for the merchant, the Access Control Server holding the issuer's authentication logic, and the scheme's Directory Server acting as switchboard.
4. The authentication request carries context — device, browser, addresses, account age, purchase type, merchant data — which is what allows an issuer to decide without asking the cardholder anything.
5. A frictionless flow is not the absence of a check; it is a check the cardholder was not asked to participate in, and on a good implementation it covers the large majority of purchases.
6. The challenge flow reaches the cardholder through a channel the issuer already trusts, and out-of-band approval in the banking app is the pattern almost all European issuers use.
7. Authentication and authorisation are two messages on two networks, and the two systems inside one issuer are frequently different products that do not always agree.
8. The field that matters is `transStatus`: Y or A ends the exchange, C begins the challenge, and the Results Request rather than the browser carries the authoritative outcome.
9. Authentication produces an authentication value, an e-commerce indicator and a directory transaction identifier, and a merchant that fails to carry all three into the authorisation has paid the friction and bought nothing.
10. Strong customer authentication is a legal obligation defined in regulation, while 3-D Secure is a delivery van; the van is not the cargo, and in jurisdictions with no authentication law the protocol is used anyway purely for the liability shift.
11. The rule is a counting rule: at least two elements from different families, independent of one another, which is why a registered phone plus a fingerprint on it became the British default.
12. Dynamic linking welds the resulting code to a specific amount and payee, so an approval obtained for one thing cannot pay for another.
13. Checking costs customers rather than money, so the regulation cuts deliberate holes in itself: low value, recurring, trusted beneficiaries, contactless, unattended terminals, corporate processes and transaction risk analysis.
14. Transaction risk analysis is a bargain rather than a permission, because the reward of less friction is withdrawn automatically once the firm's own fraud rate exceeds the reference rate for two consecutive quarters.
15. The United Kingdom's figures diverged and then partly vanished: £25 low value, sterling bands of £440, £220 and £85, and Article 11's contactless caps deleted with effect from 19 March 2026 in favour of a risk-based test.
16. The merchant requests an exemption and the issuer grants it, and the issuer's way of disagreeing is a soft decline that means "not like this" rather than "no".
17. An acquirer-claimed exemption generally leaves fraud liability with the merchant while an issuer-granted one generally does not, so two identical frictionless transactions can end in opposite places.
18. The liability shift covers fraud dispute categories only, dies where the authentication values were misrepresented or the amounts differ, and cannot separate a first-party misuse claim dressed as a fraud claim.
19. The supervisory data shows authentication working where it is applied and absent where it is not, with about 0.016 per cent fraud by value on authenticated card payments and roughly seventeen times more fraud where the counterpart sits outside the European Economic Area.
20. The direction of travel is cryptographic, device-bound authentication the customer stops noticing, with PSD3 and the migration of the United Kingdom's technical standards into FCA rules about to renumber almost every citation in this chapter.

Sources: Regulation (EU) 2018/389 and Directive (EU) 2015/2366 via EUR-Lex and legislation.gov.uk; the FCA Handbook Technical Standards on Strong Customer Authentication, instruments FCA 2020/70 and FCA

2025/62, Handbook Notice 136, SUP 15, and FCA statements of 19 December 2025 and on the 14 March 2022 e-commerce deadline; the Payment Services Regulations 2017; EBA Opinion EBA-Op-2019-06; the joint EBA-ECB Report on Payment Fraud of 15 December 2025; UK Finance's Annual Fraud Report 2026; EMVCo 3-D Secure materials; and Visa's notice discontinuing 3-D Secure 1.0.2.

PCI DSS and Its Relatives

48.0 What this chapter gives you

1. You will be able to say who writes PCI DSS, who enforces it and why the Council that owns the standard cannot fine anybody for breaking it.
2. You will be able to explain why your validation level follows transaction volume while the work you must actually do follows how close the card number gets to your own computers.
3. You will be able to explain why the same coffee chain answers 27 requirements or 139 depending on a decision a developer made in an afternoon.
4. You will be able to describe what point-to-point encryption and tokenisation do to scope, and name the conditions that silently forfeit the reduction.
5. You will be able to state the storage rules exactly: sensitive authentication data never retained after authorisation even encrypted, the account number rendered unreadable, display limited to the leading digits and the last four.
6. You will be able to say what an Attestation of Compliance proves — a date, a version, a scope and a sample — and what it therefore cannot tell a buyer.
7. You will be able to describe requirements 6.4.3 and 11.6.1, explain the browser-side attack they exist for, and say what the January 2025 revision of SAQ A did to them.
8. You will be able to state the version position as it stood in August 2026: v4.0.1 the only active version, every future-dated requirement effective since 31 March 2025, and a Request for Comments run from 3 June to 20 July 2026.
9. You will be able to place PCI DSS alongside the UK GDPR and explain how an organisation can be compliant with one and in breach of the other on the same day.
10. You will be able to tell a supplier questionnaire that Direct Debit data creates no cardholder data environment, because the account number is the defining factor and there is no account number.

The plain version

Imagine a card number is a key to a house with the address stamped on it. Anyone who copies the key can walk in. The person who loses the key is not the person who gets burgled, and the person who gets burgled is not the person who pays for the new locks. That gap between who is careless and who suffers is the whole reason the rules in this chapter exist.

Six card companies sit at the centre of the world's card payments: American Express, Discover, JCB, Mastercard, UnionPay and Visa. In the early 2000s each wrote its own list of safe-handling rules for card numbers, so a shop taking several brands had to satisfy several overlapping lists. In 2006 they gave up on that and set up a joint body to write one list instead. It is called the PCI Security Standards Council, and the list it maintains is the Payment Card Industry Data Security Standard, or PCI DSS.

Here is the part almost everybody gets wrong. The Council that writes the rules cannot punish you for breaking them. It is not a government; it has no inspectors, no courts and no power to fine anybody, and it is closer to the committee that writes the specification for a plug socket. What gives PCI DSS teeth is that when your shop signed a contract with the bank that pays you — the acquirer — that contract said you would follow PCI DSS. Break it and the person who comes after you is your bank, holding your own signature.

So how much of the rulebook applies to you? Not, as you might expect, according to how big you are. Size decides how you have to *prove* it. What you actually have to *do* depends on something much less obvious: how close the card number gets to your own computers.

Take a real-shaped example. A coffee chain has forty shops and a website, and puts 2.4 million Visa card payments through in a year. Under Visa's published programme at the time of writing, a merchant doing between 1 million and 6 million Visa transactions a year across all channels is a Level 2 merchant, which means it fills in a self-assessment questionnaire once a year rather than paying an outside assessor to write a full report. Good news, so far.

Now look at the website. The Council publishes different questionnaires for different shapes of business, and the one you may use depends on where the payment form comes from. The two numbers below are the Council's own, published in the April 2025 revision of its questionnaire guidance and accurate at time of writing.

If the checkout page drops in a little window supplied whole by the payment provider — an `iframe`, or a redirect off to the provider's own page — then every part of the page that touches the card number came from the provider, not the coffee chain. The chain uses the shortest questionnaire, SAQ A, and answers 27 requirements. If instead the chain's own website builds the payment form, so the card number is typed into a box the chain's own code drew on the screen before being sent on, the chain must use a different questionnaire, SAQ A-EP, and answer 139.

Same company. Same turnover. Same forty shops. Twenty-seven items or 139, decided years earlier by a developer choosing how to build a checkout page, probably in an afternoon, probably without knowing that was the decision being made.

The shops themselves can be shrunk the same way, and there are only two tricks that genuinely work.

The first is point-to-point encryption, and it is a locked post box. A card reader approved for it scrambles the card number inside the reader itself, in hardware, the instant the card is read, using a key the shop does not have and cannot get. What travels through the shop's tills, network and broadband is unreadable, unscrambled only much later inside the payment provider's own protected building. The shop is carrying sealed post it cannot open. Break into the till and you get nothing worth stealing, and the questionnaire the shop fills in gets dramatically shorter.

The second is tokenisation, and it is a cloakroom ticket. When a customer saves a card for next time, the shop keeps not the card number but a meaningless reference — say `tok_8813f2` — that only the payment provider can turn back into a real number, and which the shop can hand back to charge the customer again next month. A thief who steals the shop's entire customer database steals a pile of tickets for a cloakroom they cannot enter.

What is in the rulebook is less exotic than people expect: twelve numbered sections under six goals, mostly the security housekeeping any competent organisation should be doing anyway. Know what computers you have, change the default passwords, encrypt card numbers at rest and in transit, patch things, restrict access, keep logs and read them, test yourself, write it down. Three specifics are unusu-

ally concrete, and these are the figures in force at the time of writing: passwords in scope must be at least twelve characters, or eight if the system genuinely cannot manage twelve; external vulnerability scans must be run at least once every three months by a scanning company the Council has approved; and patches for critical vulnerabilities go on within one month of release.

And when it goes wrong, the money moves in a direction most people do not expect. The card schemes do not bill the breached shop. They assess the shop's acquiring bank, and the bank then comes after the shop under the contract; Visa's published rules go further and forbid the bank from telling the shop that Visa imposed the charge on it. Separately and quite independently, the Information Commissioner's Office can fine a UK organisation for losing personal data up to £17.5 million or four per cent of total annual worldwide turnover, whichever is higher — the statutory maximum in force at the time of writing. Card rules and data protection law are two different animals arriving from two directions on the same bad day.

Where the plain version stops being true

"Certified" is the wrong word, and the annual pass is not a state of being. Nothing in PCI DSS produces a certificate. What it produces is an Attestation of Compliance, a signed declaration about a single assessment covering a defined scope on a defined date. It is a photograph, not a passport. The standard's own position, and the schemes', is that compliance is a continuous obligation: Visa's rules state that a service provider and merchant must always maintain full compliance, and the standard bakes this in by requiring service providers to review at least once every three months that staff are actually performing security tasks as documented, and to reconfirm their own scope at least once every six months. An organisation can hold a perfectly valid attestation dated four months ago and be flagrantly non-compliant today, and frequently is. Passing the test is not the obligation; meeting the standard every day is, and the test is only evidence.

The club does not fine you, and nobody publishes the tariff. The members'-club analogy is wrong in two ways that matter commercially. First, the direction: the schemes assess the acquirer or issuer, who then recovers under contract, which is why "the PCI fine" is a phrase used by people who have never seen the paperwork. Second, the amounts. Circulating figures for per-card penalties and monthly non-compliance charges are not published by the Council, which does not set them, and the schemes do not put numbers on them either — the Council states plainly that any fines or penalties are a matter for each participating payment brand's own programme. Anyone quoting a precise penalty schedule is quoting marketing, a leaked contract, or an invention. One published nuance is revealing, though: Visa states that assessments may be waived where a forensic investigation demonstrates there was no evidence of PCI DSS non-compliance prior to and at the time of the breach. The attestation is not what protects you. The evidence trail behind it is.

Tokenisation and point-to-point encryption reduce scope. They do not remove you from the standard, and the reduction is conditional. This is where the cloakroom-ticket and locked-post-box analogies mislead most expensively. A merchant using a validated, PCI-listed P2PE solution still fills in a questionnaire and still has requirements to meet, and eligibility for the short version depends on conditions the merchant can silently fail: all payment processing must run through the validated solution, no account data may be received or stored electronically by any other route, retained data must be on paper not received electronically, and — the one that catches people — the merchant must have implemented all the controls in the P2PE Instruction Manual supplied by the solution provider. That manual is a set of obligations, not a leaflet. Listings also expire, and the Council is explicit that solutions on its list of expired validations are no longer considered validated. Scope reduction is a thing you maintain, not a thing you buy. Tokenisation has the mirror-image trap: the tokens may be

harmless, but the vault mapping them back to card numbers is squarely in scope for somebody, and if the merchant can retrieve the original number from the token, the merchant has escaped nothing.

You do not choose your questionnaire, and PCI DSS does not cover fraud. The Council writes the questionnaires but does not decide who may use which. That decision belongs to what the Council calls the compliance-accepting entity — in practice the acquirer, the payment facilitator or the brand. The confusion is live enough that the Council issued a bulletin on 4 August 2026 revising its FAQ 1331 specifically to clarify that merchants should always consult their compliance-accepting entities to confirm validation and reporting requirements, because an earlier May 2025 version of that guidance was being misread. Separately, the whole standard concerns the confidentiality of stored and transmitted account data. It is not about fraud. It will not tell you whether the person holding the card is the cardholder, it does not govern chargebacks, and it does not deliver strong customer authentication. Those live elsewhere, including in 3-D Secure and the regulatory authentication regime, treated below as relatives rather than parts of the same thing.

The technical version

■ What PCI DSS is, and who owns it

The Payment Card Industry Data Security Standard is a private technical standard owned and maintained by the PCI Security Standards Council, LLC, based in Wakefield, Massachusetts. The Council was formed in 2006 and is led by a policy-setting Executive Committee drawn from its Founding Members and Strategic Members: American Express, Discover Financial Services, JCB International, Mastercard, UnionPay and Visa Inc. Participating Organisation membership is open to merchants, banks, processors, hardware and software developers and point-of-sale vendors, who gain access to Request for Comments periods and drafts.

The division of labour, in the Council's own words, is the single most load-bearing statement in the subject. The Council is responsible for developing and managing the PCI Security Standards and the related qualification and listing programmes; each participating payment brand maintains its own separate compliance enforcement programme, including which entities must validate compliance, at what validation level, whether an entity may complete a Self-Assessment Questionnaire or must complete a Report on Compliance, and any fines or penalties. Visa states the same split from the other side: the Council owns, maintains and manages PCI DSS and its supporting documents, while Visa manages all data security compliance enforcement and validation initiatives.

PCI DSS is therefore not legislation anywhere. It reaches you through a chain of private contracts: scheme rules bind the acquirer, the acquirer's merchant agreement binds the merchant, and the merchant's contracts bind its suppliers. That chain is not weaker than a statute in practice — it can terminate your ability to accept cards, worse for most businesses than a fine — but it is legally a different creature, and confusing the two produces bad advice in both directions.

■ The enforcement chain, concretely

Visa's programme is Account Information Security, and the governing instrument is the Visa Core Rules and Visa Product and Service Rules. Visa's published summary, accurate at time of writing, states that issuers and acquirers are responsible for ensuring the PCI DSS compliance of their service providers and merchants, including service providers the merchant is using, and that a service provider and merchant must always maintain full compliance, citing Visa Core Rules section IDs 0002228 and 0008031. Where a service provider or merchant does not comply or fails to rectify a security issue, Visa may assess a non-compliance assessment to the issuer or acquirer, who is responsible

for paying all assessments and must not represent that Visa has imposed any assessment on the service provider or merchant, citing section ID 0001054. Visa also states that assessments may be waived where there is no evidence of PCI DSS non-compliance prior to and at the time of a data breach, as demonstrated during a forensic investigation.

Mastercard's programme is Site Data Protection. Its published process, at time of writing, has the merchant determine its level using Mastercard transaction volume from the most recent 52-week period, confirm the applicable validation requirement, engage an approved assessor where required, and submit validation documents to the acquirer, who submits an SDP Acquirer Submission and Compliance Status Form semi-annually reporting progress for Level 1 and Level 2 merchants. Level 3 and Level 4 merchants need not validate compliance to Mastercard, but the acquirer must validate to Mastercard that it has a risk management programme identifying and managing payment security risk within those portfolios. Those merchants are still required to comply; only validation is excused.

Visa also operates the Technology Innovation Programme, which at time of writing eliminates the requirement to validate PCI DSS compliance for eligible merchants where at least 75 per cent of yearly transactions originate through EMV chip-enabled terminals, a validated point-to-point encryption solution, or an integrated tokenisation solution meeting the EMVCo tokenisation specification. Note the precise scope of that relief: it removes the obligation to *validate*, not to *comply*.

■ Validation levels

Levels determine how you prove compliance, not what you must do. The two largest schemes' published criteria, accurate at time of writing, are below. They are not identical, and a merchant that is Level 2 under one may be Level 1 under another because each scheme also picks up the other's Level 1 designations.

Scheme	Level	Criteria as published at time of writing	Validation
Visa	1	Over 6 million Visa transactions annually across all channels, or global merchants identified as Level 1 by any Visa region	Annual ROC by a QSA, or by an internal resource if signed by an officer of the company, plus an AOC
Visa	2	1 to 6 million Visa transactions annually across all channels	Annual SAQ
Visa	3	Fewer than 1 million Visa e-commerce transactions annually across all channels	Annual SAQ
Mastercard	1	More than 6 million combined Mastercard and Maestro transactions annually; or meeting Visa's Level 1 criteria; or designated by Mastercard at its sole discretion	Annual ROC, with the AOC signed by a QSA, a certified ISA, or, unless prohibited by law, an executive officer of the merchant

Scheme	Level	Criteria as published at time of writing	Validation
Mastercard	2	More than 1 million but not more than 6 million combined Mastercard and Maestro transactions annually, or meeting Visa's Level 2 criteria	Annual SAQ; where the SAQ is A, A-EP or D, a QSA or certified ISA must also be engaged for validation
Mastercard	3	More than 20,000 but not more than 1 million combined Mastercard and Maestro e-commerce transactions annually, or meeting Visa's Level 3 criteria	Annual SAQ
Mastercard	4	All other merchants	Annual SAQ; validation to Mastercard not required

Visa notes that level identification is based on the corporate entity's total volume of Visa transactions, inclusive of credit, debit and prepaid, in one country or with one acquirer per year, and that volume from independently owned and operated locations such as franchisees may be excluded if the corporate entity does not process it.

Service providers are treated more strictly. Visa's published criteria at time of writing put VisaNet processors and any service provider storing, processing or transmitting over 300,000 Visa transactions annually at Level 1, requiring an annual on-site assessment and an AOC signed by both the service provider and the QSA; those below 300,000 are Level 2 and submit a signed SAQ D or an AOC including QSA signature, and QSA validation is required before any service provider can be listed on the Visa Global Registry of Service Providers. Mastercard's service provider levels are defined by role rather than volume: at time of writing all Third-Party Processors, Staged Digital Wallet Operators, Digital Activity Service Providers, Business Payment Service Providers, Token Service Providers, 3-D Secure Service Providers, Installment Service Providers and Merchant Payment Gateways are Level 1, together with AML and sanctions service providers, Data Storage Entities and Payment Facilitators above 300,000 combined Mastercard and Maestro transactions annually; those below that threshold, plus Terminal Servicers, are Level 2. Mastercard also recommends at time of writing that Level 1 and Level 2 service providers demonstrate compliance with the Designated Entities Supplemental Validation appendix. For any service provider eligible to self-assess, SAQ D for Service Providers is the only questionnaire; there is no short form.

■ Scope: account data, the CDE, and where the boundary is drawn

PCI DSS applies to all entities that store, process or transmit cardholder data or sensitive authentication data, or that could impact the security of either. Cardholder data comprises the Primary Account Number, cardholder name, expiry date and service code. Sensitive authentication data comprises full track data — magnetic stripe data or the equivalent on a chip — the card verification code, and PINs and PIN blocks. Together these are account data. The PAN is the defining factor: the Council's position is that if an entity stores, processes or transmits PAN, a cardholder data environment exists and PCI DSS requirements apply to it. The storage rules that follow are unforgiving, and are worth stating exactly as they stand at time of writing.

Element	Storage	Must be rendered unreadable when stored
PAN	Kept to a minimum per Requirement 3.2	Yes, per Requirement 3.5
Cardholder name, service code, expiry date	Kept to a minimum per Requirement 3.2 where in the same environment as PAN	No
Full track data, card verification code, PIN and PIN block	Cannot be stored after authorisation per Requirement 3.3.1	Data held until authorisation completes must be protected with strong cryptography per Requirement 3.3.2

Requirement 3.3.1 as published states that SAD is not stored after authorisation, even if encrypted, and that all SAD received is rendered unrecoverable upon completion of the authorisation process; the narrow exception for issuers and companies supporting issuing services is defined separately at Requirement 3.3.3. Requirement 3.4.1 as published states that PAN is masked when displayed, with the BIN and last four digits the maximum number of digits to be displayed, such that only personnel with a legitimate business need can see more.

The cardholder data environment is not only the systems that touch account data. As the Council defines it, the CDE comprises system components, people and processes that store, process or transmit account data, plus components that do not but have unrestricted connectivity to those that do; and PCI DSS additionally applies to components, people and processes that could impact the security of account data. That third category is where most scoping arguments happen, because it pulls in authentication servers, remote access infrastructure, logging servers, jump boxes, monitoring agents and increasingly the entire identity plane.

Scope reduction by segmentation is permitted, and the Council's test is strict: to be out of scope, a system component must be isolated from the CDE such that it could not impact the security of the CDE even if it were compromised. Requirement 12.5.2 obliges the assessed entity — not its assessor — to confirm scope accuracy at least once every 12 months, identifying all locations and flows of account data and all systems connected to or capable of impacting the CDE, and to retain the documentation; Requirement 12.5.2.1 tightens this for service providers to at least once every six months and upon significant change. Where segmentation is relied upon, Requirement 11.4.5 requires penetration testing of the segmentation controls at least once every 12 months and after any changes, and Requirement 11.4.6 requires it at least once every six months for service providers.

■ Tokenisation and P2PE as scope-reduction mechanisms

Tokenisation replaces the PAN with a surrogate value, and its scope effect turns on one question: can the entity holding tokens retrieve the PAN? If it can, the token is a key to cardholder data and the retrieval path is in scope. If it cannot, the token is not cardholder data and the systems holding it may fall outside the CDE, subject as always to whether they could impact the CDE's security by other routes. The vault performing the mapping is in scope for whoever operates it, in practice usually a service provider assessed at Level 1. Where the tokens are EMV payment tokens issued under the EMVCo Payment Tokenisation Specification Technical Framework, the issuing entity is covered by a separate PCI standard, the Token Service Provider Standard, and Mastercard classes all TSPs as Level 1 service providers at time of writing.

Point-to-point encryption is a more absolute mechanism, because it removes cleartext PAN from the merchant environment rather than substituting for it. The governing standard is PCI P2PE, current

at time of writing at v3.2 published June 2025, supported by the P2PE Program Guide v3.2 dated July 2026. A P2PE solution encrypts account data within a PCI-approved point-of-interaction device using keys the merchant never holds, and decrypts it only within the solution provider's or component provider's validated environment. The merchant's own systems carry ciphertext they cannot open.

The reporting consequence is SAQ P2PE, whose eligibility criteria as published at time of writing require that all payment processing is via a validated PCI-listed P2PE solution; that the only systems in the merchant environment storing, processing or transmitting account data are payment terminals from that solution; that the merchant does not otherwise receive, transmit or store account data electronically; that any retained account data is on paper not received electronically; and that the merchant has implemented all controls in the P2PE Instruction Manual provided by the solution provider. It applies to neither e-commerce channels nor service providers. The Council's footnote on validity matters as much as the criteria: solutions on the PCI list with expired validations are no longer considered validated per the P2PE Program Guide, and a merchant using one should ask its acquirer whether the questionnaire is still acceptable.

Encryption solutions that are not PCI-listed are a much weaker position. The Council published Assessment Guidance for Non-listed Encryption Solutions in June 2020 for exactly this case; such solutions confer no SAQ P2PE eligibility, and any scope reduction is a matter for the acquirer's judgement on the evidence, not an entitlement.

■ E-commerce scope: the SAQ A / SAQ A-EP boundary

For card-not-present channels the decisive question is where each element of the payment page originates. SAQ A is for card-not-present merchants — e-commerce or mail and telephone order — that completely outsource all account data functions to PCI DSS validated and compliant third parties, with no electronic storage, processing or transmission of account data on their systems or premises, and any retained account data on paper not received electronically. SAQ A-EP is for e-commerce merchants that partially outsource, with a website that does not itself receive account data but does affect the security of the transaction or the integrity of the page accepting the customer's account data.

The Council's own mapping, in the questionnaire guidance revision of April 2025 and accurate at time of writing:

E-commerce method	Questionnaire	Applicable PCI DSS v4.x requirements
Fully outsourced, merchant has no access to its own webpage	SAQ A	14
Fully outsourced, merchant webpage redirects customers to a compliant TPSP, for example a URL redirect	SAQ A	27
Fully outsourced, merchant webpage includes a compliant TPSP's embedded payment page or form, for example an iframe	SAQ A	27

E-commerce method	Questionnaire	Applicable PCI DSS v4.x requirements
Merchant website creates the payment form and payment data goes directly from the consumer browser to the TPSP, often called a Direct Post; or the merchant website loads or delivers scripts running in the consumer browser that support creation of the payment page or how data is transmitted	SAQ A-EP	139
All other e-commerce methods and implementations	SAQ D for Merchants	All requirements

The Council's rule of thumb is blunt: if any element of a payment page delivered to consumers' browsers originates from the merchant's own website, SAQ A does not apply.

The complete questionnaire set at time of writing is SAQ A, SAQ A-EP, SAQ B, SAQ B-IP, SAQ C, SAQ C-VT, SAQ P2PE, SAQ SPoC and SAQ D, the last in separate merchant and service provider forms. The card-present ones divide by connectivity: SAQ B for imprint machines or standalone dial-out terminals with no internet connection, SAQ B-IP for standalone PCI-listed approved PTS point-of-interaction devices with an IP connection to the processor and isolated from all other system types, SAQ C-VT for manual entry into a validated third-party virtual terminal from an isolated device, SAQ C for payment application systems connected to the internet at a single location, and SAQ SPoC for merchants using a commercial off-the-shelf mobile device with a secure card reader from the Council's list of validated SPoC solutions.

■ The SAQ A revision of January 2025

This episode is the clearest illustration of how the standard and the reporting tools can diverge, and it is dated precisely.

PCI DSS v4.x introduced two requirements aimed squarely at e-commerce script attacks — the pattern in which an attacker modifies or adds a script that a checkout page loads into the customer's browser, so that card details are copied as they are typed, before any of the merchant's own encryption applies. The merchant's servers are never touched and its logs show nothing wrong, which is precisely why the controls exist. Requirement 6.4.3 as published requires that all payment page scripts loaded and executed in the consumer's browser are managed so that a method confirms each script is authorised, a method assures the integrity of each script, and an inventory of all scripts is maintained with written business or technical justification for each. Requirement 11.6.1 as published requires a change- and tamper-detection mechanism to alert personnel to unauthorised modification — including indicators of compromise, changes, additions and deletions — to the security-impacting HTTP headers and script contents of payment pages as received by the consumer browser, evaluated at least weekly or at a frequency set by the entity's targeted risk analysis under Requirement 12.3.1.

On 30 January 2025 the Council announced that, in response to stakeholder feedback about the complexity of implementing these controls, Requirements 6.4.3 and 11.6.1 were being removed from SAQ A, along with Requirement 12.3.1 in so far as it supported 11.6.1. In their place SAQ A gained an eligibility criterion requiring the merchant to confirm that its site is not susceptible to attacks from scripts that could affect the merchant's e-commerce systems. Two versions of SAQ A were live at once for a period: the October 2024 version was retired on 31 March 2025, and the January 2025 version

took effect the same day, which was also the day the future-dated requirements became effective. Requirements 6.4.3, 11.6.1 and 12.3.1 themselves became effective on 31 March 2025 and remain in the standard, continuing to apply to merchants reporting on SAQ D and to merchants assessed on-site with a Report on Compliance. All those dates are as published and accurate at time of writing.

Read the substitution carefully: the reporting burden moved, not the security problem. A merchant attesting to the new eligibility criterion has asserted something about its own susceptibility to script attacks with no prescribed method of verifying it.

■ The version in force, and the migration position, dated

As at the time of writing, August 2026, the position is as follows, and every element of it is dated. PCI DSS v4.0 was published in March 2022. PCI DSS v3.2.1 was retired on 31 March 2024. PCI DSS v4.0.1, a limited revision containing corrections and clarifications and, in the Council's words, no additional or deleted requirements, was published in June 2024. PCI DSS v4.0 was retired on 31 December 2024, after which v4.0.1 became the only active version supported by the Council, and the Council's document library still lists it as the current standard at time of writing.

Version 4.x introduced 64 new requirements, of which 51 were future-dated. Future-dated requirements are treated as best practice until the specified date and are not validated before it; once reached, they become effective and applicable. That date was 31 March 2025, it was not changed by the v4.0.1 revision, and it has passed. There is therefore no remaining phase-in: at the time of writing every requirement of PCI DSS v4.0.1 is fully in force. The requirements below carried the "best practice until 31 March 2025" designation and are now effective; the wording is as published in v4.0.1 and accurate at time of writing.

Requirement	Substance
3.4.2	Technical controls prevent copy or relocation of PAN over remote-access technologies, except for documented, explicitly authorised personnel with a legitimate business need
3.5.1.2	Disk- or partition-level encryption renders PAN unreadable only on removable media, or on non-removable media only where PAN is also rendered unreadable by another mechanism meeting 3.5.1
4.2.1.1	An inventory of trusted keys and certificates protecting PAN in transmission is maintained
5.4.1	Processes and automated mechanisms detect and protect personnel against phishing attacks
6.4.2	Public-facing web applications are fronted by an automated technical solution continually detecting and preventing web-based attacks, actively running, up to date, generating audit logs, and configured to block or alert for immediate investigation
6.4.3	Payment page scripts authorised, integrity-assured and inventoried with justification
8.3.6	Passwords used as authentication factors are at least 12 characters, or eight where the system cannot support 12, and contain both numeric and alphabetic characters

Requirement	Substance
8.4.2	MFA is implemented for all non-console access into the CDE
8.6.2	Passwords for application and system accounts usable for interactive login are not hard-coded in scripts, configuration files or source code
10.4.1.1	Automated mechanisms are used to perform audit log reviews
11.5.1.1	Service providers only: intrusion detection or prevention techniques detect, alert on or prevent, and address covert malware communication channels
11.6.1	Change- and tamper-detection mechanism for payment page headers and scripts, at least weekly or per targeted risk analysis
A1.1.1, A1.1.4	Multi-tenant service providers: logical separation such that neither provider nor customer can reach the other's environment without authorisation, confirmed at least once every six months by penetration testing

Others in the same set cover encryption of SAD held before authorisation completes, keyed hashing of the full PAN, software and component inventories, targeted-risk-analysis frequencies for malware scanning and log review, management of lower-ranked vulnerabilities, awareness training on acceptable use of end-user technologies, and multi-tenant support for customer penetration testing.

The forward-looking position, also dated, is that the Council has begun the next iteration. From 3 June to 20 July 2026 it ran a six-week Request for Comments on the currently published v4.0.1, expressly to gather feedback shaping the standard's evolution and specifically inviting comment on opportunities supporting future technology and artificial intelligence innovation. As at the time of writing no successor version has been published and no publication date has been announced. Anything predicting the content or timing of a next version is speculation; the Council's pattern with v4.0 was a multi-year development cycle, then a transition period with two versions simultaneously active, then a separate phase-in for future-dated requirements, but it has not committed to repeating that shape.

■ The relatives

PCI DSS is one member of a family; the others govern devices, keys, software, mobile acceptance, card manufacture and authentication. The versions below are those current in the Council's document library at time of writing.

Standard	Current version at time of writing	Governs
PCI DSS	v4.0.1, June 2024	Environments where account data is stored, processed or transmitted
PCI PTS POI Modular Security Requirements	v7.0, May 2025	Point-of-interaction devices: PIN terminals, POS devices, encrypting PIN pads, unattended payment terminals
PCI PTS HSM Modular Security Requirements	v5.0, May 2026	Hardware security modules across their lifecycle

Standard	Current version at time of writing	Governs
PCI PIN Security Requirements and Testing Procedures	v3.1, March 2021	Management, processing and transmission of PIN data at ATMs and attended and unattended POS
PCI P2PE	v3.2, June 2025	P2PE solutions, components and applications
PCI 3DS Core Security Standard	v1.0, October 2017	Environments where specified 3-D Secure functions are performed
PCI 3DS SDK Security Standard	v1.1, December 2018	3-D Secure software development kits
PCI MPoC	v1.1, November 2024	Mobile payment acceptance on commercial off-the-shelf devices
PCI Secure Software and Secure SLC	Software Security Framework	Secure design and maintenance of payment software
PCI Card Production and Provisioning, Logical and Physical	Two complementary standards	Manufacture, transport and personalisation of payment cards
PCI Token Service Provider Standard	Current	TSPs issuing EMV payment tokens

Three sunsets are in progress at time of writing, and they matter because they change what a supplier can validly claim. The Council announced a formal sunset period of 1 May to 31 October 2026 for the PCI 3DS Software Development Kit Standard, and the same window for both the Software-based PIN Entry on COTS and the Contactless Payments on COTS Standards, the latter two superseded in substance by MPoC. The Payment Application Data Security Standard was retired earlier, on 28 October 2022, superseded by the Secure Software and Secure Software Lifecycle Standards; a vendor still advertising PA-DSS validation is advertising something that has not existed for years.

Two of these deserve a paragraph rather than a table row.

PCI PTS POI makes P2PE and PIN entry possible at all, because it defines what a trustworthy card reader is. Approved devices are evaluated against physical and logical requirements including tamper responsiveness: the device is built so that interference with its enclosure or internal signals causes it to erase its cryptographic keys and stop working, which is why the correct response to a suspected-tampered terminal is to take it out of service rather than open it up. This is also why PCI DSS Requirement 9.5.1 exists — periodic inspection of POI devices for tampering and unauthorised substitution — and why Requirement 9.5.1.2.1 lets the entity set inspection frequency and type by targeted risk analysis rather than imposing one interval. Approvals expire; a lapsed device does not stop working but does stop counting for eligibility purposes.

PCI PIN is a separate compliance regime with its own assessor qualification, the Qualified PIN Assessor, whose Qualification Requirements stood at v1.2 in March 2025 and Program Guide at v1.2 in May 2025, with the PIN Attestation of Compliance template at v3.3 dated February 2026 — all as published at time of writing. One wrinkle trips people up: Visa announced the sunset of its own PIN Security compliance programme effective 1 October 2023, while stating that clients, processors and service providers would still be required to comply with PCI PIN security requirements. The programme that collected the paperwork went away; the obligation did not.

■ What certification does and does not prove

There is no PCI DSS certificate. There are two reporting outputs and one attestation form that accompanies either. A Report on Compliance is the long form, produced against the Council's ROC Template — at time of writing v4.0.1 revision 3 dated January 2025, and mandatory for QSAs documenting any assessment as a ROC — recording the entity's environment, the samples the assessor selected, and how each requirement was assessed. A Self-Assessment Questionnaire is the short form, available only where the entity meets the eligibility criteria in that questionnaire and its compliance-accepting entity permits it. Either way, the Attestation of Compliance is the covering declaration: a statement of the assessment's results, signed by the assessed entity and, where one was involved, the QSA company.

What an AOC proves is narrow, and the narrowness is the point. It proves that on a stated date, against a stated version of the standard, for a stated scope, on a stated sample, an assessor or the entity itself concluded the applicable requirements were met. Change any one of those variables and the document tells you nothing. The commonest failure in third-party due diligence is not reading the scope section: an AOC covering a supplier's hosting platform in one region does not cover the product you are buying, and one covering "the tokenisation service" does not cover the support tooling your account managers will actually use.

It does not prove the entity is compliant now, because the obligation is continuous and the document is not. Nor does it prove every requirement was met by the defined method. Requirements may be met by the defined approach, by a compensating control documented on a worksheet, or by the customised approach, in which the entity designs its own control to meet a stated Customized Approach Objective and the assessor writes bespoke testing procedures for it. The Council's guidance for compensating controls and the customised approach, most recently revised in June 2026 at time of writing, exists because these are the two mechanisms through which an assessment can pass while looking nothing like the standard's text. Neither is illegitimate; both require you to read further before relying on the attestation.

It also does not discharge your own third-party risk, and PCI DSS pushes that back onto you explicitly. Requirement 12.8.4 requires a programme to monitor third-party service providers' compliance status at least once every 12 months, and Requirement 12.9.2 requires TPSPs to support their customers' requests for information to meet 12.8.4 and 12.8.5 by providing compliance status information and stating which requirements are whose responsibility. The Council also notes that a service provider holding an AOC is expected to provide it to customers on request. If a supplier will not hand over an AOC and a responsibility matrix, that refusal is itself the finding — and two scheme-published registries give an independent cross-check, the Visa Global Registry of Service Providers and the Mastercard SDP Compliant Registered Service Provider List.

Finally, there is the programme nobody wants to meet. The PCI Forensic Investigator programme qualifies firms and individuals, who must already be QSAs, to investigate actual or suspected compromises affecting card transactions or cardholder data, under a single set of qualification requirements accepted across the participating brands rather than the brand-by-brand arrangements that preceded it. A PFI investigation is where the attestation and the reality are compared in writing by someone paid to find the difference, and Visa's waiver language quoted earlier is why that matters: the forensic report decides whether there was evidence of non-compliance before and at the time of the breach.

■ Where PCI stops and the law starts

PCI DSS is contract. Data protection is statute, and the two run on separate rails to separate destinations.

In the United Kingdom, a card number combined with anything identifying a person is personal data, and losing it engages the UK GDPR and the Data Protection Act 2018 regardless of PCI DSS status. The Information Commissioner's Office can impose fines up to the higher maximum of £17.5 million or four per cent of the undertaking's total annual worldwide turnover in the preceding financial year, whichever is higher — the figures in force at time of writing. A valid AOC is evidence of care that may bear on the ICO's assessment; it is not a defence, and the ICO is not bound by the schemes' conclusions or vice versa. An organisation can be compliant with PCI DSS and in breach of data protection law, or the reverse.

The boundary also runs the other way. Bank-transfer and direct debit data — sort code and account number, mandate references, payer details — is not account data as PCI DSS defines it, because the PAN is the defining factor for cardholder data and there is no PAN. Handling UK Direct Debit therefore does not create a cardholder data environment; the obligations that do apply come from the Bacs scheme rules operated by Pay.UK and from the sponsoring bank, separate instruments with separate accreditation routes. Treating one regime as evidence for the other is a category error that turns up regularly in supplier questionnaires.

Strong customer authentication is a third rail again. PCI 3DS Core governs the security of environments where specified 3-D Secure functions are performed; it does not mandate the use of 3-D Secure, and says nothing about when authentication is legally required or what exemptions apply. Three separate documents therefore govern the same checkout page: PCI DSS for the confidentiality of the card number, PCI 3DS Core for the security of the authentication infrastructure where it is in scope, and financial regulation for whether the customer must be authenticated at all.

The final observation is the one to take into a negotiation. PCI DSS is unusually explicit about the limits of its own authority. The Council states that it does not define compliance requirements for any organisation and does not set compliance validation responsibilities; that those are set by brands, acquirers and payment facilitators; and that organisations with questions must consult their compliance-accepting entity — a clarification it repeated as recently as its bulletin of 4 August 2026. When someone tells you PCI DSS requires a particular outcome for your business, the accurate response is that PCI DSS says what it says, and your acquirer decides what you must prove. Those are two different conversations, and only one can be settled by reading the standard.

48.98 Common wrong ideas

Wrong: We are PCI certified. **Right:** Nothing in PCI DSS produces a certificate; what it produces is an Attestation of Compliance covering one scope on one date against one version, which is a photograph rather than a passport.

Wrong: Last year's attestation means we are compliant. **Right:** Compliance is a continuous obligation, which is why service providers must review quarterly that staff are performing the documented tasks and reconfirm their scope at least every six months.

Wrong: The PCI Security Standards Council fined us, and there is a published tariff of per-card penalties. **Right:** The Council sets no compliance requirements and levies nothing, the schemes assess the acquirer who then recovers under contract, Visa's rules forbid the acquirer from representing that Visa

imposed the assessment, and nobody publishes a penalty schedule, so a precise figure is marketing, a leaked contract or an invention.

Wrong: Buying a tokenisation or point-to-point encryption product takes us out of the standard.

Right: It reduces scope conditionally, and eligibility for the short questionnaire is lost by an expired listing, by receiving account data through any other route, or by not implementing every control in the P2PE Instruction Manual.

Wrong: Tokens are harmless, so the whole token estate is out of scope. **Right:** If the entity holding tokens can retrieve the account number, the retrieval path is in scope, and the vault performing the mapping is in scope for whoever operates it.

Wrong: We choose which questionnaire we complete. **Right:** The Council writes the questionnaires but the compliance-accepting entity decides who may use which, a point it restated by revising FAQ 1331 on 4 August 2026.

Wrong: PCI DSS makes us safe from card fraud. **Right:** The standard concerns the confidentiality of stored and transmitted account data; it will not tell you whether the person holding the card is the cardholder, does not govern chargebacks and does not deliver strong customer authentication.

Wrong: Removing 6.4.3 and 11.6.1 from SAQ A removed the script-attack problem. **Right:** The reporting burden moved and the requirements remain in the standard for merchants on SAQ D and on-site assessment, while SAQ A merchants now assert their own non-susceptibility with no prescribed method of verifying it.

Wrong: Visa's Technology Innovation Programme, or being a Mastercard Level 4 merchant, means we need not comply. **Right:** Both remove the obligation to validate, not the obligation to comply, and Mastercard still requires the acquirer to run a risk management programme over those portfolios.

Wrong: A supplier's Attestation of Compliance discharges our third-party risk. **Right:** Requirement 12.8.4 obliges you to monitor supplier compliance status at least annually and 12.9.2 obliges the supplier to tell you which requirements are whose, and the commonest due-diligence failure is not reading the scope section.

48.99 Chapter summary in 20 lines

1. A card number is a key with the address stamped on it, and the person who loses it is not the person who gets burgled or pays for the locks, which is the gap the whole chapter addresses.
2. Six card companies wrote six overlapping rulebooks until 2006, when they founded the PCI Security Standards Council to maintain one.
3. The Council develops the standards and runs the qualification and listing programmes, and each payment brand separately runs its own enforcement, validation and penalty programme.
4. The standard is therefore not legislation anywhere; it reaches a business through a chain of private contracts and its ultimate sanction is the loss of card acceptance.
5. Validation level follows transaction volume, and the two largest schemes' criteria do not agree, so a merchant can be Level 2 under one and Level 1 under another.
6. What you must actually do follows how close the card number gets to your own systems, not how big you are.
7. On a website that difference is 27 requirements under SAQ A or 139 under SAQ A-EP, decided by whether the payment form came whole from the provider or was drawn by the merchant's own code.

8. The Council's rule of thumb is blunt: if any element of the payment page delivered to the browser originates from the merchant's own website, SAQ A does not apply.
9. Point-to-point encryption removes cleartext account numbers from the merchant environment using keys the merchant never holds, and tokenisation substitutes a value only the vault can reverse.
10. Both reduce scope conditionally, and the reduction is something you maintain rather than something you buy.
11. Scope covers account data, everything with unrestricted connectivity to it, and everything that could impact its security, which is the category that pulls in identity, logging and remote-access infrastructure.
12. Segmentation removes a component only where it could not impact the cardholder data environment even if compromised, and the segmentation controls must be penetration tested annually, or six-monthly for service providers.
13. The storage rules are unforgiving: sensitive authentication data must not survive authorisation even encrypted, the account number must be rendered unreadable at rest, and display is limited to the leading digits and the last four.
14. Requirements 6.4.3 and 11.6.1 exist for script attacks that copy card details in the shopper's browser, where the merchant's servers are untouched and its logs show nothing wrong.
15. On 30 January 2025 the Council removed both from SAQ A in favour of a self-asserted eligibility criterion, which moved the reporting burden and not the security problem.
16. As at August 2026, v4.0.1 is the only active version, all 51 future-dated requirements became effective on 31 March 2025, and the Council ran a Request for Comments on the successor from 3 June to 20 July 2026.
17. There is no certificate: there is a Report on Compliance or a Self-Assessment Questionnaire, and an Attestation of Compliance covering one date, one version, one scope and one sample.
18. A requirement may have been met by the defined approach, by a compensating control or by the customised approach, so the attestation alone does not tell you what the control actually is.
19. PCI DSS is one of a family covering devices, hardware security modules, PINs, encryption solutions, 3-D Secure, mobile acceptance, software and card production, with sunsets in progress that change what a supplier may validly claim.
20. PCI DSS is contract and data protection is statute, running on separate rails to separate destinations: the Information Commissioner's Office can fine up to £17.5 million or four per cent of worldwide turnover, an attestation is evidence of care rather than a defence, and your acquirer decides what you must prove.

Sources: PCI SSC document library, standards and blog, including PCI DSS v4.0.1, the v4.x Quick Reference Guide and Prioritized Approach, the SAQ Instructions and Guidelines v4.0.1 r1, and the FAQ 1331 bulletin of 4 August 2026; Visa Account Information Security and the Visa Core Rules; Mastercard Site Data Protection; ICO data protection fining guidance. All verified as published in August 2026.

Chargebacks

49.0 What this chapter gives you

1. You will be able to explain why the money leaves a merchant's account before anybody has decided who was right, and why that ordering is deliberate rather than a defect.
2. You will be able to name the four parties in a card dispute and say which of them the cardholder actually has a claim against.
3. You will be able to read a dispute's reason code as the question being asked, and choose evidence that answers that question rather than the customer's grievance.
4. You will be able to work out, for a given sale, whether contesting a dispute pays, using the dispute fee, the countered fee, staff time and a realistic win rate.
5. You will be able to explain why a merchant refunds a case it would almost certainly win, and why that is arithmetic rather than cowardice.
6. You will be able to compute a VAMP ratio and a Mastercard Excessive Chargeback Programme ratio, and say why the two produce different answers for the same business.
7. You will be able to distinguish a fraud report, filed as TC40 or through SAFE, from a dispute, and explain why an event that moved no money can still end a merchant account.
8. You will be able to say what chargeback gives a United Kingdom consumer that section 75 does not, and what section 75 gives that chargeback cannot.
9. You will be able to explain why a disputed sale can come back at a different amount from the one that settled, and name at least three causes.
10. You will be able to state what a merchant must record at the moment of sale in order to use Compelling Evidence 3.0 eighteen months later.

The plain version

Picture a school summer fete where nobody carries cash. Twenty stalls, and every one is paid the same way: you tell the desk by the gate what you bought and for how much, the desk writes it down, and at the end of the week the desk pays each stall what it is owed. You settle up with the desk separately.

The desk is not a bystander. It is the thing that actually moved the money. So if you come back on Monday and say the cake stall took four pounds off you and handed you an empty box, the desk has a lever nobody else at the fete has. It can decline to pay the cake stall, or, having already paid, take four pounds out of next week's takings.

That is a chargeback, and your bank is the desk. When you tap a card, your bank pays the shop and bills you afterwards. Your money never travels to the shop; your bank's money does. Your bank therefore has both the standing and the plumbing to reverse the payment.

The part that surprises people is the order of events. The money moves first and the argument happens afterwards. You ring your bank, your bank forms a view in a few minutes, and the shop's money is taken. It is not a hearing followed by a verdict; it is a seizure followed by an appeal.

Here is what that looks like with real numbers. Dee runs a small online shop in Leicester selling headphones and turntables. In March she sells a pair of headphones for £145 to a customer in Swansea and posts them the next morning. Her cost was £85, so on that sale she made £60 before overheads.

Eleven weeks later, in the last week of May, an email arrives from her payment provider. The £145 has been disputed. The reason given is that the cardholder says he did not make the payment. Dee's account has already been debited: £145 for the sale, and £20 on top as a fee for having received a dispute at all. That £20 figure is what a mainstream UK payment provider charged at the time of writing in August 2026, and it does not come back whatever happens next.

So before Dee has read the second paragraph of the email, she is £165 down and the headphones are in Swansea.

Now she has a decision to make, and it is not the one most people expect.

She can fight. The industry word for fighting is **representation**, which literally means presenting the transaction a second time: here is the sale again, and here is why it was good. To do that she gathers what she has. The order, the delivery address, the courier's tracking record, the time and date the parcel was signed for, the customer's email address, and the fact that the same email address and the same card bought a turntable from her in January and never disputed it. She uploads all of it. Filing costs another £20 at the time of writing, refundable only if she wins.

And then she waits. Not days. The bank that raised the dispute typically takes two to three months to read the file and decide.

Two things about that wait matter. Dee never speaks to her customer again; she is submitting a file to a stranger at a bank she has no relationship with, who will read it alongside a few hundred others. And the question that stranger answers is not "was Dee treated fairly?" It is a much narrower question, fixed at the very beginning by the box the customer's bank ticked.

That box has a number, and the number is the most important thing in this chapter.

If the box says "I did not make this payment", the only question is whether the cardholder made the payment. Delivery proof helps only insofar as it suggests the real cardholder received the goods at his own address. If the box says "the goods never arrived", the only question is delivery, and a signature at the door ends the argument in a sentence. If the box says "the goods were not as described", a signature at the door is worth nothing at all, because nobody disputes that the parcel arrived; they dispute what was in it.

Merchants lose disputes they should have won by answering the grievance instead of the box.

Suppose Dee wins. The £145 comes back, and so does the £20 filing fee. The £20 dispute fee does not. She has spent three hours and waited eleven weeks to end up £20 and three hours worse off than if the sale had gone through.

Suppose she loses. She is out £145, out both £20 fees, and out the £85 the headphones cost her: £270 in total. On a net margin of four and a half per cent, which is realistic for a small online retailer, £270 is the entire profit on about £6,000 of sales. One disputed pair of headphones eats a fortnight.

If the story ended there, chargebacks would be an expensive nuisance. It does not end there, because the card companies count them.

Every month, Visa and Mastercard divide the number of disputes a shop received by the number of sales it made, and compare the answer with a threshold. Cross it and the shop is put into a monitoring programme. That means monthly fines starting at around a thousand US dollars and climbing, at the top of the schedules in force in August 2026, to two hundred thousand US dollars a month. It means a written remediation plan. And if the numbers do not come down, the shop's bank ends the relationship, at which point its name goes onto a shared list that other banks consult before taking on a new customer. For a business that sells online, being unable to accept cards is not a setback. It is the end of the business.

Dee sells about 9,000 items a month. On Mastercard's threshold in force at the time of writing, a shop of that size is at risk once it takes 135 chargebacks in a month and its ratio reaches 1.5 per cent. Not 135 lost chargebacks. 135 chargebacks. Whether she wins them makes no difference whatsoever to the count, because the count is taken when the dispute is filed, months before anyone knows who was right.

Which produces the most counter-intuitive behaviour in the business, and the thing this chapter exists to explain. A shop with a strong, defensible, almost certainly winnable case will very often refund the customer instead of fighting. Not because it thinks it would lose, but because winning costs a non-refundable fee, several hours of somebody's week, an eleven-week wait, and does nothing at all to the number that can get its ability to take payments withdrawn. Refunding the customer before a dispute is raised costs the sale price and keeps that number down. The shop is not conceding the argument. It is paying to keep a counter from ticking.

That is why the industry has quietly built an entire layer in front of the dispute process: services that let a bank ask the shop about a suspicious transaction before raising anything, and services that let the shop pre-authorise an automatic refund the instant a dispute is about to be filed. Both exist for the same reason. A refund that arrives before the dispute is a refund. A refund that arrives afterwards is a chargeback, and chargebacks are counted.

Where the plain version stops being true

Chargeback is not your legal right, and you are not a party to it. The plain version made it sound as though Dee's customer was pursuing a claim. He was not. He asked his bank to pursue one, and his bank chose to. In the United Kingdom, at the time of writing in August 2026, chargeback has no statutory basis at all: it is a private contractual mechanism created by the card schemes and binding only on their licensees. UK Finance states it in exactly those terms: unlike section 75, chargeback is not a legal right. The Financial Ombudsman Service is as plain, saying a bank does not have to raise a chargeback, though where valid reasons exist it can be good practice, and that what it examines is whether it was fair not to raise one or to discontinue one. That is a duty owed by the bank to its customer, not a right of the customer against the shop. The legal instruments sit alongside chargeback and behave differently: section 75 of the Consumer Credit Act 1974 makes a credit card issuer jointly and severally liable with the supplier for breach of contract or misrepresentation, applying only to credit cards, only where the cash price of the item is over £100 and no more than £30,000, and with a six-year window rather than a scheme deadline; and regulation 76 of the Payment Services Regulations 2017 requires a payment service provider to refund an unauthorised transaction no later than the end of the business day following the day it becomes aware of it, which is a statutory obligation on the bank that exists whether or not any chargeback is ever raised or ever succeeds. The bank refunds because the law tells it to. The chargeback is how the bank goes looking for the money afterwards.

A chargeback is not the reversal of the original payment, and it does not always come back the same size. The fete analogy invites you to picture the same four pounds travelling backwards. It is not the same four pounds and it is often not four pounds. A chargeback is a new financial transaction with its own message, its own clearing cycle, its own value date and its own currency conversion, and any system built on the assumption that the disputed amount equals the original amount will eventually produce a reconciliation break nobody can explain. The commonest cause is foreign exchange: a payment converted at January's rate and disputed in April comes back at April's rate, so a sale that settled as 100 euros can return as 107.86 euros. Three other causes are routine. An issuer can bundle several months of a disputed subscription into one chargeback filed against a single charge. A cardholder can dispute part of an order. And a cardholder can dispute the full amount of a charge already partly refunded, leaving the merchant out the refund and the chargeback at once and having to prove the overlap.

Winning is neither final nor free, and it does not undo the count. The plain version implied a scoreboard with two outcomes. There are more. A successful representment can be met by the issuer filing a further stage, so cases marked won are routinely reversed weeks later and the money leaves the merchant's account twice. Money also moves outside the lifecycle altogether when an issuer credits or debits to correct an amount, which is why a lost case occasionally turns into a late win months after everyone stopped looking. The fee for receiving a dispute is normally non-refundable however the case ends. And the monitoring counters are indifferent to outcomes by design, not by oversight: waiting months for every outcome would make monitoring useless, and what the schemes measure is how well a merchant prevents disputes rather than how well it litigates them. Refunds do not remove a dispute from the count either, so a merchant who refunds after the dispute exists has paid twice for nothing.

Most disputes filed as fraud are not fraud, and the reason code is frequently the wrong one. The plain version treated the box as a neutral label chosen by an informed person. It is neither. A dispute filed as card-absent fraud asserts that the cardholder did not authorise the transaction, and that is the assertion the merchant must answer, regardless of what the customer's actual grievance was. Very often the actual grievance is that the subscription was harder to cancel than to start, or that a family member used the card, or that the statement descriptor was a holding-company name nobody recognised. The industry calls this first-party misuse, and it is a large enough share of card-absent disputes that both major schemes have built dedicated programmes to attack it. Two consequences follow. A merchant holding perfect delivery proof can lose a dispute where delivery was never in question, because it answered the customer instead of the code. And merchants get undeserved wins in the other direction, when a case really about quality is filed as non-delivery and defeated by a tracking number, leaving the customer's genuine complaint unadjudicated by anybody.

The technical version

Everything in this section is stated as accurate at the time of writing, in August 2026, and this is the most perishable chapter in the most perishable volume of this book. Scheme rulebooks are revised at least twice a year, one threshold below moved in April 2026 having moved in 2025, and the fee schedules are contractual documents not published in full. Check the current rulebook and your own acquirer's schedule before relying on any figure here.

■ Vocabulary, and who the parties actually are

Visa's rulebook calls the process a dispute; Mastercard's calls it a chargeback. Most acquirers treat the words as synonyms, and so does this chapter, while keeping one distinction sharp: a **pre-dispute** event is one where no money has moved and no counter has incremented, and a **dispute** proper is one where funds have been debited from the acquirer and passed to the issuer.

The parties are those of the four-party model. The **cardholder** complains. The **issuer** decides whether to raise the case and carries the statutory obligations towards its customer. The **scheme** provides the message formats, the deadlines, the reason codes and, at the end, the adjudication. The **acquirer** is debited by the scheme and passes the debit on to the **merchant**. That last link matters more than merchants realise: the acquirer's right to take money out of a merchant's account comes from the contract between them, not from the scheme, and the acquirer carries the credit risk if the merchant has no money left. That single fact explains rolling reserves, delayed settlement for high-risk sectors, and why an acquirer exits a merchant long before the scheme forces it to.

Two data feeds run alongside the dispute system and are constantly confused with it. Visa's **TC40** and Mastercard's **SAFE** are fraud reports: an issuer that believes a transaction was fraudulent must report it, and the report reaches the acquirer as what most gateways call an early fraud warning. A fraud report is not a dispute and no money moves, but it counts in Visa's monitoring ratio, which is the most misunderstood fact in merchant risk management. Analysis published by one large processor at the time of writing indicates that around 80 per cent of early fraud warnings become a fraud dispute if the merchant does nothing.

■ Stage zero: the pre-dispute layer

Before a dispute exists there is a layer designed to stop one existing, and it explains most of the merchant behaviour that looks irrational from outside. Visa's is operated by Verifi. **Order Insight** pushes enriched transaction detail, the merchant name a human would recognise, the item, the delivery address, into the issuer's call-centre screen and banking app at the moment the cardholder queries the charge, so that "I don't recognise this" is answered before it becomes anything; an April 2026 update lets merchants share Compelling Evidence 3.0 data at this stage. **Rapid Dispute Resolution** goes further: the merchant writes rules in advance, up to ten scenarios keyed on attributes such as amount and dispute condition, and a matching case about to be raised is refunded automatically in near real time. Mastercard's equivalent alerting network is **Ethoca**.

The commercial logic is one sentence: cases resolved through these pre-dispute products are excluded from the merchant's dispute count for monitoring purposes, as Visa's own programme documentation states. A merchant paying an alert fee and refunding a sale in full is buying the removal of a tick from a counter, and that is often worth several times the value of the sale.

American Express and Discover retain a formal **inquiry** stage, sometimes called a retrieval or request for information, where the issuer asks before it takes; Visa and Mastercard no longer use inquiries this way. Failure to answer an inquiry is treated as implicit acceptance and produces a chargeback that is, in practice, unwinnable.

■ The lifecycle, scheme by scheme

Visa has run two workflows since Visa Claims Resolution took effect in April 2018. The **allocation** workflow covers fraud, dispute condition group 10, and authorisation, group 11: Visa allocates liability up front from data it already holds, such as whether the transaction was authenticated or authorised, so there is no classic representment and the merchant's evidence is carried into a pre-arbitration case filed by the acquirer. The **collaboration** workflow covers processing errors, group 12, and con-

sumer disputes, group 13, and works the way most people picture it: chargeback, representment by the acquirer, then, if the issuer is unpersuaded, an issuer-initiated pre-arbitration, then arbitration.

Mastercard runs one flow: first chargeback, second presentment by the acquirer, pre-arbitration by the issuer, arbitration case filed with Mastercard's Dispute Resolution Management team. The mechanics are executed in Mastercom, and a party that fails to respond within its window is treated as having accepted liability.

American Express is a three-party scheme and its own issuer, so there is no inter-bank adjudication to have: inquiry, chargeback, a single opportunity to respond, and the decision is final. There is no arbitration stage.

■ Reason codes, and what each one actually demands

The reason code is the question. Evidence that does not answer it is not weak evidence but irrelevant evidence, and submitting a great deal of it is the commonest self-inflicted loss in the trade.

Visa's conditions and issuer time limits, as published at the time of writing in August 2026:

Code	Name	Issuer window	Workflow
10.1	EMV liability shift, counterfeit fraud	120 days	Allocation
10.2	EMV liability shift, non-counterfeit fraud	120 days	Allocation
10.3	Other fraud, card-present environment	120 days	Allocation
10.4	Other fraud, card-absent environment	120 days	Allocation
10.5	Visa Fraud Monitoring Programme	120 days	Allocation
11.1	Card recovery bulletin	75 days	Allocation
11.2	Declined authorisation	75 days	Allocation
11.3	No authorisation	75 days	Allocation
12.2 to 12.5	Incorrect transaction code, currency, account number or amount	120 days	Collaboration
12.6.1	Duplicate processing	120 days	Collaboration
12.6.2	Paid by other means	120 days	Collaboration
12.7	Invalid data	75 days	Collaboration
13.1	Merchandise or services not received	120 days	Collaboration
13.2	Cancelled recurring transaction	120 days	Collaboration
13.3	Not as described or defective	120 days	Collaboration
13.4	Counterfeit merchandise	120 days	Collaboration
13.5	Misrepresentation	120 days	Collaboration
13.6	Credit not processed	120 days	Collaboration
13.7	Cancelled merchandise or services	120 days	Collaboration

Code	Name	Issuer window	Workflow
13.8	Original credit transaction not accepted	120 days	Collaboration
13.9	Non-receipt of cash at an ATM	120 days	Collaboration

Mastercard's principal codes, with the windows published at the time of writing:

Code	Name	Issuer window
4808	Authorisation-related chargeback	90 days
4834	Point-of-interaction error	90 days
4837	No cardholder authorisation	120 days
4841	Cancelled recurring or digital goods transaction	120 days
4849	Questionable merchant activity	Programme-driven
4853	Cardholder dispute	120 days, start date varies
4863	Cardholder does not recognise, potential fraud	120 days
4870	Chip liability shift	120 days
4871	Chip and PIN liability shift	120 days

The evidence that answers each family is different in kind.

For **fraud in a card-absent environment**, Visa 10.4 and Mastercard 4837 and 4863, the merchant must show that the person who benefited was the cardholder: delivery to the cardholder's own billing address, a signature, an account with a login history predating the order, prior undisputed purchases from the same device or address, and any usage log showing the account was used after the purchase. Nothing about product quality is relevant, and a successful 3-D Secure authentication usually disposes of the case before any of it matters.

For **non-delivery**, Visa 13.1 and Mastercard 4853, the merchant must show delivery, or that the promised delivery date has not yet passed. A courier record placing the parcel at the cardholder's address, ideally with a signature or a photograph and a geolocation, ends it; for digital goods, an access log with timestamps and IP address.

For **not as described or defective**, Visa 13.3, the burden is descriptive: the listing as it stood on the day, photographs, specifications, the terms the customer accepted, and evidence that what shipped matched the listing. Proof of delivery is worthless here.

For **cancelled recurring or subscription** disputes, Visa 13.2 and Mastercard 4841, the merchant must produce the cancellation policy as displayed, the record that the customer agreed to it, the absence of a cancellation request before the billing date, and, if a trial converted, evidence of the reminder sent before conversion. For **credit not processed**, Visa 13.6, the answer is either the refund record with its date and amount or the published policy showing no refund was due. For **processing errors**, Visa group 12 and Mastercard 4834, the merchant is defending its own systems: authorisation and settlement records, the currency indicator submitted, and proof that a supposed duplicate was two genuine purchases.

■ The clock

Deadlines are calendar days, and they are enforced by software rather than by people. A missed deadline is a loss with no appeal.

Stage	Visa	Mastercard	American Express
Issuer raises dispute	120 days for most conditions, 75 for authorisation conditions	120 days for most codes, 90 for 4808 and 4834	120 days
Where goods or services are supplied later	Clock runs from the expected delivery or performance date	Clock runs from the delivery, cancellation or service date depending on code	Extended for certain conditions
Merchant or acquirer responds	30 days	45 days	20 days
Issuer escalates	30 days to pre-arbitration	45 days to pre-arbitration	Not applicable
Acquirer answers escalation	30 days	30 days	Not applicable
Arbitration	Filed after the escalation window	10 days to escalate after rejection	No arbitration stage

Two corrections to that table. The merchant almost never gets the scheme's window, because the acquirer or gateway must reserve time to format and transmit the response, so the internal deadline a merchant sees is typically 7 to 21 days by network and provider. And the issuer's own decision takes far longer than any response window: 60 to 75 days is normal, and the full lifecycle commonly runs two to three months, with no way to accelerate it other than to give up.

■ Compelling Evidence 3.0, and its Mastercard counterpart

The single most useful development for merchants in this decade is Visa's Compelling Evidence 3.0, effective 15 April 2023 and still in force at the time of writing. It applies to dispute condition 10.4, fraud in a card-absent environment, and it lets a merchant defeat the dispute with history rather than argument.

To qualify, the merchant must identify two previous transactions by the same cardholder that were not disputed, that are at least 120 days old and no more than 365 days old measured from the dispute date, and that share with the disputed transaction at least two of four data elements: user or account identifier, IP address, delivery address, and device identifier or fingerprint. At least one of the two matching elements must be the IP address or the device identifier. Meet that and liability moves back to the issuer.

Two design points matter. This is a rule about instrumentation rather than advocacy, so a merchant that does not record device fingerprints and IP addresses against orders cannot use it however honest its business; and Visa excludes qualifying fraud reports from the monitoring ratio, so the mechanism protects the counter as well as the money.

Mastercard's answer is the **First-Party Trust** programme, launched in the United States and, per Mastercard's announcement of June 2025, expanding to Canada, Latin America and the Caribbean, and Asia Pacific. Merchants supply enriched evidence either during authorisation or afterwards during the dispute, including purchase history, device details, delivery information, identity elements and location, so that an issuer can test a cardholder's claim at the point it is made. Mastercard's framing

of the problem, citing its 2025 State of Chargebacks report, is that the global cost of chargebacks is forecast to reach 42 billion US dollars by 2028, with nearly half reported as fraudulent.

■ Who pays, at each stage

Follow the money for a single £145 card-absent dispute at a UK merchant, on the published fee schedules in force at the time of writing in August 2026.

Event	Merchant	Acquirer or PSP	Issuer
Original sale	+£145 less scheme and acquirer fees	Earns processing margin	Funds the purchase, bills cardholder
Dispute raised	-£145 and -£20 dispute received fee	Debited by scheme, passes on	Credits cardholder, often already required to by law
Representment filed	-£20 dispute countered fee	Transmits and formats the case	Reviews for 60 to 75 days
Merchant wins	+£145 and +£20 countered fee	Fee income retained	Bears the loss or pursues the cardholder
Merchant loses	Nothing further moves; total cost £185 plus £85 of stock	Fee income retained	Retains the funds
Pre-arbitration	Further acquirer handling fee	Charges per case	Files at its own cost
Arbitration	Scheme fees on the losing side	Passes them through	Scheme fees on the losing side

The £20 received and £20 countered figures are the published UK dispute fees of one large payment provider, effective 17 June 2025 and unchanged at the time of writing; the same provider charges 15 US dollars for each in the United States and 20 euros in the euro area. Traditional UK acquirers typically charge £15 to £25 per chargeback. The countered fee is returned on a win; the received fee, at that provider and most others, is not.

Arbitration is where the economics become absurd on purpose. The schemes' fee schedules are contractual documents and are not published, which is easy to demonstrate: Visa's own public rulebook, the Visa Core Rules and Visa Product and Service Rules in the edition of 18 April 2026, sets out the whole dispute resolution machinery, arbitration and compliance included, without stating a single amount, referring the reader instead to the applicable fee schedule, and Mastercard's Chargeback Guide reaches members through its technical resource centre rather than open publication. The shape of the charging is nonetheless consistent across every account of it: a filing fee paid by the party that escalates, a substantially larger fee borne by the party that loses, and a separate penalty for technical violations such as filing outside the rules or in the wrong format. Industry sources consistently put the minimum exposure on a losing party at several hundred US dollars, with case filing and ruling amounts commonly quoted between 500 and 600 US dollars and technical violation penalties of 100 to 250 US dollars, and those sources disagree with one another at the edges. Treat the range as the settled view of the trade rather than as a verified figure, because no publicly available primary specification confirms it.

The consequence is not in dispute even where the numbers are. Spending six hundred dollars to contest £145 is irrational, so arbitration is almost never used at ordinary retail values. It exists as a threat that disciplines the earlier stages, and as a remedy for large-ticket disputes: travel, freight, professional services, luxury goods.

■ Ratios, and the programmes that end merchant accounts

This is the part that decides whether a business survives, and the arithmetic differs between schemes in ways that catch people out.

Visa replaced the old Visa Dispute Monitoring Programme and Visa Fraud Monitoring Programme with a single **Visa Acquirer Monitoring Programme**, VAMP, which took effect on 1 April 2025 with an advisory period that Visa's published fact sheet ended on 30 September 2025.

The VAMP ratio is the count of TC40 fraud reports plus the count of TC15 disputes, divided by the count of settled transactions, measured monthly on card-absent volume. Three features of that formula matter enormously. Fraud reports count even though no money moved and no dispute was ever raised. A transaction appearing in both the TC40 and TC15 feeds is counted twice. And the denominator is settled transactions, so declining suspicious authorisations does not flatter the ratio, though it does keep fraud out of the numerator. Excluded from the count are cases resolved through pre-dispute products and fraud reports qualifying for Compelling Evidence 3.0.

On the thresholds in force at the time of writing in August 2026, a merchant is identified as excessive at a VAMP ratio of 1.5 per cent in Asia Pacific, Canada, the European region, the United States and Latin America and the Caribbean, and at 2.2 per cent in Central Europe, the Middle East and Africa. The 1.5 per cent figure is recent: Visa cut it from 2.2 per cent on 1 April 2026. A minimum count applies before a merchant is assessed at all, 1,500 combined items a month in most regions, or 150 items and 75,000 US dollars in CEMEA. At acquirer portfolio level the thresholds are far tighter: 0.5 per cent above standard, 0.7 per cent excessive.

Visa monitors enumeration, meaning card-testing traffic, separately: excessive monitoring applies at an enumeration count of 300,000 and a ratio of 20 per cent, with no fines attached.

Visa assesses a per-item fee on identified merchants but does not publish the schedule in its public fact sheet. Industry reporting at the time of writing puts the fee at 8 US dollars per fraud report or dispute at the excessive tier and 4 US dollars at the above-standard tier, with a three-month grace period for a first identification in a rolling twelve months. Treat those as industry reports rather than primary citations.

Mastercard operates the Excessive Chargeback Programme in two tiers, and calculates its ratio differently: chargebacks received in the current month divided by the transaction count of the **preceding** month. That definition punishes sharply seasonal merchants, whose January disputes are divided by a December that has ended.

Programme	Chargeback count	Ratio	Effect
Excessive Chargeback Merchant	100 to 299 in the month	1.5 to 2.99 per cent	Escalating monthly fines from month two
High Excessive Chargeback Merchant	300 or more in the month	3 per cent or more	Steeper fines, plus issuer recovery

The assessment schedules published at the time of writing, in US dollars or euros per month, count from the first month above threshold:

Months above threshold	Excessive Chargeback Merchant	High Excessive Chargeback Merchant
1	none	none
2	1,000	1,000

Months above threshold	Excessive Chargeback Merchant	High Excessive Chargeback Merchant
3	1,000	2,000
4 to 6	5,000	10,000
7 to 11	25,000	50,000
12 to 18	50,000	100,000
19 and beyond	100,000	200,000

From month four an issuer recovery assessment is added: 5 US dollars or euros for each chargeback above 300 in the month, so a merchant identified in month four with 400 chargebacks pays 5,500 on top of nothing else. Exit requires three consecutive months below the threshold, and the schedule then resets to month one rather than resuming where it stopped. Mastercard counts a chargeback regardless of any prior refund, of liability shift, and of the outcome.

Mastercard's separate **Excessive Fraud Merchant** programme targets a different failure: high fraud combined with an unwillingness to authenticate. On the criteria published at the time of writing, all of the following must hold in a month: at least 1,000 Mastercard e-commerce payments; net fraud chargeback volume above 50,000 US dollars, or 15,000 Australian dollars in Australia; a fraud chargeback rate above 0.5 per cent, or 0.2 per cent in Australia; and 3-D Secure used on no more than 10 per cent of Mastercard payments in non-regulated countries or 50 per cent in regulated ones. Fines run from 500 US dollars in month two to 100,000 US dollars from month nineteen. Where both are breached, this programme takes precedence but the Excessive Chargeback Programme clock keeps running underneath.

Two regional regimes deserve naming. In the United States only, Visa's **Secure Excessive Fraud Programme** captures merchants whose fraud on domestic 3-D Secure authenticated Visa transactions reaches 75,000 US dollars and 0.9 per cent in a month; there is no fine, but the merchant loses the liability shift on domestic authenticated transactions until it fully exits, which is heavier than any fine. In Australia, the AusPayNet card-not-present framework works quarterly from 50,000 Australian dollars of fraud chargebacks and a 0.2 per cent ratio, escalating to a mandate that every card-not-present transaction be sent for strong customer authentication.

■ The lists, which are the real sanction

Fines are survivable. The terminated merchant files are not, and they operate on thresholds far below the monitoring programmes.

Mastercard's **MATCH**, the Mastercard Alert to Control High-risk Merchants system, is a shared database of merchants whose acquiring relationship was terminated. Every processor must check it when onboarding, and must add a terminated merchant meeting the criteria within one business day. Eleven of the reason codes are qualitative, covering laundering, collusion, illegal transactions, PCI DSS non-compliance and the rest. Two are quantitative, and as published at the time of writing they are stark. Reason code 4, **excessive chargebacks**, means Mastercard chargebacks in any single calendar month exceeding 1 per cent of Mastercard sales transactions in that month and totalling 5,000 US dollars or more, with no minimum chargeback count. Reason code 5, **excessive fraud**, means a fraud-to-sales dollar volume ratio of 8 per cent or more in a calendar month with ten or more fraudulent transactions totalling 5,000 US dollars or more. Note the 1 per cent: below every monitoring threshold in this chapter, and indifferent to whether the chargebacks were later reversed or won. Records remain for five years and are then purged automatically; a processor can remove an entry only if it

listed the merchant in error, or where the listing was for PCI DSS non-compliance since remedied. Mastercard does not assess the accuracy of listings.

Visa's equivalent is the **Visa Merchant Screening Service**. Its two quantitative reason codes, as published at the time of writing, are excessive disputes, meaning a dispute count of 1,000 and a 1.8 per cent dispute-to-sales amount ratio in any single month, and excessive fraud, meaning 250,000 US dollars of fraud at the same 1.8 per cent ratio; both are qualified by the merchant having failed to remediate adequately.

Neither list is supposed to be a blacklist. Both function as one, because most acquirers decline an application that produces a hit.

■ Why a merchant refunds a case it would win

Everything above assembles into two decisions taken at different moments, and only the second is about winning. Confusing them is why merchant behaviour looks like cowardice from outside.

The first decision comes before any dispute exists, when an early fraud warning or a pre-dispute alert arrives. Refunding costs the whole sale, £145, with certainty. Doing nothing risks the dispute, and on the processor analysis cited earlier roughly 80 per cent of ignored early fraud warnings become disputes: a non-refundable £20 fee, the £145 if lost, the stock, and one tick on a counter that no later victory removes. Hence the published guidance of one large processor at the time of writing, which is to refund proactively where the charge is roughly at or below your dispute fee and generally not to bother above about 135 per cent of it. For a UK merchant paying £20, that means refunding the small ones and standing and fighting on the rest.

The second decision comes after the dispute exists, and by then the counter has already moved. Refunding now achieves nothing: the merchant cannot refund outside the dispute process while it is open, and a refund would not remove the dispute from any monitoring count. So the only remaining question is money. Fighting the £145 case costs the £20 countered fee, returned on a win, plus perhaps £30 of loaded staff time to assemble the evidence properly. At a 40 per cent win rate, realistic for a well-evidenced delivery case, the expected recovery is £66 against about £42 of expected cost, and fighting pays. At a 15 per cent win rate, realistic for a subscription dispute with no usage logs, the expected recovery is £25 against £47 of cost, and it does not.

The behaviour looks irrational because observers watch the second decision while the merchant is making the first. A business inside a monitoring programme, or near one, is not optimising individual cases at all. Its objective is the three consecutive months below threshold that end the programme, and the only lever on the numerator is stopping disputes before they are filed. Hence alerts that cost more than the sale they cancel.

Three circumstances justify fighting regardless. Where the value is large enough that the fee stack is immaterial. Where the case is a template, a cancellation policy or a delivery process that will generate hundreds of identical disputes, so winning once fixes evidence practice for all of them. And where the merchant is comfortably below every threshold, so the counter is not binding and money is the only consideration.

■ The United Kingdom layer

British merchants and issuers operate under scheme rules and statute simultaneously, and the two do not align.

Regulation 76 of the Payment Services Regulations 2017, quoted earlier for its next-business-day deadline, also requires the account to be restored to the state it would have been in, with a value date no

later than the original debit. That obligation is owed by the bank to its customer and is independent of the chargeback. Regulation 77(6) allocates the loss further along: where strong customer authentication was required under regulation 100 but the payee or the payee's payment service provider did not accept it, that party must compensate the payer's payment service provider for the sums it paid out under regulation 76. That is the statutory backbone of what merchants experience as the 3-D Secure liability shift, and it is why a UK merchant that suppresses authentication is not merely accepting scheme liability but a regulatory one.

Section 75 of the Consumer Credit Act 1974 is the other statutory route, and often the better one: joint and several liability of the credit card issuer with the supplier, on the terms set out earlier, with a six-year limitation period rather than a 120-day scheme deadline. Chargeback reaches debit cards and small values that section 75 does not; section 75 reaches consequential loss and long-dead purchases that chargeback cannot. A bank handling a complaint properly considers both.

The Financial Ombudsman Service sits above all of it. Its approach for consumers is described earlier in this chapter; for a merchant complaining about its acquirer, its stated position is that where the acquirer acted in accordance with the merchant services agreement the complaint is unlikely to be upheld, which returns the merchant to the contract it signed.

For scale, UK Finance's Annual Fraud Report published on 15 June 2026 recorded £1.28 billion of payment fraud losses in 2025, up 4 per cent, of which £703.4 million was unauthorised fraud, down 5 per cent. Remote purchase fraud, the card-absent category that generates most fraud disputes, was £423.5 million across 3.2 million cases: a 3 per cent rise in value on a 13 per cent rise in cases. Falling average values on rising case counts is precisely the pattern that pushes dispute counts up while losses hold steady, and it is counts that the monitoring programmes measure.

■ What is moving under this chapter

Three things will date this chapter first, and a reader in 2028 should check all three before quoting any of it.

The thresholds are moving downwards. Visa's merchant excessive threshold fell from 2.2 per cent to 1.5 per cent on 1 April 2026, having been set only a year earlier, and the acquirer-level figures moved on their own schedule. There is no reason to expect that to be the last revision.

The adjudication is being automated. On 1 April 2026 Visa announced six dispute services: the Visa Dispute Resolution Network for pre-dispute handling, in pilot with general availability planned for late 2026; Visa Dispute Recovery Manager, which automates representment with generative responses and win-prediction scoring; Dispute Intelligence and Dispute Doc Analyzer, which summarise and score cases for issuers and acquirers; Visa Dispute Case Manager; and the Order Insight update that carries Compelling Evidence 3.0 into the pre-dispute conversation. Visa's stated context was 106 million disputes processed globally in 2025, a 35 per cent increase since 2019. When both sides of a dispute are assembled and assessed by models, the evidence that wins will be structured data captured at the time of the order rather than prose written afterwards, which is a change in what merchants must build, not merely in what they must write.

And the ground is shifting under the word fraud. Compelling Evidence 3.0 and First-Party Trust both let history speak for the merchant, and both reward instrumentation over argument. The merchant that survives the next revision of these rules is not the one with the best dispute writer. It is the one that recorded the device, the address, the login and the delivery at the moment of sale, and can prove eighteen months later that the same person came back.

49.98 Common wrong ideas

Wrong: A chargeback is a legal right the cardholder exercises against the shop. **Right:** In the United Kingdom it has no statutory basis at all; it is a private contractual mechanism of the card schemes, the cardholder asks their bank to raise one, and the bank decides whether to.

Wrong: A chargeback reverses the original payment, so the same amount comes back. **Right:** It is a new financial transaction with its own clearing cycle, value date and conversion, so a sale that settled as 100 euros can return as 107.86 euros, and issuers can also bundle, split, or overlap with a refund already given.

Wrong: Winning the dispute puts the merchant back where it started. **Right:** The fee for receiving a dispute is normally non-refundable, the issuer can escalate and take the money a second time, and no victory removes the tick from the monitoring counter.

Wrong: Good delivery proof wins most disputes. **Right:** Proof of delivery decides non-delivery cases, helps in card-absent fraud only insofar as it places the goods with the cardholder, and is worth nothing at all under “not as described”, where nobody disputes that the parcel arrived.

Wrong: Refunding the customer once the dispute exists gets the chargeback off the books. **Right:** The count is taken when the dispute is filed, the merchant cannot refund outside the process while it is open, and a refund at that stage pays twice for nothing.

Wrong: Disputes filed as fraud are fraud. **Right:** A large share of card-absent fraud disputes are first-party misuse, a subscription harder to cancel than to start, a family member, an unrecognised statement descriptor, which is why both schemes have built programmes aimed squarely at it.

Wrong: Declining suspicious authorisations improves the VAMP ratio. **Right:** The denominator is settled transactions, so declines do not flatter the ratio; they only keep fraud out of the numerator.

Wrong: The fines are the thing to fear. **Right:** Fines are survivable; the terminated merchant files are not, and MATCH reason code 4 bites at 1 per cent of monthly Mastercard sales totalling 5,000 US dollars, below every monitoring threshold and indifferent to whether the chargebacks were later won.

Wrong: Arbitration is where a merchant with a strong case gets justice. **Right:** With case filing and ruling amounts commonly quoted between 500 and 600 US dollars, it is irrational at ordinary retail values and exists as a threat that disciplines the earlier stages.

Wrong: The merchant gets the scheme’s full response window. **Right:** The acquirer or gateway must reserve time to format and transmit the file, so the internal deadline is typically 7 to 21 days, and a missed deadline is a loss with no appeal.

49.99 Chapter summary in 20 lines

1. When you tap a card your bank pays the shop and bills you afterwards, so it is the bank’s money that moved and the bank that has both the standing and the plumbing to reverse it.
2. A chargeback therefore runs backwards from the ordinary intuition: the money is seized first and the argument happens afterwards.
3. For the merchant the first news is a debit of the sale plus a non-refundable dispute fee, arriving weeks or months after the goods were posted.
4. Contesting a dispute is called representment, costs a further fee refundable only on a win, and is decided by a stranger at the issuing bank two to three months later.

5. The question that stranger answers is fixed by the reason code the issuer chose, not by the merchant's account of what happened.
6. Evidence that does not answer the reason code is irrelevant rather than weak, and submitting a great deal of it is the commonest self-inflicted loss in the trade.
7. Visa runs allocation and collaboration workflows, Mastercard a single flow ending in arbitration, and American Express, being its own issuer, decides once and finally.
8. Deadlines are calendar days enforced by software, and the merchant's internal window is shorter than the scheme's because the acquirer must reserve time to transmit.
9. Chargeback in the United Kingdom is a scheme mechanism with no statutory basis; section 75 of the Consumer Credit Act 1974 and regulation 76 of the Payment Services Regulations 2017 sit alongside it and behave differently.
10. Because a chargeback is a new transaction, foreign exchange, bundled subscriptions, partial disputes and earlier refunds routinely make the returning amount differ from the original.
11. Winning is neither final nor free: issuers escalate, fees stay spent, and money moves outside the lifecycle through issuer credits and debits months later.
12. Every month the schemes divide a merchant's disputes by its transactions and compare the answer with a threshold.
13. Visa's VAMP ratio adds fraud reports to disputes over settled card-absent transactions, so events that moved no money count, and a transaction appearing in both feeds counts twice.
14. Mastercard divides this month's chargebacks by last month's transaction count, which punishes sharply seasonal merchants.
15. The counters are indifferent to outcomes by design, because what the schemes measure is prevention and waiting months for every result would make monitoring useless.
16. Crossing a threshold brings escalating monthly fines, a written remediation plan, and eventually the end of the acquiring relationship.
17. The real sanction is the terminated merchant lists, MATCH and the Visa Merchant Screening Service, whose quantitative triggers sit below every monitoring threshold and whose entries last five years.
18. That is why a merchant refunds a case it would win: a refund arriving before the dispute is a refund, and only one arriving afterwards is counted.
19. Hence the pre-dispute layer, Order Insight, Rapid Dispute Resolution and Ethoca, whose commercial value is the removal of a tick from a counter rather than the rescue of a sale.
20. And hence the direction of travel: Compelling Evidence 3.0 and First-Party Trust reward data captured at the moment of the order rather than prose written afterwards.

Sources, all consulted in August 2026: Visa's Acquirer Monitoring Programme fact sheet, its Compelling Evidence 3.0 merchant readiness paper and its press release of 1 April 2026 on dispute resolution services; Verifone Order Insight and Rapid Dispute Resolution; Mastercard's Chargebacks Made Simple guide, its Excessive Chargeback Programme guide and its June 2025 First-Party Trust announcement; Stripe's public documentation on disputes, monitoring programmes and high-risk merchant lists, which quotes the MATCH and VMSS standards verbatim, and its June 2025 dispute pricing notice; regulations 76, 77 and 100 of the Payment Services Regulations 2017 and section 75 of the Consumer Credit Act 1974 on legislation.gov.uk; UK Finance on chargeback and section 75 and its Annual Fraud Report of 15 June 2026; Financial Ombudsman Service pages on disputed transactions and goods bought on credit; and trade references from Chargeback Gurus, Chargeflow, Braintree and the Merchant Risk Council where scheme documents are not public, used only where consistent and flagged where not.

Reconciliation

50.0 What this chapter gives you

1. You will be able to explain why three records of the same week disagree, and why forcing them to agree is not the job.
2. You will be able to take a till total, a processor payout and a bank statement and account for every penny of the difference, naming fees, refunds, timing and breaks separately.
3. You will be able to tell a reason from a question, so that an explained difference is closed and an unexplained one is opened as a break with a date of birth and an owner.
4. You will be able to age a break register by value as well as by count, and say why the age of the oldest open item is the most diagnostic number in it.
5. You will be able to name the five assertions a reconciliation makes, completeness, existence, accuracy, cut-off and valuation, and show why equal totals test only the weakest of them.
6. You will be able to design a ledger that hands you the break list rather than making you hunt for it, using an immutable event log and clearing accounts held per counterparty and per currency.
7. You will be able to list the five identifiers that must be persisted at the moment of the event, because none of them can be reconstructed afterwards.
8. You will be able to assign an unexplained item to a known category quickly, whether it is fee netting, an authorisation never captured, a partial capture, a retranslation movement or a Direct Debit return.
9. You will be able to state what CASS 15 requires of a United Kingdom safeguarding institution on each reconciliation day, and what must happen by the end of the day a shortfall is found.
10. You will be able to explain why young payments firms die of being unable to prove where the money is rather than of losing it.

The plain version

Ravi runs a bakery on a side street in Coventry. Hardly anybody pays him in cash any more, so there are three separate written records of what happened in his shop last week, kept by three parties who have never spoken to one another.

The first is Ravi's own. His till logs every sale as it happens: the time, the amount, the item. That is a record of what he believes he sold.

The second belongs to the company that supplies his card machine. Each day they total what they collected on his behalf, subtract what they charge for the service, and send the remainder onward, publishing all of it in a weekly report. That is a record of what they believe they processed.

The third is his bank statement. It shows money landing in his account: a figure and a date, nothing more. That is a record of what actually arrived.

Three notebooks describing one week. Here is the thing that startles everybody the first time they meet it. The three notebooks will not agree. They will not agree this week, they will not agree next week, and no amount of diligence will make them agree on the first pass. The job is not to force them to agree. The job is to explain, line by line, precisely why they differ.

That job has a name. It is reconciliation, and it is simultaneously the least glamorous and the most load-bearing activity in finance.

Let us walk through Ravi's week, because the shape of the answer matters more than the numbers. His till says he took 386 card payments between Monday and Sunday, adding up to £4,812.60.

Day	Card sales	Value taken
Monday	49	£612.40
Tuesday	47	£588.15
Wednesday	56	£701.90
Thursday	52	£655.25
Friday	73	£912.60
Saturday	88	£1,104.30
Sunday	21	£238.00
Total	386	£4,812.60

The card company's report says three things the till does not. First, it charges him: one and a half per cent of each sale plus five pence per transaction. Second, it records a refund of £38.50, handed back on Wednesday to a customer whose wedding cake order was cancelled. That customer paid three weeks earlier, so the £38.50 appears nowhere in this week's sales but comes out of this week's money. Third, and this is the smallest and most interesting difference, the report shows 385 payments, not 386. One sale, £18.75 on Saturday afternoon, is simply not there.

So the report reads: £4,793.85 processed across 385 sales, less a refund of £38.50, less fees of £91.16, leaving £4,664.19 to be paid over.

Finally, the bank statement. Money from a card machine does not arrive the instant the card is tapped. It arrives a couple of working days later, and a weekend's takings are bundled into the Monday run. These are the deposits that landed.

Day money arrived	Amount	What it was for
Monday	£498.11	The previous week's Thursday
Tuesday	£2,043.77	The previous week's Friday, Saturday and Sunday
Wednesday	£600.76	This Monday
Thursday	£576.98	This Tuesday
Friday	£650.07	This Wednesday, less the £38.50 refund
Saturday	none	
Sunday	none	
Total	£4,369.69	

Three numbers for one week: £4,812.60, £4,664.19 and £4,369.69. All three are correct. Reconciling means turning that spread into arithmetic that closes.

Take the first gap. £4,812.60 minus £4,664.19 is £148.41, and Ravi can account for every penny: £91.16 of fees, £38.50 of refund, £18.75 for the sale that vanished. The gap is now explained. It is not a problem. It is a set of reasons, each with a name.

The second gap needs one more piece of information: what each day was worth after the card company took its cut. Monday's £612.40, less £11.64 of fees, is £600.76, which is precisely what landed on Wednesday. Do the same for the rest of the week and the seven days come to £4,702.69 net, which less the £38.50 refund is £4,664.19, exactly what the card company says it is paying.

Now the second gap closes. The card company paid £4,664.19; the bank shows £4,369.69 arriving; the difference is £294.50, and it is pure timing. Four days of this week's money had not landed when Sunday ended: Thursday at £642.82 net, Friday at £895.26, Saturday at £1,064.92 and Sunday at £233.38, which is £2,836.38 in transit. Against that, £2,541.88 of last week's money landed on Monday and Tuesday. Subtract the second from the first and you get £294.50. Explained.

That is the entire discipline in miniature. Ravi did not make the numbers match. He wrote down a reason for every difference, and the reasons added up.

Except for one. The fees are a reason. The refund is a reason. The timing is a reason. The missing £18.75 is not a reason; it is a question. A customer bought two loaves and a coffee cake, the machine said approved, and no money ever moved. The likeliest explanation is that the payment was reserved against the card but never claimed before the machine closed its day, so the reservation expired and the card released the money back. Ravi is out the bread and out the cash.

An unexplained difference of this kind is called a break, and the most important fact about a break is not what caused it but how old it is.

A break one day old is normal. Money is in transit, files arrive late, somebody will look at it tomorrow. A break ninety days old is a different animal: a standing statement in Ravi's accounts that he owns £18.75 which does not exist, quietly making him richer on paper than in life. Do that a thousand times and the gap between the books and the bank stops being a rounding error and becomes a hole.

This is why serious operations do not merely count their breaks; they age them. Every unexplained difference gets a date of birth, an owner with a name, and a bucket: under two days, under a week, under a month, over a month. A business with a flat ageing profile knows where its money is; a business with a lengthening tail does not.

Then scale it. Ravi had one break in 386 payments, roughly a quarter of one per cent, which is an unremarkable rate. At a hundred thousand payments a day, the same rate produces around 260 new breaks daily and 1,800 a week, arriving forever. At ten minutes of attention each, that is more than a full working week of somebody's time every week just to stand still.

Which brings us to the sentence this chapter exists to deliver. Young financial technology companies do not usually die of fraud, and they do not usually die of being hacked. They die because they could not prove where the money was. Not because they lost it. Because they could not prove it. The two are different, and the second is fatal on its own.

Where the plain version stops being true

There are not three notebooks, and “the processor report” is not one document. A real payments business reconciles between six and fifteen sources at once: an order system, a gateway, one or more acquirers, the card schemes, a wallet provider, a foreign exchange counterparty, a Bacs bureau or sponsoring bank for Direct Debits, an operating bank account, and, if the firm is regulated, a safeguarding account whose contents are legally not its own. Worse, a single processor issues several reports covering the same period, deliberately keyed on different dates, which will never tie to one another without work. Stripe’s public documentation, at the time of writing in August 2026, shows the seam plainly: a balance transaction carries an `available_on` date, when its net funds become available in the Stripe balance, and separately an `automatic_payout_effective_at` date, when the payout is expected to reach your bank account. Reconcile on the first and you get one population; on the second, another, both correct. Adyen’s settlement details report, again at the time of writing, changes shape with the acquirer: where the acquirer supplies transaction information at interchange level, the Commission column is empty and the Markup, Scheme Fees and Interchange columns are populated instead; where it does not, Commission carries the total and those three are blank. Identical economics, different structure, and a reconciliation written against one shape will silently produce nonsense against the other.

Matching is not the control. Explanation is. Two systems can match on amount and reference and both still be wrong, because both were fed by the same faulty upstream event. Two errors of opposite sign can cancel, producing a difference of zero across a population containing a five-hundred-pound overstatement and a five-hundred-pound understatement. This is why mature reconciliation runs on gross, signed values and never on the net, why a reported difference of nil is the beginning of the check rather than the end, and why the useful metric is not “what proportion of lines matched” but “what is the total absolute value of items I cannot yet explain, and how old is the oldest one”.

“As at close of business” describes a moment that never existed. The bakery assumed a single instant at which all three records can be photographed. There is none: the United Kingdom’s payment systems settle at fixed and different times of day. On the Bank of England’s published summary of the Real-Time Gross Settlement daily timetable, as it stood at the time of writing in August 2026, CHAPS settlement starts at 6am, Faster Payments settles three times each working day at 7am, 1pm and 5pm, Bacs settles at 9.30am, LINK at 11am, Mastercard at 11.15am, Visa at 12.15pm and the Image Clearing System at 4.30pm, with CHAPS closing for customer payments at 5.40pm and for financial institution payments at 6pm. A payment can therefore be irrevocable in one system, provisional in another and invisible in a third at the same wall-clock second. A bank statement line also carries at least two dates, booking and value, frequently not the same day. A reconciliation that does not define in writing which clock and which date field it cuts on produces a different answer depending on when it is run, which is the same as producing no answer.

Once the money belongs to somebody else, this stops being bookkeeping and becomes law. In the bakery the money was Ravi’s, so a break was his problem alone. The moment a firm holds funds for customers, the arithmetic acquires a statutory deadline and a regulator at the other end of it. In the United Kingdom, at the time of writing in August 2026, authorised payment institutions and electronic money institutions must safeguard relevant funds under regulation 23 of the Payment Services Regulations 2017 and regulations 20 to 22 of the Electronic Money Regulations 2011, and since 7 May 2026 those requirements have been supplemented by a new chapter of the FCA’s Client Assets sourcebook, CASS 15, which converts reconciliation from good practice into a rule with a same-day remedy attached. Under CASS 15 as it stands at the time of writing, a shortfall found by an internal safeguarding reconciliation must be paid in by the end of the day on which the reconciliation was performed,

out of the firm's own money if customer money cannot be reached. No version of that obligation can be met with a spreadsheet somebody updates on Fridays.

The technical version

■ On dating

Everything below is stated as accurate at the time of writing, in August 2026, and this corner of the subject moves fast. The FCA's Supplementary Regime for safeguarding came into force on 7 May 2026, three months before this was written, following Policy Statement PS25/12 of 7 August 2025. In that statement the FCA declined, for now, to proceed with the end-state proposals it had consulted on, rebranding them the Post-Repeal Regime, so the rule references below have a known revision ahead of them. The Payment Systems Regulator still existed as a separate body as this was written: the Government announced in March 2025 that it intended to abolish the PSR and consolidate its functions primarily within the FCA, consulted in September 2025, and on 21 April 2026 published a response confirming full consolidation and noting that this requires primary legislation, to be brought forward as soon as Parliamentary time allows. Check the Handbook, the Bank of England's timetable and your own scheme documentation before relying on any figure here.

■ What a reconciliation actually asserts

A reconciliation is not a comparison. It is an assertion, and most weak processes fail by never having stated the assertion at all. It has five parts. Completeness: every event that occurred appears in the record. Existence: every event in the record occurred, and nothing has been invented or duplicated. Accuracy: each event is recorded at the right amount, in the right currency, against the right counterparty. Cut-off: each event falls in the right period, on a defined and documented clock. Valuation: where an amount has been translated between currencies, or is subject to later adjustment, the value carried is the right one on the right basis.

Notice that "the two numbers are the same" tests a weak form of accuracy and nothing else. A process producing a matching total while a hundred transactions are absent from both sides has passed a test never worth taking, which is why a processor restatement of a closed period must be posted as a versioned, dated correction referencing the original rather than an edit to it.

■ The three legs, named properly

Leg one is the internal ledger. In a serious build this is not the order table. It is a double-entry sub-ledger fed by an append-only, immutable event log, in which a payment is a state machine with explicit transitions: initiated, authorised, captured in full or in part, settled, refunded in full or in part, disputed, represented, written off. Each transition posts a balanced journal. Critically, the design includes a clearing account, sometimes called cash-in-transit or settlement suspense, held per counterparty and per currency, debited when a payment is captured and credited when the cash lands. Its balance is, by construction, the exact population of open items. Designed this way, you do not hunt for breaks; the ledger hands you the list.

Leg two is the processor or acquirer record. It arrives as files or API objects, and it is where the commercial complexity lives. Under an interchange-plus-plus arrangement it decomposes the cost of a transaction into interchange, paid to the card issuer; scheme fees, paid to Visa or Mastercard; and the acquirer's own markup. In the United Kingdom, at the time of writing, the UK Interchange Fee Regulation caps interchange on consumer card transactions where the merchant, the acquirer and the card issuer are all within the UK, at 0.2 per cent of transaction value for consumer debit cards including prepaid cards and 0.3 per cent for consumer credit cards, with some American Express cards

outside the regime. Commercial cards and most cross-border transactions fall outside those caps, and UK to European Economic Area card-not-present interchange rose substantially after the UK left the European Union, prompting a Payment Systems Regulator market review of cross-border interchange fees. That is why a fee reconciliation built on one blended assumed rate drifts the moment a business acquires foreign or corporate customers. On a blended contract the same economics arrive as one undifferentiated deduction: simpler to reconcile, impossible to audit.

Leg three is the bank record. In modern form that means ISO 20022: a camt.053 end-of-day statement, a camt.052 intraday report, a camt.054 debit or credit notification. Legacy estates still consume SWIFT MT940 and MT942, or BAI2 in North American systems. All distinguish booking date from value date, and all carry remittance information of wildly varying quality, which is why bank-side matching is harder than processor-side matching. A processor gives you its own reference on every line. A bank frequently gives you eighteen characters of free text typed by a stranger.

A reconciliation is only as strong as the weakest identifier in the chain, so five keys must be persisted as each event happens: the internal payment identifier and idempotency key, the processor's reference, the modification or capture reference for anything amending an earlier event, the payout or settlement batch identifier, and the bank's end-to-end reference on the credit. None can be reconstructed later, and the break a missing one causes will be unclosable rather than merely open.

■ Why they never match: a taxonomy

These causes account for the overwhelming majority of differences in a card-and-bank-transfer business, and they are worth learning as a taxonomy, because break investigation is mostly the work of assigning an unexplained item to a known category quickly.

Cause	Where it shows	Typical age	How it clears
Fees netted at settlement	Ledger gross, payout net	Same day	Post fee journal from itemised fee data
Settlement timing	Processor paid, bank silent	1 to 3 days	Clears on arrival; monitor the tail
Authorised but never captured	Ledger sale, no processor line	Days to weeks	Write off or re-attempt; fix capture logic
Partial capture	Amounts differ on one reference	Same day	Model capture separately from authorisation
Refund crossing a period	Payout smaller than sales	1 to 30 days	Link refund to original charge, not to the day
Chargeback and representment	Money out, then possibly back	30 to 180 days	Separate dispute sub-ledger
Foreign exchange	Amounts differ by a few per cent	Ongoing	Retranslate; recognise the difference explicitly
Duplicate submission	Two ledger events, one payout line	Same day	Idempotency keys; deduplicate at source
Processor restatement	A closed period changes	Weeks	Versioned correction, never an overwrite
Lump-sum bank credit	One credit, many payouts	Same day	Match many-to-one on batch reference
Direct Debit unpaid	Collection made, then returned	2 to 10 days	Consume the returns file as a first-class event
Missing or partial file	A whole day absent	Hours	File-level completeness control

Four deserve expansion.

Fee netting. Fees break reconciliations not because they exist but because they are deducted at a different point from where they are incurred. Adyen documents that per-transaction payment fees are booked when captured funds are settled, while its own processing fees are invoiced and adjusted monthly, with the adjustment booked separately. A business reconciling on one monthly total therefore finds two fee populations arriving on two cadences, and the difference between them is not an error. Accrue from itemised data, not an assumed rate.

Authorisation versus capture. These are separate events with separate lifetimes, and treating them as one is probably the commonest structural defect in an early-stage ledger. An authorisation places a hold on a customer's available balance; a capture claims it. Authorisations expire. Captures can be partial, so an authorisation for £120 may be captured at £97.40 when an item is out of stock, and a system recognising revenue on authorisation is permanently overstated by the difference. Some schemes and acquirers support incremental authorisation and over-capture within a tolerance, so a captured amount can legitimately exceed the amount authorised, and a validation rule forbidding that will reject good data.

Foreign exchange. A single cross-border card payment can involve four amounts: the transaction currency the customer sees, the presentment currency in which the scheme presents it to the acquirer, the settlement currency in which the acquirer pays, and the functional currency of the books. Each conversion carries its own rate and its own rounding to the minor unit, struck at different moments. Under IAS 21, or Section 30 of FRS 102 for most UK small and medium-sized companies, monetary balances are retranslated at the closing rate with the resulting exchange differences taken to profit or loss. So an unreconciled receivable in euros changes value every month without anybody touching it, and a reconciliation that does not separate genuine breaks from retranslation movement produces an unreadable break register within a quarter. Refunds are worse: a sale converted at January's rate and refunded in April is refunded at April's rate, so a full refund of a full sale does not net to zero and never will.

Direct Debits. For anyone operating on Bacs there is a whole reconciliation layer with no card equivalent, and it is the one most often bolted on late. Bacs runs a fixed three-working-day cycle: day one is input day, when the service user submits the file, the transmission window closing at 22.30 for processing on the following processing day; day two is processing day, when Bacs validates and distributes; day three is entry day, when the payer's account is debited and the service user's credited. Nothing has moved on days one or two, so a ledger recognising cash on submission is wrong for two days on every collection. Then the returns arrive. ARUDD, the Automated Return of Unpaid Direct Debits, reports collections that failed and are being returned, with a reason code attached; ARUCS does the equivalent for unapplied credits; ADDACS carries amendments to and cancellations of instructions; AWACS advises that account details used for a credit were wrong; and indemnity claims, notified through DDICA, reverse a collection under the Direct Debit Guarantee at the payer's bank's initiative, long after the event. Each is a file, each is a financial event, and each must be consumed as a first-class ledger entry rather than read by a human and typed into a spreadsheet. A firm that reconciles Bacs collections but not Bacs returns will show cash it has already given back.

■ Break investigation and ageing

A break register is a small database and should be treated as one, not as a tab in a workbook. Every open item carries a unique identifier that never changes, the date first detected, the source systems on both sides, the signed amount and currency, a category from a controlled list, a named individual owner, an expected clearing date, a timestamped investigation trail, and a resolution code recorded when it closes.

Three rules separate a register that works from one that is theatre. First, a break with no named owner is not a break, it is a rumour: ownership sits with an individual rather than a team, and the register reports the count of unowned items, which should be zero. Second, ageing is measured gross, and by value as well as by count, because ten thousand two-pence breaks and one four-hundred-thousand-pound break are different problems and a register reporting only counts hides the second behind the first; workable buckets are nought to one day, two to five days, six to thirty days, thirty-one to ninety days and over ninety days, in reconciliation days rather than calendar days. Third, explained is not the same as cleared. An item can be fully understood and still open, because the cash has not arrived. An item can be cleared by a manual journal and still not understood, which is far more dangerous, because the journal removes the symptom and leaves the defect. Every clearing journal should require a second approver and should be counted, and a rising count is among the earliest reliable signs that a reconciliation process has begun to fail.

Ageing thresholds should trigger action, not merely reporting. A serviceable ladder runs: over five days, the owner records a root cause hypothesis; over thirty days, the item escalates to the finance lead and a provision is considered; over ninety days, the item is written off to profit or loss but deliberately kept open in the register, because writing an item off is an accounting decision and closing it is an operational lie. Items written off but held open are the raw material for root cause analysis. Items closed on write-off vanish, and the defect that generated them runs forever.

Reconciliation is also a fraud control, and quietly one of the best available, because internal misappropriation almost always leaves an anomaly in exactly this data: a payout corresponding to no settled transaction, a refund with no matching original, a clearing account that never returns to a plausible balance. The defensive value comes from the segregation of duties around it rather than from the matching itself. Whoever can create a payment must not be whoever can approve a reconciliation adjustment, and neither should be able to amend the break register without leaving a record.

■ Where the law arrives

For a UK firm holding customer money this is a regulatory obligation with named rules, deadlines and notification duties. What follows is the position at the time of writing, in August 2026.

Authorised payment institutions, small payment institutions that elect to safeguard, electronic money institutions and credit unions issuing electronic money are, as safeguarding institutions, subject to CASS 15, in force since 7 May 2026. Firms providing only payment initiation or account information services do not hold relevant funds and are outside it. CASS 15.8 is the records, accounts and reconciliations chapter, and its architecture is worth knowing precisely.

An internal safeguarding reconciliation compares what the firm's own records say it owes customers, the safeguarding requirement, against what those records say it holds, the safeguarding resource. CASS 15.8.19R requires this as frequently as necessary and no less than once each reconciliation day. That term was a late and practically significant change during the consultation: the Handbook glossary defines a reconciliation day, at the time of writing, as a business day that is not a Saturday or Sunday, Christmas Day, Good Friday or a bank holiday in any part of the United Kingdom, and, where a

reconciliation requires anything to be done by reference to a market outside the United Kingdom, not a day on which that market is not normally open for business. The standard method also includes a forward-looking comparison of the firm's D+1 segregation requirement against its D+1 segregation resource. A firm relying on an insurance policy or guarantee unlimited in cover need not perform an internal safeguarding reconciliation at all, but must record a daily calculation of its safeguarding requirement.

An external safeguarding reconciliation compares the firm's internal records against statements or confirmations from the third parties actually holding the money. Under CASS 15.8.42R it too must be performed as frequently as necessary and no less than once each reconciliation day, and as soon as reasonably practicable after the date to which it relates.

The consequences of a discrepancy are where the rules bite hardest. Where an internal reconciliation reveals a difference between resource and requirement, CASS 15.8.50R requires the firm to determine the reason and then ensure either that any shortfall is paid into a relevant funds bank account or invested in relevant assets as soon as possible and in any case by the end of the day on which the reconciliation is performed, or that any excess is withdrawn. CASS 15.8.51R applies the same end-of-day deadline to a deficiency in the D+1 segregation resource and goes further: where relevant funds cannot remedy it, the firm must use its own, even where that produces a surplus, with the position corrected by later reconciliations. Where an external reconciliation reveals a discrepancy, CASS 15.8.56R requires investigation and all reasonable steps to resolve it without undue delay, unless it arises solely from timing differences between the two parties' accounting systems. And CASS 15.8.57R provides that where the firm cannot immediately resolve such a discrepancy, and one set of records indicates a need for a greater amount of relevant funds or assets than another, it must assume until the matter is finally resolved that the record showing the greater amount is accurate, and pay its own funds in accordingly. The regulatory default is to assume you are short.

Then the notifications. CASS 15.8 requires a safeguarding institution to inform the FCA in writing without delay if its internal records and accounts of relevant funds are materially out of date, inaccurate or invalid; if it will be unable to, or materially fails to, conduct an internal or an external safeguarding reconciliation in compliance with the rules; if it will be unable to, or materially fails to, pay in a shortfall or withdraw an excess after an internal reconciliation; if it will be unable to, or materially fails to, identify and resolve discrepancies after an external reconciliation; or if it becomes aware that at any time in the preceding twelve months the amount of relevant funds safeguarded was materially different from the total aggregate amount it was required to safeguard. That last limb deserves reading twice. Discovering historic under-safeguarding is itself notifiable, so an honest clean-up of a legacy break population is a conversation with the regulator, not a quiet weekend of journals.

Around the reconciliation sit the monitoring obligations. Under SUP 16.14A a safeguarding institution must submit a safeguarding return to the FCA within fifteen business days of each month end, and that return asks directly whether internal safeguarding reconciliations were carried out every reconciliation day during the period. An annual safeguarding audit by a qualified auditor is required, with an exemption at the time of writing for firms not required to safeguard more than £100,000 of relevant funds at any point over a period of at least 53 weeks; the report is due within four months of the period covered, extended to six months for the first report under the new regime. Under SUP 3A.9 the auditor must confirm whether the firm has maintained systems adequate to enable compliance with the relevant funds regime. Oversight must be allocated to a single named director or senior manager.

Investment firms holding client money live under the older parallel regime in CASS 7.15, which requires an internal client money reconciliation each business day under CASS 7.15.15R, based on the

records as at close of business on the previous business day. CASS 15 borrows heavily from CASS 7, which is useful rather than merely tidy: the enforcement history and audit methodology built up around CASS 7 over more than a decade are largely transferable.

The direction of travel is visible elsewhere. The Bank of England's draft Code of Practice for systemic sterling-denominated stablecoin issuers, published in 2026, requires at paragraph 3.5 that a systemic issuer carry out an internal safeguarding reconciliation as regularly as necessary and at least once each business day. Where money is held for others, daily reconciliation is becoming the floor across UK financial regulation. And the Payment and Electronic Money Institution Insolvency Regulations 2021 created a special administration regime whose objectives include returning relevant funds to customers as soon as is reasonably practicable, exactly the objective unreconciled records defeat.

■ Why young firms actually die here

This part of the chapter is an argument rather than a description, and it should be made without hedging.

The FCA's own published evidence is stark. In consultation paper CP24/20, published in September 2024, it reported that twelve payments firms which safeguarded relevant funds became insolvent between the first quarter of 2018 and the second quarter of 2023, that the weighted average shortfall in client funds across those cases was 65 per cent, and that seven of the twelve showed shortfalls over 50 per cent. Its table of selected averages put mean client funds owed at £11.3m against mean total assets available of £4.0m, a shortfall of £7.3m; for electronic money institutions alone, £16.3m owed against £3.3m available, a shortfall of £13.0m or 80 per cent. Of the six cases in which some funds had been distributed, a first distribution took on average over two years. The FCA was careful to say these shortfall figures are likely underestimates, because the data available did not record safeguarded funds and total assets at insolvency had to serve as a proxy. The same paper recorded that in 2023 it opened supervisory cases about safeguarding arrangements at approximately 15 per cent of the firms that safeguard.

Read those numbers with the reconciliation lens on. A weighted average shortfall of 65 per cent is not embezzlement; embezzlement on that scale across twelve firms would have been a national scandal. It is what happens when firms do not know, at any given moment, what they owe. The requirement side drifted from the resource side over months and years, nobody computed the difference daily, and the gap was first measured on the day an administrator arrived to measure it. The two years to a first distribution is the same failure from the other end: an insolvency practitioner cannot pay customers until they know who is owed what, and the records did not say.

The proximate causes are almost always mundane. Corporate and customer money commingled in one operating account, because opening a safeguarding account was slow and the founders were busy. Revenue recognised on authorisation rather than capture, so the profit and loss account told a happier story than the bank did. Fees accrued at a blended estimate rather than from itemised data, so the estimate drifted into a suspense account nobody owned. A break register that was a spreadsheet, so its oldest item was whatever the last person to open the file had not deleted. Customer float treated as working capital, because float behaves exactly like money right up until the customers ask for it back at the same time.

The legal machinery makes the consequences harder rather than easier. In the administration of the electronic money institution Ipagoo, the Court of Appeal held in 2022 that the Electronic Money Regulations do not impose a statutory trust over funds received from e-money holders, while also holding that a shortfall in the customer asset pool should be topped up. A subsequent judgment following the insolvency of Allied Wallet, [2022] EWHC 1877 (Ch), extended that reasoning to the Payment Services

Regulations, so relevant funds held with authorised payment institutions are likewise not subject to a statutory trust. The FCA's assessment in CP24/20 was that this leaves practitioners seeking directions from the court on the treatment of the safeguarded asset pool, raising costs and delaying returns. All of that litigation is downstream of one fact: the records did not make it obvious which money was whose. Wirecard illustrates the same point at scale. When it collapsed in June 2020 the missing €1.9 billion, said to be held in trustee accounts in Asia, turned out not to exist: an asset asserted on a balance sheet that no independent third-party record ever confirmed.

So why do founders under-invest in this, consistently, across an industry that has watched other firms die of it?

Because reconciliation has no demo. It cannot be shown to an investor, and when it works perfectly its entire visible output is a report saying nothing is wrong. Every hour spent on it is an hour not spent on something a customer can see.

Because its cost curve is deceptive. At a hundred transactions a day, reconciliation genuinely is a spreadsheet and half an hour on a Monday, and the founder who says "we will do this properly later" is right that it is not urgent today. At a hundred thousand transactions a day it is a team, a data pipeline and a rules engine, and the transition does not happen gradually. It happens in the eight weeks after a large partner goes live, on top of a data model designed when none of it mattered. Migrating from "our ledger is the payments table" to an immutable double-entry event log with clearing accounts per counterparty and currency is a rewrite of the most sensitive code in the business, invariably attempted while that code is under maximum load.

Because the early signals look benign. A small unexplained surplus reads as a pleasant surprise rather than a defect, and a growing business generates surpluses constantly through timing alone. Inflows exceed outflows, so the safeguarding account never looks empty, and the difference between what is held and what is owed is masked by growth for as long as growth continues. When growth stops the hole is instantly visible, and it has been there for two years.

Because the first person who would notice has usually not been hired. Reconciliation sits in the seam between engineering, which owns the data and not the balance sheet, and finance, which owns the balance sheet and cannot query the data. It is nobody's specialism until somebody is made responsible by name, which under CASS 15, at the time of writing, is a rule.

And because the consequences are back-loaded and then arrive at once: a qualified audit opinion; a due diligence process that stalls because the acquirer's accountants cannot tie revenue to cash; a notification to the FCA under CASS 15.8 that a reconciliation was not performed; a monthly safeguarding return that has to answer "no"; or a report from a skilled person, required under section 166 of the Financial Services and Markets Act 2000 and paid for by the firm. Any one of those can consume a management team for six months. All are cheaper to prevent than to survive.

■ What competence looks like

The automatic match rate at T+1, by count and by value, is the headline health metric; in a well-instrumented card business it should sit comfortably above 99 per cent by value on the day after settlement. The absolute value of unexplained items in each ageing bucket matters far more than the net figure. The age of the oldest open item is the most diagnostic number in the register and belongs in the board pack. Mean time to clear, by break category, tells you which defect to fix next. The count of manual clearing journals, and their aggregate value, tells you how much of the match rate is real. And the number of reconciliation days on which the reconciliation was not completed should be zero, because for a safeguarding institution a failure to perform it is itself notifiable.

Around those numbers sit four practices. Persist every identifier at the moment of the event, because none can be reconstructed afterwards. Use idempotency keys on every mutating call, because duplicate submission is the cheapest break to prevent and among the most expensive to unwind. Reconcile file-level completeness before line-level content, because a missing file produces a thousand phantom breaks and one real problem. And write the reconciliation before the product ships, because the data model that makes reconciliation possible is the same one that makes the product correct.

■ What will date this chapter

Three things, and a reader in 2029 should check all of them before quoting any of the above.

The CASS 15 rule set is explicitly staged. The Supplementary Regime in force since 7 May 2026 is, in the FCA's own framing, an intermediate step. The end-state proposals consulted on in CP24/20, including a statutory trust over relevant funds and a requirement to receive funds directly into a designated safeguarding account, were not taken forward in that form; renamed the Post-Repeal Regime, they await further consultation once HM Treasury has repealed and replaced the Payment Services Regulations 2017 and the Electronic Money Regulations 2011. Rule numbers will move, and the substance may move further.

The regulatory perimeter is being redrawn. Abolition of the Payment Systems Regulator and transfer of its functions to the FCA was confirmed as government policy in the April 2026 consultation response, but requires primary legislation not enacted at the time of writing, so any sentence here naming the PSR as a live regulator has a foreseeable expiry date.

And the plumbing is changing underneath. The Bank of England confirmed in February 2026 that it intends to move the start of CHAPS settlement from 6am to 1.30am in September 2027, and has since consulted on further steps towards near twenty-four-hour, seven-day settlement on its renewed RTGS platform, with Bacs, Faster Payments, the Image Clearing System, Visa, Mastercard and LINK expected to keep their existing timelines under that first extension. The timetable quoted earlier is therefore a snapshot of a moving object. Migration to ISO 20022 continues to widen the structured data on the bank leg, which will make bank-side matching materially easier than the account given here. When the timetable moves, every cut-off assumption built on it moves too, which is a useful reminder that a reconciliation is not a piece of software but a standing relationship with external clocks you do not control.

50.98 Common wrong ideas

Wrong: Reconciliation means making the numbers match. **Right:** It means explaining, line by line, exactly why they differ; Ravi's three totals were all correct, and the discipline is writing a reason against each difference until the reasons add up.

Wrong: There are three records to reconcile. **Right:** A real payments business reconciles between six and fifteen sources, and a single processor issues several reports covering the same period deliberately keyed on different dates, which will never tie to one another without work.

Wrong: A reported difference of nil means the reconciliation passed. **Right:** Two errors of opposite sign cancel, so nil is the beginning of the check; mature processes run on gross signed values and measure the total absolute value of items they cannot yet explain.

Wrong: "As at close of business" is a single moment at which all the records can be photographed. **Right:** CHAPS, Bacs, Faster Payments, LINK, Visa, Mastercard and the Image Clearing System settle at fixed and different times of day, and a bank line carries booking and value dates, so a reconciliation that does not define its clock in writing gives a different answer depending on when it is run.

Wrong: Authorisation and capture are one event. **Right:** They have separate lifetimes; authorisations expire, captures can be partial, and some schemes permit over-capture within a tolerance, so recognising revenue on authorisation overstates permanently and a rule forbidding over-capture rejects good data.

Wrong: A full refund of a full sale nets to zero. **Right:** Across currencies the sale converts at one date's rate and the refund at another's, so it does not net to zero and never will.

Wrong: An item cleared by a manual journal is an item resolved. **Right:** Explained is not the same as cleared and cleared is not the same as understood; the journal removes the symptom and leaves the defect, which is why every clearing journal needs a second approver and a rising count of them is an early sign of failure.

Wrong: A break that has been written off can be closed. **Right:** Writing an item off is an accounting decision and closing it is an operational lie; items written off but held open are the raw material for root cause analysis.

Wrong: A small unexplained surplus is a pleasant surprise. **Right:** A growing business generates surpluses constantly through timing alone, which masks the gap between what is held and what is owed for exactly as long as growth continues.

Wrong: A legacy safeguarding shortfall can be tidied up quietly over a weekend. **Right:** CASS 15.8 makes it notifiable if at any time in the preceding twelve months the amount safeguarded was materially different from the amount required, so an honest clean-up is a conversation with the regulator.

50.99 Chapter summary in 20 lines

1. A card-accepting business has several separate records of the same week, kept by parties who have never spoken to one another, and they will not agree on the first pass.
2. Reconciliation is not the work of making them agree; it is the work of explaining every difference by name.
3. Ravi's till, his card company's report and his bank statement differ by fees, a refund from an earlier period, settlement timing and one sale that never became money, and all three totals are correct.
4. A difference with a name is a reason and is closed; a difference without one is a break and is opened.
5. The most important fact about a break is not its cause but its age, which is why serious operations age their register by value as well as by count.
6. The plain picture understates the problem, because real firms reconcile between six and fifteen sources and one processor issues several reports of a period keyed on different dates.
7. Matching is not the control, since two systems fed by the same faulty upstream event can agree while both are wrong, and errors of opposite sign cancel.
8. There is no single moment called close of business, because the United Kingdom's payment systems settle at fixed and different times and bank lines carry two dates.
9. A reconciliation asserts completeness, existence, accuracy, cut-off and valuation, and equal totals test only a weak form of accuracy.
10. Done properly, the internal leg is a double-entry sub-ledger over an immutable event log, with clearing accounts per counterparty and currency whose balance is by construction the population of open items.
11. The processor leg carries the commercial complexity, interchange, scheme fees and markup, which is why a fee accrual built on one blended assumed rate drifts the moment foreign or corporate cards arrive.

12. The bank leg is the hardest to match, because a processor gives you its own reference on every line and a bank frequently gives you free text typed by a stranger.
13. Five identifiers must be persisted as each event happens, because a missing one makes the resulting break unclosable rather than merely open.
14. Most differences fall into a known taxonomy, of which fee netting, authorisation versus capture, foreign exchange and the Bacs returns files deserve particular attention.
15. Once the money belongs to somebody else this stops being bookkeeping and becomes law, and United Kingdom safeguarding institutions have fallen under CASS 15 since 7 May 2026.
16. CASS 15.8 requires internal and external safeguarding reconciliations on each reconciliation day, with any shortfall paid in by the end of the day the reconciliation was performed, from the firm's own money if customer money cannot be reached.
17. Where a discrepancy cannot be resolved the firm must assume the record showing the greater amount is accurate, and a whole list of failures, including historic under-safeguarding, is notifiable to the FCA.
18. The FCA's own evidence records twelve insolvent safeguarding firms with a weighted average shortfall of 65 per cent and a first distribution taking on average over two years, which is what not knowing what you owe looks like from the outside.
19. Founders under-invest because reconciliation has no demo, its cost curve jumps rather than climbs, its early signals look benign, and nobody owns it until somebody is named.
20. Young financial technology firms rarely die of fraud or of being hacked; they die because they could not prove where the money was, and being unable to prove it is fatal on its own.

Sources, all consulted in August 2026: the FCA Handbook, CASS 15.8, CASS 7.15, SUP 3A.9, SUP 16.14A and the glossary definition of reconciliation day; FCA Policy Statement PS25/12 and Consultation Paper CP24/20, including Table 3 and its methodology caveat; the Payment Services Regulations 2017, the Electronic Money Regulations 2011 and the Payment and Electronic Money Institution Insolvency Regulations 2021; the Bank of England's summary of the RTGS daily timetable, its 2026 papers on extending RTGS and CHAPS settlement hours, and its 2026 draft Code of Practice for systemic sterling stablecoin issuers; Bacs Payment Schemes Limited and Pay.UK on the Bacs cycle, the 22.30 cut-off and the ARUDD, ARUCS, ADDACS, AWACS and DDICA services; the Payment Systems Regulator on the UK Interchange Fee Regulation; HM Treasury's April 2026 response on payment systems regulation; Ipagoo LLP in the Court of Appeal and Allied Wallet [2022] EWHC 1877 (Ch); Stripe and Adyen developer documentation; and IAS 21 and Section 30 of FRS 102.

Fraud Detection

51.0 What this chapter gives you

1. You will be able to explain why a fraud system makes two kinds of mistake and why only one of them ever generates an incident report.
2. You will be able to compute the decline threshold for a given basket from the loss given fraud and the margin forgone, and say why it lands near one in three rather than nine in ten.
3. You will be able to show how a rule with seventy-five per cent recall and five per cent precision can make a business nearly eight times worse off while appearing all the while to work.
4. You will be able to distinguish precision, recall and specificity, and explain why at a fraud rate of 0.4 per cent almost all the leverage sits on the false positive rate.
5. You will be able to name the four different problems the word fraud covers, and say which of them no device signal will ever help with.
6. You will be able to separate signals supplied by the party being judged from evidence that party does not produce, and layer the two accordingly.
7. You will be able to say who bears the loss at each of the five decision points in a card transaction, and how United Kingdom strong customer authentication moves it.
8. You will be able to explain why a fraud rate below the reference rate is a commercial asset under the transaction risk analysis exemption, and what happens after two consecutive quarters above it.
9. You will be able to size a manual review queue by economics rather than headcount, and say when a human is worth roughly two pounds an order.
10. You will be able to explain why a model trained only on approved traffic grows confident about a world it has stopped observing, and what a holdout costs and buys.

The plain version

Imagine a doorman on a busy Friday night. He has three answers, not two. He can wave you in, he can turn you away, or he can say “hang on a second” and check something. That third answer is the interesting one, and we will come back to it.

He cannot ask everybody for identification, because the queue would stretch round the block and half of it would go somewhere else. He cannot let everybody in, because after enough bad nights the venue loses its licence. So he reads people. He notices that you arrived alone but keep glancing back at a group, that this is the fourth time in twenty minutes somebody has come to the door in the same brand-new jacket, that you have been here on eleven previous Saturdays and never caused trouble. None of those observations is proof. Every one shifts the odds.

Now here is the part everybody gets wrong. The doorman makes two completely different kinds of mistake and only one is visible. If he lets a troublemaker in, there is a fight and everyone remembers.

If he turns away an ordinary person who came out for a quiet drink, nothing happens at all. No incident, no report, no smashed glass. That person goes somewhere else, tells their friends the place is unwelcoming, and never comes back. The mistake is real, expensive and completely silent.

Fraud detection is that doorman, wired into a computer, running thousands of times a second. Now put money on it.

Nadia sells running shoes online from a unit outside Leicester, about ten thousand orders a month at an average of eighty pounds. The shoes cost her forty-eight pounds a pair, so she keeps thirty-two pounds of gross margin on a typical order.

About forty orders a month are placed with stolen card details. She ships them, and six to ten weeks later the real cardholder notices the charge and disputes it. Three things then go wrong at once. The eighty pounds is taken back out of her account. The shoes are gone. And her payment provider charges her an administration fee for handling the dispute, somewhere around fifteen pounds, set by her contract rather than by any regulator.

So one successful fraud costs Nadia forty-eight pounds of stock plus fifteen pounds of fee, which is sixty-three pounds. Forty of those a month is two thousand five hundred and twenty pounds. That is her fraud problem, whole and entire. One wrongly refused genuine customer costs her thirty-two pounds, the margin she did not make.

Those two numbers are the whole chapter. Sixty-three against thirty-two. Stopping one fraud is worth roughly two lost customers. Not two hundred. Two.

Nadia's first instinct is the one everybody has. Most of the fraudulent orders go to a delivery address that does not match the billing address on the card, so she writes a rule: if the two differ, refuse the order. It works, in the sense that it catches fraud. Thirty of the forty frauds are stopped, one thousand eight hundred and ninety pounds of loss prevented, and if that were the only number on the report she would think she had had a very good month.

But the rule does not only flag fraudsters. It flags the man buying trainers for his nephew's birthday, the student whose card is registered to her parents' house, everybody who has moved recently, everybody shipping to an office. Six hundred genuine orders a month trip the same rule. Six hundred times thirty-two pounds is nineteen thousand two hundred pounds of margin, destroyed silently, with no incident report anywhere.

Add it up. Before the rule, fraud cost her two thousand five hundred and twenty pounds a month. After it, she still loses six hundred and thirty pounds to the ten frauds that got through and has additionally thrown away nineteen thousand two hundred pounds of good business. The rule made her situation nearly eight times worse, and it felt the whole time as though it were working, because the only number anybody looked at was the fraud she stopped.

There are two words for what went wrong and a business reader should know both. **Precision** asks: of everything I stopped, how much was actually bad? Nadia's rule stopped six hundred and thirty orders and thirty were fraud. Precision of about five per cent, meaning nineteen out of every twenty people she refused had done nothing wrong. **Recall** asks: of all the bad things out there, how much did I stop? She stopped thirty of forty, so recall of seventy-five per cent. Her rule had good recall and catastrophic precision, and precision was where the money was.

Now watch what happens when she asks four questions at once. She keeps the address mismatch, but only acts on it when three other things are also true: the account was created less than ten minutes before the order; the same card number was submitted more than once in the session with different

expiry dates; and that delivery postcode has been used by three or more different accounts in the last thirty days.

Each is innocent alone. New accounts are made constantly, people mistype expiry dates, blocks of flats share postcodes. It is the coincidence that is not innocent. The combination fires on forty-five orders a month and twenty-eight are fraud.

Precision has gone from five per cent to sixty-two, recall has slipped from seventy-five to seventy. She now prevents one thousand seven hundred and sixty-four pounds of fraud while destroying five hundred and forty-four pounds of good business, so her total monthly cost is one thousand three hundred pounds against two thousand five hundred and twenty if she did nothing. She is making money on the arrangement for the first time, and she is also letting twelve frauds through every month on purpose.

That is not failure. That is the answer.

Because there is a way to get fraud to exactly zero and everybody knows what it is: refuse every order. Nadia would lose three hundred and twenty thousand pounds of margin a month and never be defrauded again. Zero fraud is trivially available and nobody wants it. So the real question was never how to stop fraud. It was where, on the dial between letting everything through and letting nothing through, the total cost bottoms out.

The dial has a setting, it falls out of those two numbers, and it is genuinely surprising. Nadia should refuse an order when she thinks there is more than about a **one in three** chance it is fraudulent. Not a nine in ten chance. Not “probably fraud”. Thirty-four per cent. At that point the sixty-three pounds she saves two-thirds of the time exactly balances the thirty-two pounds she throws away one-third of the time. Above it she should refuse; below it she should ship the shoes and accept the losses, because refusing costs more than being defrauded. Cheap shoes on a fat margin move the dial one way and expensive shoes on a thin margin move it the other. It is a business decision wearing a technical costume.

So what does the machine look at, if not the address?

Velocity. Not “how many”, but “how many, in how long, from the same thing”. One card tried once is a customer; one card tried nineteen times in four minutes across eleven accounts is not. One delivery address on two orders is a household; one delivery address on sixty orders in a week, all to different names, is a drop. Velocity is the oldest and cheapest signal there is and still one of the most useful, because criminals have to repeat themselves to make money and honest people mostly do not.

Device fingerprinting. When your browser talks to Nadia’s shop it says a great deal about itself without being asked: which fonts are installed, how big the screen is, which time zone, how the graphics chip draws a particular curve. Combine forty of those and you get something often unique. It is not you. It is the machine you are sitting at, which matters because a fraudster with a thousand stolen cards is often sitting at one machine.

Behavioural signals. How you type, not what you type. Whether you typed your own name fluently, the way somebody does who has written it ten thousand times, or slowly, the way somebody does who is copying it off a screen. Whether you were on a call the whole time you were paying, which is one of the strongest warnings in the business, because someone talking a stranger through a bank transfer is usually stealing from them.

The network. Nadia alone sees one strange order. Nadia plus four hundred other shops on the same platform sees the same device try eleven of them in ninety seconds, a pattern invisible to any one of

them. The biggest advantage a large processor has over a small one is not better mathematics. It is that it has seen the card before.

And finally the doorman's third answer. Some orders go into a pile for a person to look at. A trained reviewer costs an employer something like forty thousand pounds a year once you count national insurance, pension, a desk and a supervisor, about two pounds per order at a dozen an hour. Two pounds is not nothing when the order only makes thirty-two, so the rule for the review pile is simple and almost never followed: only put an order in front of a human if a human might change the answer.

Where the plain version stops being true

The doorman can see the person. The machine can only see what the person chooses to send it. Every signal above arrives over a wire from the party being judged. The device fingerprint, the browser headers, the claimed time zone, the typing rhythm: all of it is client-supplied, and a determined attacker can influence any of it. This is not a reason to abandon the signals, which remain useful because spoofing consistently is much harder than spoofing once. It is a reason to notice which signals do not come from the client at all. The issuer's view of the account, the network's view of the card across thousands of merchants, a cryptogram generated inside a chip, a passkey bound to hardware: these are expensive to forge precisely because the attacker is not producing them. The defence is layering: cheap client signals to sort the ordinary from the odd, expensive non-client evidence to decide the odd cases.

"Fraud" is at least four different problems and the arithmetic above fits only one. Nadia's case is third-party fraud: a criminal using a real cardholder's details. Beside it sit three others with different physics. First-party misuse, sometimes called friendly fraud, is where the genuine cardholder makes a genuine purchase and then disputes it, and no device signal will ever help because everything about the transaction is authentic. Account takeover is where the criminal is inside the real customer's account, so the history that normally exonerates now vouches for the attacker. And authorised push payment scams invert the problem: the victim makes the payment themselves, correctly authenticated, fully intending to, while being manipulated by someone on the phone. In the United Kingdom that fourth category is now comparable in scale to the first. UK Finance's Annual Fraud Report 2026, published on 15 June 2026 and covering calendar year 2025, reported unauthorised fraud losses of £703.4 million against authorised push payment losses of £576.4 million, the second figure rising nineteen per cent year on year. In APP fraud the model is not looking for a criminal in the session. It is looking for a victim. The signal set, the intervention and the legal consequences are all different.

The plain version put the whole loss on the merchant, and in the United Kingdom that is frequently wrong. Who pays depends on who authenticated, and it moved substantially with strong customer authentication. Under the Payment Services Regulations 2017, as they stood at the time of writing in August 2026, regulation 76 requires a payment service provider to refund an unauthorised transaction no later than the end of the business day following the day it becomes aware of it, unless it has reasonable grounds to suspect fraud by the customer and has notified the authorities under the tipping-off provisions of the Proceeds of Crime Act 2002. Regulation 77, again as at August 2026, caps the payer's own liability at a maximum of thirty-five pounds for losses arising from a lost, stolen or misappropriated payment instrument, and removes even that where strong customer authentication was required but the payer's provider did not apply it, or where the instrument was used in connection with a distance contract. So a false negative and a false positive have different owners at every point in the chain, and "the loss ratio" means one thing to a merchant, another to an acquirer and another to a card issuer. Any conversation about tuning that has not established whose money is at stake is a conversation about nothing.

A score is not a probability, and the threshold is not an engineering decision. Nadia's thirty-four per cent rule only works if the model's output is a genuine probability, meaning that among all orders scored at thirty-four per cent, roughly thirty-four per cent really do turn out to be fraud. Most models out of training do not have that property, and a raw score of eight hundred and fifty on a scale to a thousand means nothing until somebody has mapped it onto reality. Worse, the ground truth arrives late and incomplete: a dispute may not surface for months, and blocked orders never generate any outcome at all, so the system learns only from traffic it approved and becomes confident about a world it has stopped observing.

The technical version

■ On dating

Everything below is stated as accurate at the time of writing, in August 2026, and this is among the fastest-moving corners of the subject. Three things shifted in the eight months before it was written and are named in place: the FCA's removal of the prescribed contactless limits, effective March 2026; the replacement of Article 22 of the UK GDPR with new Articles 22A to 22D on 5 February 2026; and the continuing revision of Visa's acquirer monitoring thresholds. Scheme programmes are distributed to acquirers rather than published, so treat every scheme figure here as a prompt to check your acquirer's current bulletin.

■ Where the decision actually gets made

There is no single fraud check. There are five decision points in a card transaction, each with different data, a different latency budget and a different party carrying the loss.

Point	Who decides	Data available	Latency	Who bears the loss
Pre-authorisation screening	Merchant or gateway	Basket, history, device, session behaviour	Tens to hundreds of milliseconds	Merchant, in lost margin
Authentication	Issuer, via EMV 3-D Secure	Merchant-supplied 3DS data plus issuer's own view	Seconds, with a challenge	Shifts on successful authentication
Authorisation	Issuer	Card, account, balance, issuer models, network signals	Milliseconds in the network round trip	Issuer, or merchant if unauthenticated
Fulfilment hold	Merchant	Everything above, plus time and manual work	Minutes to hours	Merchant
Dispute	Issuer, then scheme arbitration	The full record, months later	Days	Per scheme rules and liability shift

EMVCo publishes the authentication specification, the EMV 3-D Secure Protocol and Core Functions Specification, which has run through version 2.3.1.1, covered by EMVCo Specification Bulletin 279 of 11 August 2025. The point of 3DS 2.x is that it carries merchant-supplied risk data into the issuer's decision, permitting a frictionless outcome in most cases rather than a password screen for everybody.

■ Rules, scorecards and learned models

Practitioners sometimes talk as though machine learning replaced rules. It did not, and a stack with no rules layer usually cannot implement policy. Rules carry what are not statistical judgements: sanctions and blocklist hits, hard regulatory constraints, contractual prohibitions, an instruction from an investigator to freeze a specific device. These must be deterministic, auditable and changeable at three in the morning by somebody who is not a data scientist. A rule engine edited under change control, with every version retained, is a compliance artefact as much as a detection tool.

Gradient-boosted decision tree ensembles trained on tabular features remain the workhorse for scoring the grey zone. They handle mixed data types and missing values, they are fast enough at inference, and their feature attributions are legible enough to survive a model validation review.

Sequence and graph models sit above both. A sequence model treats activity as an ordered series rather than a bag of counters, which lets it notice that a session is shaped wrongly rather than that a single field is wrong. A graph model works over the entity network, with cards, devices, addresses and beneficiaries as nodes and transactions as edges, and earns its keep on organised fraud and mule networks, where the individual transactions are unremarkable and only the topology gives it away.

They fail differently: rules by going stale, tree ensembles by drifting as the population moves, graph models by being expensive and lagging real time. Running all three, with the rules able to override the models, is defence in depth.

■ Velocity, entity resolution and the feature store

A velocity feature is a count or sum over a sliding window, keyed on an entity. The design questions are which entity, which window, which aggregate. Cards per device in the last hour. Distinct billing postcodes per card in the last day. Declined authorisations per merchant per BIN range in the last five minutes.

Two engineering problems dominate. The first is entity resolution: a counter is only as good as the identity it is keyed on, and identities are messy, since the same card is a different token at every merchant and the same device changes after an update. Real systems maintain an identity graph stitching these together with explicit confidence, and the quality of that graph sets a ceiling on everything downstream.

The second is point-in-time correctness. If a training feature is computed using data that would not have been available at the moment of the decision, the model learns from the future and looks superb in backtest and mediocre in production. This is leakage, and it is the commonest reason a fraud model performs worse live than on paper.

Card testing, also called enumeration, is the canonical velocity problem. From the defender's side it presents as a burst of low-value authorisation attempts with an abnormally high decline rate, often against a narrow range of card numbers, at merchants with a low-friction checkout. The controls exist because of it. Visa's Account Attack Intelligence produces a VAAI Score used to confirm enumerated authorisations, and Visa's Acquirer Monitoring Program carries a separate enumeration ratio. Mastercard's Security Rules and Procedures, Merchant Edition, in the edition dated 11 February 2025 which was current at the time of writing, provides at section 3.12.4 that in the case of a BIN attack at a card-not-present merchant or payment facilitator, the acquirer, its service providers or the affected merchant must begin to capture and transmit CVC 2 values in all card-not-present authorisation request messages within seventy-two hours of detection or notification.

■ Device and session signals, and their legal wrapper

Device identification in a browser draws on the properties a client exposes: user agent and header ordering, characteristics of the TLS handshake, screen and viewport metrics, font enumeration, canvas and WebGL rendering differences, declared time zone and language, and the coherence between them. On mobile the signal set is richer and more stable, because an app can read hardware and operating system attributes a browser cannot.

Two caveats belong here. First, the entropy available in browsers is deliberately shrinking as privacy changes across the major engines reduce the distinguishing power of these signals, so fingerprint-based identification is a depreciating asset. Second, this is regulated processing. In the United Kingdom, at the time of writing in August 2026, storing information on or gaining access to information stored on a user's terminal equipment engages regulation 6 of the Privacy and Electronic Communications (EC Directive) Regulations 2003, which requires consent subject to a narrow exemption where the access is strictly necessary to provide a service the subscriber has explicitly requested. Firms generally rely on that exemption for payment fraud checks, and on legitimate interests as the lawful basis under Article 6(1)(f) of the UK GDPR, supported by recital 47, which states expressly that processing strictly necessary for the purposes of preventing fraud constitutes a legitimate interest. Neither position survives inspection if it has not been documented, scoped and assessed.

Behavioural signals are a distinct family. Passive behavioural biometrics observe keystroke dynamics, pointer or touch trajectories, hesitation and correction patterns, and device motion from the accelerometer and gyroscope. The interpretive value is comparative: does this session look like this account's previous sessions, and like a person performing a familiar task or an unfamiliar one. In APP scam detection this family carries the load, because there is nothing wrong with the credentials or the device, and the signals that matter indicate coercion or coaching: a long dwell on the payee details, evidence of an active call during the session, hesitation inconsistent with the customer's baseline. UK Finance's 2026 report quoted BioCatch's global advisory director to this effect, that detection must move to identifying social engineering in real time before funds leave the account. One caution fraud teams underweight: behavioural biometrics can incidentally reveal information about a person's health, motor control or disability, and where processing reveals data concerning health it becomes special category data under Article 9 of the UK GDPR, which matters both for the lawful basis and for the automated decision rules below.

■ The network effect and the mechanics of sharing

The largest structural advantage in fraud detection is population. A model trained on one merchant's traffic sees that merchant's fraud. A model trained across a consortium sees the same device, card or beneficiary as it moves between merchants, which is where organised fraud is visible. Sharing runs through four mechanisms in the United Kingdom and they are not interchangeable.

Scheme-level merchant risk lists are the oldest. Mastercard operates MATCH, the Mastercard Alert to Control High-risk Merchants system, now MATCH Pro, into which acquirers must add terminated merchants under defined reason codes and against which they must query prospective ones; Visa operates an equivalent. These are compliance obligations under the schemes' security rules, not optional intelligence products, and being listed is close to fatal to a merchant's ability to obtain acquiring.

Industry fraud databases operate at consumer and account level. Cifas maintains the National Fraud Database, into which members file confirmed fraud markers and against which they screen, and UK Finance operates fraud intelligence sharing among its members. Pay.UK's Confirmation of Payee, mandated by the Payment Systems Regulator in phases, checks the name a payer types against the

name on the destination account before the payment is sent, a control aimed squarely at APP fraud and invoice redirection.

Statutory gateways sit underneath. The Economic Crime and Corporate Transparency Act 2023 created a direct disclosure gateway allowing certain businesses to share customer information to prevent economic crime, with protection from civil liability for breach of confidence where the conditions are met. That addresses the confidentiality problem, not the data protection one, which still requires a lawful basis, purpose limitation analysis, transparency and an impact assessment. Commercial consortium models are what most merchants buy, and the question usually left unasked in procurement is what happens to the merchant's own data on exit.

■ Precision and recall, properly

The confusion matrix has four cells and every one has a price.

	Model says fraud	Model says legitimate
Actually fraud	True positive: loss avoided	False negative: chargeback, goods gone, fee
Actually legitimate	False positive: lost margin and customer	True negative: normal business

Precision is true positives divided by everything flagged. Recall, also called the detection rate, is true positives divided by all actual fraud. Specificity, the true negative rate, is rarely quoted and is where the money hides.

The reason is the base rate, and it deserves a worked example because business readers are routinely misled by accuracy figures. Take ten thousand orders with forty frauds, a fraud rate of 0.4 per cent, and a detector with ninety per cent recall and a one per cent false positive rate. It flags thirty-six real frauds and one hundred false alarms, giving precision of twenty-six per cent: three out of four people you refuse have done nothing wrong. Now improve only the false positive rate, from one per cent to 0.2 per cent, changing nothing about recall. The flagged set becomes thirty-six real frauds and twenty false alarms, and precision jumps to sixty-four per cent.

Nothing about the model's ability to detect fraud changed. All the leverage was on the negative class, because the negative class is two hundred and fifty times bigger. At low base rates, small movements in the false positive rate dominate everything else, and any metric averaging over both classes will hide it. That is also why area under the receiver operating characteristic curve is a poor headline metric for fraud, since it is computed across the full range of false positive rates, almost all of which are commercially unusable. Precision-recall curves, or simply recall at a fixed operating point such as detection rate at a one per cent flag rate, tell you what you need to know.

Two refinements separate professional measurement from amateur. Count-weighted and value-weighted metrics diverge, so a system catching ninety per cent of fraud attempts by count while missing the three largest by value has failed; value detection rate belongs beside the count figures on every report. And the labels are wrong. Chargebacks arrive with a lag of weeks to months, and blocked transactions produce no outcome at all, so the false positive rate cannot be measured without deliberately letting some flagged traffic through as a holdout. A team that has never run one does not know its own false positive rate and should not claim to.

■ Tuning to a loss ratio

The industry expresses fraud performance in basis points of processed volume: fraud losses divided by transaction value, times ten thousand. The operating discipline hangs off that ratio rather than off an absolute number of incidents.

The context for what counts as a normal ratio comes from the regulators. The joint EBA and ECB Report on Payment Fraud published on 1 August 2024, the first in that series, found EEA payment fraud of €4.3 billion in 2022 and €2.0 billion in the first half of 2023, with card fraud on EEA-issued cards running at 0.031 per cent of total value and 0.015 per cent of total number in the first half of 2023. It also found fraud rates on card payments around ten times higher where the counterparty was outside the EEA, where strong customer authentication is not legally required, and that seventy-one per cent of card fraud by value in that period was cross-border. Later editions report the same direction with a wider extra-EEA gap.

That ratio is not merely a management preference. It is a regulatory threshold in two directions. The first is the transaction risk analysis exemption. Under Article 18 of the UK version of the strong customer authentication regulatory technical standards, as they stood at the time of writing in August 2026 and as reflected in the FCA's Technical Standards chapter last updated on 19 March 2026, a payment service provider may decline to apply strong customer authentication to a remote payment it has identified as low risk, provided its fraud rate for that transaction type is at or below the reference fraud rate in the Appendix, the value is within the corresponding exemption threshold value, and a real-time risk analysis has not identified abnormal spending or behaviour, unusual device or software information, malware in the authentication session, a known fraud scenario, an abnormal payer location or a high-risk payee location. That clause turns a fraud team's performance into a commercial asset: below the reference rate you may waive authentication on low-risk payments, which materially improves conversion; above it you may not.

The reference fraud rates in the equivalent EU instrument, Commission Delegated Regulation (EU) 2018/389, remain as follows at the time of writing, with the UK Appendix carrying sterling exemption threshold values rather than euro ones.

Exemption threshold value	Reference fraud rate, remote card	Reference fraud rate, remote credit transfer
500	0.01%	0.005%
250	0.06%	0.01%
100	0.13%	0.015%

Article 19, at the time of writing, defines the calculation precisely: the total value of unauthorised or fraudulent remote transactions, whether or not funds were recovered, divided by the total value of all remote transactions of that type, whether authenticated or executed under any exemption, on a rolling ninety-day basis, subject to the audit review required by Article 3(2). Article 20 sets the consequence: a provider must report immediately to the FCA when a monitored fraud rate exceeds the applicable reference rate, must cease using the exemption for that transaction type and threshold band if the rate is exceeded for two consecutive quarters, and may not resume until the rate has been at or below the reference rate for one quarter and it has notified the FCA with evidence. Article 21 requires quarterly monitoring of fraud values, fraud rates, average transaction values and exemption usage.

Two other UK figures under the same instrument are worth stating with their dates. The low-value remote transaction exemption in Article 16, at the time of writing in August 2026, applies where the

individual transaction does not exceed £25 and either the cumulative value since the last strong customer authentication does not exceed £85 or there have been no more than five consecutive such transactions. And the contactless exemption in Article 11 no longer carries a prescribed limit at all: the FCA announced on 19 December 2025, in a press release accompanying Handbook Notice 136, that it was removing the fixed contactless limits and allowing providers with strong fraud controls to set their own, with the rule changes taking effect in March 2026. The Article 11 text as it stood in August 2026 simply permits the exemption where the provider identifies the transaction as posing a low level of risk. The FCA expected most providers to keep their existing limits, so a firm's operational limit and the regulatory limit are now separate facts to be checked separately. UK Finance reported contactless fraud losses of £46.8 million in 2025, up eight per cent.

The second regulatory direction is the schemes' own monitoring, which binds merchants and acquirers rather than issuers.

Visa's Acquirer Monitoring Program consolidated the older fraud and dispute programmes into a single count-based ratio. The numerator combines card-not-present fraud reports, filed as TC40 records, with non-fraud disputes filed as TC15 records; the denominator is settled card-not-present transactions. A separate enumeration ratio measures enumerated authorisation attempts, approved and declined, against the total count of authorisation attempts, approved and declined. The programme identifies both acquirer portfolios and individual merchants, applies minimum monthly counts before a party enters it, and assesses fees per fraudulent or disputed transaction levied on the acquirer, which passes them on by contract. The thresholds themselves are published by Visa rather than confined to acquirer bulletins: its fact sheet "Visa Acquirer Monitoring Program Overview", on the corporate Visa site, puts an acquirer portfolio above standard at a VAMP ratio of fifty basis points and excessive at seventy, sets the excessive merchant threshold in Asia-Pacific, Canada, the European Union and the United States at two hundred and twenty basis points falling to one hundred and fifty on 1 April 2026, requires fifteen hundred combined fraud and dispute events in a month before a party is identified at all, and puts the enumeration criteria at a ratio of two thousand basis points on a minimum of three hundred thousand enumerated authorisations. Those figures have been revised more than once since the programme launched and are reported inconsistently by third parties, so check them against Visa's current fact sheet and against your acquirer's own bulletin before relying on any of them.

Mastercard operates an Excessive Chargeback Program and an Excessive Fraud Merchant programme of similar structure, keyed on fraud chargeback basis points, absolute fraud value and, distinctively, the merchant's take-up of EMV 3-D Secure. Separately and verifiably, the Mastercard Security Rules and Procedures, Merchant Edition, in the 11 February 2025 edition current at the time of writing, provides at section 3.12.4 that an acquirer must ensure that each merchant which has exceeded one hundred gross basis points in fraudulent card-not-present transactions for two consecutive calendar months captures and transmits the CVC 2 value for all mail order and telephone order transactions, and for e-commerce either captures and transmits CVC 2 or becomes Mastercard Identity Check enabled, within two months following the second trigger month.

So there are two floors under the loss ratio. There is the commercial floor, where fraud losses plus lost good business plus operating cost are minimised, and a regulatory or scheme floor, above which you lose exemptions, lose the frictionless authentication path and start paying per-transaction penalties. The second is often higher, which means a merchant optimising purely on expected value can find itself out of compliance while making money.

The expected-value calculation is the one from the plain version, stated properly. Let the calibrated probability of fraud be p , the loss given fraud be L , and the contribution margin forgone on a wrongly

declined order be M . Decline when pL exceeds $(1 - p)M$, that is when p exceeds M divided by $(M + L)$. With the shoe retailer's L of £63 and M of £32 the optimal threshold is 32 divided by 95, approximately 33.7 per cent. Three consequences follow and all three are routinely missed. The threshold moves with the basket, not just the customer: a high-margin, low-cost digital good has a small L and a large M , so the threshold rises. Calibration is not optional, because if the score is not a probability the threshold is meaningless, and calibration decays faster than discrimination. And L is not just the goods; it includes the dispute fee, the operational cost of the case, the effect on the scheme monitoring ratio and any downstream penalty. Firms modelling L as the transaction value alone systematically under-block.

■ Manual review queues and their economics

A review queue is a third decision bucket and should be sized by economics rather than headcount, which is the reverse of how it usually happens. A reviewer on a thirty-thousand-pound salary costs an employer roughly forty thousand fully loaded once employer national insurance, pension, workspace, tooling and supervision are counted. Against roughly seventeen hundred productive hours a year that is about twenty-four pounds an hour, and at twelve cases an hour the marginal cost of a review is about two pounds.

Now test whether the queue earns it. Take three hundred orders a month in the grey band, of which twenty are fraud and two hundred and eighty genuine, on the shoe retailer's numbers. Auto-approving the band costs twenty times sixty-three pounds, or £1,260. Auto-declining it saves that but destroys two hundred and eighty times thirty-two pounds, a net loss of £7,700, so declining is clearly wrong. Reviewing all three hundred at a reviewer accuracy of ninety per cent on fraud and ninety-seven per cent on genuine orders costs six hundred pounds of review time, plus two missed frauds at £126, plus about eight wrongly declined genuine orders at £269, totalling roughly £995. The queue saves about two hundred and sixty-five pounds a month. It is worth running, barely. Raise the average order value to four hundred pounds on the same margin percentage and the same queue saves thousands, which is why review is standard in travel, jewellery and electronics and marginal in low-value retail: the economics are driven overwhelmingly by order value.

Four operating principles follow. Queue only where a human has information the model lacks: a reviewer seeing the same features on the same screen is an expensive, less consistent copy of the model, whereas one who can make an outbound call or verify a document earns their cost. Measure reviewer precision and recall as you would a model's, by seeding the queue with blind duplicates of resolved cases; reviewers drift, tire and develop idiosyncratic heuristics, and on cases where the model was already confident they are frequently worse than it is. Treat latency as a cost, because conversion decays measurably with delay. And do not let capacity set the threshold: the commonest failure in real operations is a threshold chosen so the queue is exactly as long as the team can clear, which means the risk appetite of the business is set by its recruitment budget.

■ The legal position on automated decisions

Fraud decisions refuse people service, sometimes close accounts, and sometimes place a marker other institutions will see. That engages data protection law directly, and the UK position changed recently enough that many internal policies had not caught up at the time of writing.

Section 80 of the Data (Use and Access) Act 2025 replaced Article 22 of the UK GDPR with new Articles 22A to 22D, brought into force on 5 February 2026 by the Data (Use and Access) Act 2025 (Commencement No. 6 and Transitional and Saving Provisions) Regulations 2026. The old general prohibition on solely automated decisions with legal or similarly significant effects has become a safeguards regime. Under Article 22A a decision is based solely on automated processing where there is no meaningful human involvement, and is significant where it produces a legal effect or a similarly significant one.

Under Article 22B a significant decision based entirely or partly on special category data may not be taken solely automatically unless explicit consent applies, or the decision is necessary for a contract or required or authorised by law and the substantial public interest condition in Article 9(2)(g) is met. Under Article 22C, where a significant decision is taken solely automatically, the controller must have safeguards giving the individual information about the decision and enabling them to make representations, to obtain human intervention and to contest it.

For a fraud team this converts into four obligations: know which decisions are significant, know which are solely automated, demonstrate that any human in the loop is meaningfully involved rather than clicking through, and have a working route for a customer to contest an outcome. A rubber-stamp review queue is not a compliance control; after these provisions it is arguably worse than no queue at all, because it asserts a human involvement that does not exist.

Two adjacent points. In the European Union the AI Act classifies AI systems used to evaluate the creditworthiness of natural persons as high risk under Annex III, but with an express exception: systems used for the purpose of detecting financial fraud are carved out. The carve-out attaches to the fraud detection purpose, not to any model a fraud team happens to run. And in the United Kingdom, banks, building societies and PRA-designated investment firms holding internal model approval for regulatory capital have been subject to the PRA's Supervisory Statement SS1/23 on model risk management principles since it took effect on 17 May 2024, whose five principles on identification and tiering, governance, development and use, independent validation and mitigants apply across model types rather than only to capital models. Firms outside that scope adopt the framework voluntarily, because the alternative when a fraud model misbehaves is having no answer to who validated it.

■ What breaks

Fraud models degrade for reasons ordinary models do not. The data-generating process is an adversary who observes your decisions and adapts, so performance decay is not drift but response. Champion-challenger deployment, where a shadow model scores the same traffic without acting on it, is the minimum; holdout groups are the only way to measure true precision, and their cost is the price of knowing.

Two failure modes are structural rather than technical. The first is the feedback loop: a model trained on approved traffic loses sight of the population it declines, grows more confident about a shrinking world, and drifts without any metric moving. The second is metric substitution: a team measured on fraud losses will reduce fraud losses, and the only thing stopping it destroying the business is that somebody else is measured on approval rate. Fraud detection is one of the few disciplines where the correct governance structure is two people with opposed incentives and a shared number.

It is worth closing on what all this exists to change. The Payment Systems Regulator's dashboard updated on 30 July 2026 recorded that in the eighteen months from 7 October 2024 to 31 March 2026, eighty-eight per cent of in-scope APP scam losses, some £316 million, were reimbursed to victims, with eighty-two per cent of claims closed within five business days and the maximum level of reimbursement set at £85,000 per claim under Specific Direction 20, following the PSR's Policy Statement PS24/7. Those figures are current at the time of writing in August 2026 and will be superseded. The structure they describe, in which detection quality decides both how much is stolen and how much a firm must hand back, will not be.

51.98 Common wrong ideas

Wrong: The goal is to stop fraud. **Right:** Zero fraud is trivially available by refusing every order, and nobody wants it; the goal is the setting of the dial at which fraud losses plus destroyed good business plus operating cost are lowest.

Wrong: A rule that catches most of the fraud is a good rule. **Right:** Nadia's address-mismatch rule had seventy-five per cent recall and about five per cent precision, and it destroyed nineteen thousand two hundred pounds of margin to prevent one thousand eight hundred and ninety pounds of loss.

Wrong: You should decline an order when it is probably fraudulent. **Right:** Decline when the calibrated probability exceeds M divided by $(M + L)$, about thirty-four per cent on the shoe retailer's numbers, because refusing a good order costs money too.

Wrong: The loss given fraud is the value of the transaction. **Right:** It is the stock, the dispute fee, the operational cost of the case, the effect on the scheme monitoring ratio and any downstream penalty, and firms modelling it as transaction value alone systematically under-block.

Wrong: A score of eight hundred and fifty out of a thousand means the order is very likely fraud. **Right:** A score is not a probability until somebody has mapped it onto outcomes, and calibration decays faster than discrimination.

Wrong: Machine learning replaced rules. **Right:** Rules carry what are not statistical judgements, sanctions and blacklist hits, hard regulatory constraints, an investigator's freeze, and they must stay deterministic, auditable and changeable at three in the morning by somebody who is not a data scientist.

Wrong: Device fingerprinting identifies the person. **Right:** It identifies the machine, its available entropy is deliberately shrinking as browser privacy changes land, and in the United Kingdom it engages regulation 6 of the Privacy and Electronic Communications Regulations and needs a documented lawful basis.

Wrong: The merchant bears the fraud loss. **Right:** Who pays depends on who authenticated, and regulations 76 and 77 of the Payment Services Regulations 2017 place much of it on the payer's provider, so any tuning conversation that has not established whose money is at stake is a conversation about nothing.

Wrong: You know your false positive rate. **Right:** Blocked orders generate no outcome at all, so it cannot be measured without deliberately letting some flagged traffic through as a holdout, and a team that has never run one does not know its own rate.

Wrong: A manual review queue is a safety net worth having. **Right:** It earns its cost only where the human holds information the model lacks, and since Articles 22A to 22D took effect a rubber-stamp queue is arguably worse than none, because it asserts a human involvement that does not exist.

51.99 Chapter summary in 20 lines

1. A fraud system is a doorman with three answers, wave in, turn away, or hold for a look, running thousands of times a second with money on it.
2. It makes two kinds of mistake, and only the fraud it lets through leaves a trace; the good customer it refuses simply goes elsewhere and says nothing.
3. For Nadia the two numbers are sixty-three pounds lost to a successful fraud and thirty-two pounds of margin lost to a wrongly refused order.

4. Stopping one fraud is therefore worth about two lost customers, not two hundred, and every design decision in the chapter follows from that ratio.
5. Her address-mismatch rule caught three-quarters of the fraud and left her nearly eight times worse off, because nineteen of every twenty orders it stopped were innocent.
6. Precision asks how much of what you stopped was actually bad, recall asks how much of the bad you stopped, and the money was in precision.
7. Combining four signals that are each innocent alone lifted precision from five per cent to sixty-two and turned a loss into a gain, while deliberately letting twelve frauds a month through.
8. Zero fraud is available to anyone willing to refuse every order, so the real question was always where the total cost bottoms out.
9. The optimal decline threshold is M divided by $(M + L)$, about thirty-four per cent here, and it moves with the basket rather than only with the customer.
10. The workhorse signals are velocity, device fingerprinting, behavioural patterns and the network, and the largest structural advantage in the field is simply having seen the card before.
11. Every client-supplied signal can be influenced by the party being judged, so the defence is layering cheap client signals against expensive evidence the attacker does not produce.
12. Fraud is at least four problems, third-party fraud, first-party misuse, account takeover and authorised push payment scams, the last now comparable in United Kingdom scale to the first.
13. There are five decision points in a card transaction, each with different data, latency and loss-bearer, and United Kingdom law decides where the loss lands at each.
14. Rules, tree ensembles and graph models fail differently, so running all three with the rules able to override the models is defence in depth.
15. Velocity features are only as good as the entity resolution beneath them, and point-in-time correctness is what separates a model that works live from one that merely backtests well.
16. At a fraud rate of 0.4 per cent, moving the false positive rate from one per cent to 0.2 per cent lifts precision from twenty-six to sixty-four while changing nothing about detection.
17. Labels arrive late and incomplete, so value-weighted metrics belong beside count-weighted ones and a holdout is the only honest measure of false positives.
18. The loss ratio has two floors, the commercial one where total cost is minimised and a regulatory and scheme one below which exemptions, frictionless authentication and freedom from per-transaction penalties are retained.
19. A review queue is a third bucket to be sized by arithmetic, and the commonest operational failure is letting the recruitment budget set the risk appetite.
20. Fraud models decay because the data-generating process is an adversary that watches your decisions, and the correct governance is two people with opposed incentives and a shared number.

Sources, all consulted in August 2026: the FCA Handbook Technical Standards on Strong Customer Authentication, Chapter 3, as updated 19 March 2026; the FCA press release of 19 December 2025 on contactless limits and Handbook Notice 136; Commission Delegated Regulation (EU) 2018/389; the Payment Services Regulations 2017 and the Data (Use and Access) Act 2025 section 80 with SI 2026/82; the PSR's APP scams reimbursement dashboard of 30 July 2026, PS24/7 and Specific Direction 20; UK Finance's Annual Fraud Report 2026; the joint EBA and ECB Report on Payment Fraud of 1 August 2024; Mastercard Security Rules and Procedures, Merchant Edition, 11 February 2025; EMVCo Specification Bulletin 279; PRA Supervisory Statement SS1/23; and the EU AI Act Annex III.

Anti-Money-Laundering

52.0 What this chapter gives you

1. You will be able to explain why what gets a criminal caught is almost never the money itself but the story attached to it.
2. You will be able to work through how a cash business launders by over-reporting its takings, and say why paying income tax on money nobody earned is a cost of laundering rather than a barrier to it.
3. You will be able to say what job each onboarding question does, so that an eleven-day account opening reads as a set of purposes rather than as obstruction.
4. You will be able to keep sanctions and money laundering apart in your head: suspicion and reasonableness on one side, absolute prohibition and strict liability on the other.
5. You will be able to explain why a screening run producing 120 possible matches and two real ones is working correctly, and why better software does not fix it.
6. You will be able to say why the risk-based approach governs due diligence but never sanctions screening.
7. You will be able to explain the “more than 25 per cent” beneficial ownership threshold and at least three structures that defeat it while leaving control untouched.
8. You will be able to describe what happens to a customer’s money after a refused Defence Against Money Laundering request, and how long it can lawfully stay immobilised.
9. You will be able to explain why a bank exits a whole category of customers whose individual members are fine, and why that makes the underlying problem worse rather than better.
10. You will be able to read the SARs and enforcement figures and say what the regime actually is: a privately funded intelligence-gathering exercise measured on process rather than on outcomes.

The plain version

Take a fifty pound note out of your pocket and look at it. There is nothing on it that says where it came from. It does not know whether it was handed over for a week’s wages, given at a birthday, or taken out of a till at knifepoint. A note is a note.

Hold on to that, because everything else follows from it. What gets a criminal caught is almost never the money itself. It is the story attached to the money.

If you turn up at a car showroom with £30,000 in a sports bag, nobody has to test the notes. They just have to ask one question: where did that come from? And you have to have an answer that survives being asked twice. Money laundering, stripped of every long word ever attached to it, is the manufacture of believable answers to that question. Anti-money-laundering is the business of asking it properly.

Here is how the manufacturing works, with real numbers, because the arithmetic is the trick. Imagine a hand car wash on a busy road: five people with sponges and a jet hose, £8 a wash, open seven days. Genuinely, they wash about ninety cars a week. That is £720 of honest takings, nearly all in cash, which is completely ordinary for a car wash.

Now suppose the owner also has £1,680 a week in cash from something illegal. They do not need to hide it. They need to *explain* it. So each week they write down not ninety washes but three hundred, and bank £2,400 instead of £720. The bank sees a car wash on a busy road banking £2,400 a week. That is believable. Nobody counts cars.

Over a year that is £87,360 of criminal cash with a birth certificate. It has been taxed. It appears in accounts. It can buy a flat. And this is the part people find hardest to accept: the owner will happily pay income tax on money they never earned, because a tax bill is the receipt that proves the story. Paying tax is a cost of laundering, not a barrier to it.

People who do this professionally break it into three moves. First, get the cash into the system at all, the risky part, because cash is heavy, conspicuous, and the one moment where a person and a pile of money are in the same room. Second, move it around enough that the trail becomes hard to follow: car wash to supplier, supplier to a company abroad, company to a loan, loan to a property. Third, spend it on something ordinary, so from then on the money looks like the proceeds of a house sale rather than of a crime.

Every anti-money-laundering rule you have ever been annoyed by exists to interfere with one of those three moves. Which brings us to why opening a bank account takes so long.

Anaya is starting a small import business. The bank asks her for what feels like an absurd amount. Passport. Proof of address. What the business actually does. Expected turnover. Which countries she will trade with. Who else owns shares. Where her opening deposit came from. It takes eleven days, and twice she is asked for something she has already sent.

None of that is the bank being difficult for sport. Each question does a job. Who are you is the easy one, and the least important. What does the business do gives the bank a shape to compare against later: an importer of textiles moving £40,000 a month is unremarkable, and the identical account moving £40,000 a month with no goods attached is not. Who else owns it matters because the person who signs the forms is often not the person who benefits, and the whole point of a shell company is that it puts a stranger's name on the door. And where the money comes from matters because a plausible source is exactly what a launderer cannot produce without lying, and a lie written down is evidence.

Then the bank runs her name against a list. Not a list of criminals: a list of people and organisations governments have ordered everyone to freeze out entirely. This is a different thing from money laundering and worth separating in your head now. Money-laundering rules say be careful and tell someone if it looks wrong. Sanctions say do not, under any circumstances, move this person's money, and there is no version of that where being careful is good enough.

The list check is where the system's most human failure lives, and the arithmetic is unkind. Say the bank takes on 4,000 new customers a week. Its screening software does not look for exact matches, because anyone wanting to evade an exact-match check could do it by dropping a letter. So it looks for near matches: same name spelled differently, name written back to front, a name that sounds the same aloud, a name transliterated from another alphabet where there was never one correct English spelling to begin with. Set the software to be sensitive and it catches more; set it to be strict and it misses things.

Set sensibly, it produces roughly 120 possible matches out of those 4,000. Perhaps two are real. The other 118 are people who share a name, or half a name, with somebody the software has heard of.

That is not a bug you can fix by buying better software. It is a property of names. Hundreds of millions of people share a few thousand very common names, and the software is deliberately built to over-report because the cost of missing one is enormous and the cost of stopping an innocent person for ten minutes is small. So a person with a common name will be stopped, and stopped again at the next bank, and will reasonably conclude the system has something against them personally. It does not. It cannot tell them apart from someone else, and it is not allowed to guess. At eight minutes each, those 118 wrong matches cost about sixteen hours a week to find two real ones. Multiply that across every bank in the country and you have an industry.

Once Anaya's account is open, the watching starts. Software compares what she actually does against what she said she would do, and against what businesses like hers normally do. A hundred thousand customers might generate two thousand alerts a month. The great majority are nothing: a bonus, an inheritance, a wedding, a house sale, a business that grew. Perhaps forty end up reported to the authorities. And when a bank does decide something looks wrong, it must not tell the customer, because warning someone that they have been reported is itself a criminal offence in the United Kingdom. So the customer gets a delay with no explanation, from staff legally forbidden to explain it. Almost every furious story you have heard about a bank refusing to say why has this at the bottom of it.

The last idea, and the one that governs the rest, is that you cannot do all of this equally to everybody. There is not enough time in the world. So you spend effort where the risk is: light checks on a teenager opening a savings account, heavy checks on a company with owners in four countries and a bank account in a fifth. That is the risk-based approach, and it sounds like common sense because it is. It is also the hardest thing in this chapter to do honestly, because it requires a firm to write down, in advance, which of its own customers it thinks are dangerous.

Where the plain version stops being true

The system is not primarily a net for catching launderers, and its own numbers say so. The car wash story implies that reports lead to arrests. Some do. But the regime measures firms on process, not outcomes: a firm complies by having systems, applying them and reporting suspicion, and it can do all three impeccably while the money still gets away. The United Nations Office on Drugs and Crime has long estimated that 2 to 5 per cent of global GDP is laundered annually, roughly \$800 billion to \$2 trillion, with its 2011 study putting the 2009 figure at around \$1.6 trillion, or 2.7 per cent. Against that, the United Kingdom Financial Intelligence Unit reported in its SARs Annual Report 2025, covering April 2024 to March 2025, that £382.6 million was denied to suspected criminals through Defence Against Money Laundering requests, a record and a 59 per cent rise on the previous year's £240.1 million. In the United States, the Bank Policy Institute's 2018 study *Getting to Effectiveness* found that surveyed institutions reviewed approximately 16 million alerts in 2017 and filed over 640,000 suspicious activity reports and more than 5.2 million currency transaction reports, and that among institutions recording the data, a median of 4 per cent of those reports drew any follow-up inquiry from law enforcement. Whatever the regime is, it is not a filter with a high hit rate. It is a very large intelligence-gathering exercise, paid for privately, from which a comparatively small number of interventions are extracted.

Identity is the easy half; ownership is the hard half, and the analogy skips it entirely. Verifying that Anaya is Anaya is close to a solved problem. Establishing who actually benefits from a company is not. At the time of writing in August 2026, regulation 5 of the Money Laundering, Terrorist Financing and Transfer of Funds (Information on the Payer) Regulations 2017 defines the beneficial owner of a body

corporate as an individual who exercises ultimate control over its management, or who ultimately owns or controls, directly or indirectly and including through bearer share holdings or by other means, more than 25 per cent of its shares or voting rights, or who otherwise controls it. That “more than 25 per cent” is a bright line, and bright lines are designed around. Four people at 25 per cent each disclose nobody. A chain of five companies across three jurisdictions dilutes an owner below the threshold at every layer while leaving control untouched. Nominee shareholders, discretionary trusts and companies limited by guarantee each break the arithmetic differently. This is why practitioners say ultimate beneficial ownership rather than ownership, and why the honest answer to who owns this company is often a judgement rather than a fact.

“False positive” is the wrong frame for screening, and screening is not part of the risk-based approach at all. A wrong name match is not an error. It is the intended output of a system deliberately tuned to over-report, because the two kinds of mistake are not symmetrical: missing a genuine sanctions target is a strict-liability breach, while stopping an innocent customer is an inconvenience. Every screening system therefore sits on a dial between precision and recall, and the dial is set where the penalties push it. More importantly, the risk-based approach does not apply here. A firm may decide a low-risk customer needs less due diligence. No firm may decide a low-risk customer needs less sanctions screening, because an asset freeze is an absolute prohibition. At the time of writing, section 146 of the Policing and Crime Act 2017, as amended by the Economic Crime (Transparency and Enforcement) Act 2022, allows the Office of Financial Sanctions Implementation to impose a civil monetary penalty where it is satisfied on the balance of probabilities that a breach occurred, and for breaches on or after 15 June 2022 any requirement that the person knew, suspected or believed anything is disregarded. Intention is irrelevant. That is a different legal world from money laundering, where suspicion and reasonableness are everything.

De-risking is not a failure of the risk-based approach; it is a rational response to how the penalties are shaped. A bank does weigh a customer’s risk against their value, but the two sides of that scale are measured in different units. The revenue from a small money-transfer business serving one remittance corridor is bounded and known. The downside is an enforcement action, a skilled person’s report, a criminal investigation and a number in the press: the Financial Conduct Authority stated when fining Nationwide Building Society £44,078,500 in December 2025 that since 2021 it had imposed thirteen fines totalling £300,767,526 on banks for anti-money-laundering systems and controls failings. Against an unbounded and asymmetric downside, exiting an entire customer category is the cheapest available answer even when the individual customers are fine. The consequence lands hardest on those with least ability to complain, and it is a policy problem rather than a compliance one.

The technical version

■ On dating

Every figure, threshold, list and rule below is stated as accurate at the time of writing, in August 2026, and this is the fastest-moving corner of regulation covered in this book. Three things were in motion as it was written: the United Kingdom’s Money Laundering and Terrorist Financing (Amendment) Regulations 2026 came into force on 30 June 2026 with two later commencement dates still ahead of it, the European Union’s single rulebook does not apply until 10 July 2027, and OFSI had announced an intention to double its maximum civil penalty without the statutory cap yet changing. Check the primary sources before relying on any number here.

■ The offences that drive the machine

In the United Kingdom the regulatory architecture rests on criminal law, and the criminal law came first. Part 7 of the Proceeds of Crime Act 2002 creates the principal offences: concealing, disguising, converting, transferring or removing criminal property under section 327; entering into or becoming concerned in an arrangement facilitating another's acquisition, retention, use or control of criminal property under section 328; and acquiring, using or possessing criminal property under section 329. Criminal property is defined by reference to a person's benefit from criminal conduct, with no minimum amount and no separate list of predicate offences: if the underlying conduct would be an offence in the United Kingdom, its proceeds are criminal property, however small.

Two further offences shape day-to-day behaviour more than the principal ones. Sections 330 and 331 create the offence of failing to disclose knowledge or suspicion arising in the course of regulated-sector business, with equivalents at sections 19 and 21A of the Terrorism Act 2000. Section 333A creates the tipping-off offence and section 342 the offence of prejudicing an investigation, with equivalents at sections 21D and 39 of that Act. The National Crime Agency states, at the time of writing, that the penalties on conviction on indictment for failing to disclose are up to five years' imprisonment, a fine, or both. This is why "we are unable to discuss the reason for this decision" is a legal formula rather than rudeness.

The defence matters as much as the offence. A person who would otherwise commit a principal offence can obtain a defence by making an authorised disclosure and obtaining appropriate consent: in practice, the Defence Against Money Laundering request, or DAML. Under section 335 the National Crime Agency has seven working days beginning with the first working day after disclosure in which to refuse consent. If it refuses, a moratorium period of 31 days runs, during which the prohibited act must not be done. Section 336A allows a court, on application by a senior officer, to extend that moratorium by up to 31 days at a time where an investigation is being conducted diligently and expeditiously, but subsection (7) caps further extensions so the total cannot exceed 186 days beginning with the day after the end of the initial 31-day period. A customer's funds can therefore be lawfully immobilised for over half a year on suspicion, with no obligation on anyone to explain why.

Section 339A sets a threshold below which deposit-taking bodies, electronic money institutions and payment institutions may operate an account without seeking a DAML. That threshold was raised from £250 to £1,000 by legislation introduced in January 2023, and the UKFIU attributes the subsequent flattening of DAML volumes partly to that change.

■ The regulations that make it a system

The criminal law tells a firm what it must not do. The Money Laundering, Terrorist Financing and Transfer of Funds (Information on the Payer) Regulations 2017, now maintained under sections 49 and 54(2) of, and Schedule 2 to, the Sanctions and Anti-Money Laundering Act 2018, tell it what it must build.

At the time of writing, the most recent substantial change is the Money Laundering and Terrorist Financing (Amendment) Regulations 2026, S.I. 2026/621, made on 9 June 2026, most of it in force 21 days later on 30 June 2026. Two parts commence later: regulation 20, inserting a new regulation 34A on enhanced due diligence for cryptoasset exchange providers and custodian wallet providers in correspondent relationships, on 1 February 2027; and regulation 37, substituting Schedule 6B on changes in control of registered cryptoasset businesses, for most purposes on 25 October 2027. The instrument implements the Government's response to HM Treasury's 2024 consultation *Improving the effectiveness of the Money Laundering Regulations*, published on 17 July 2025, and its direction is broadly deregulatory.

Its most visible change is arithmetical: every euro-denominated threshold on the face of the Regulations was converted to sterling, and not always at par.

Provision	Subject	Before	After
Reg. 27(1)(b)	Occasional transaction, currency exchange or money transmission	€1,000	£800
Reg. 27(2)	Occasional transaction, general threshold	€15,000	£12,000
Reg. 27(3)	Occasional transaction, high value dealers and others	€10,000	£10,000
Reg. 27(5), (7)	Casino and related transactions	€2,000	£2,000
Reg. 14	High value dealers, casinos, auction platforms, art market participants	€10,000	£10,000
Reg. 38(1), (2)	Electronic money, simplified due diligence and redemption limits	€150, €50	£150, £50
Regs. 64C(4), 64G(1)(b)	Transfers of funds, information requirements	€1,000	£800

Three other changes matter more than they look. The phrase “complex or unusually large” in regulations 19 and 33 became “unusually complex or unusually large in each case given the nature of the transaction”, a deliberate narrowing: a transaction is no longer high risk merely because it is structurally complicated. A new regulation 30ZA lets a credit institution open an account for a customer of a failed bank, and allow transactions from it, before completing due diligence, provided the customer is identified under regulation 28(2)(a) and the relationship begins within 30 days of the bank insolvency order. And a new block at regulation 29(10) to (18) governs pooled client accounts, requiring the firm to understand the account’s purpose and proposed use, be satisfied this is consistent with what it knows of the customer, assess and mitigate the resulting risk, and demonstrate all of it to its supervisor. It also binds the customer, who must on request identify the persons whose money is held and their beneficial owners, keep written records of every payment in and out for five years, and provide information to law enforcement, subject to legal professional privilege.

■ Customer due diligence and beneficial ownership

Regulation 27 sets the triggers: establishing a business relationship, an occasional transaction above the applicable threshold, suspicion of money laundering or terrorist financing, and doubts about documents or information previously obtained. Regulation 28 sets the content: identify the customer and verify their identity from a reliable source independent of the customer; identify any beneficial owner and take reasonable measures to verify their identity so the firm is satisfied it knows who they are; where the customer is a legal person, understand its ownership and control structure; and assess the purpose and intended nature of the relationship. Regulation 28(11) requires ongoing monitoring, including scrutiny of transactions for consistency with knowledge of the customer and keeping information up to date.

Two misreadings deserve naming. Verification is not document collection: a passport photocopy verifies nothing unless the firm has satisfied itself the document is genuine and belongs to the person presenting it. And know your customer is not a one-off event at onboarding; regulation 28(11)'s ongoing obligation is what enforcement actions most often find missing, because the burden of periodic review across a large book is enormous and the incentive to defer it constant. Record-keeping runs under regulation 40: five years from the date the firm knows or has reasonable grounds to believe the relationship ended, after which regulation 40(5) requires deletion of the personal data unless another legal basis applies.

On ownership, regulation 5 covers bodies corporate and regulation 6 partnerships and trusts. The 25 per cent threshold in regulation 5(1)(b) is the number practitioners quote, but the limbs either side of it matter as much: an individual exercising ultimate control over management, and an individual who otherwise controls the body, are beneficial owners regardless of shareholding. Where none can be identified after all reasonable measures, the senior person responsible for managing the body may be treated as such, a fallback that should be rare and is not.

The public registers are the second leg. The United Kingdom's register of people with significant control, at Companies House, uses a threshold aligned to 25 per cent, and regulation 30A requires relevant persons to report discrepancies between what a customer tells them and what the register says. Register quality was for years the weak point, and the Economic Crime and Corporate Transparency Act 2023 is the response: at the time of writing, identity verification became a legal requirement for directors and people with significant control from 18 November 2025, with new people with significant control required to verify within 14 days of their appointment being notified. This does not solve the ownership problem, but it changes the cost of lying about it from nothing to something.

■ Enhanced due diligence

Regulation 33 sets out when enhanced customer due diligence and enhanced ongoing monitoring are mandatory, and the 2026 amendment made one materially important change. Previously regulation 33(1)(b) bit on relationships and transactions involving a "high-risk third country" as separately defined. From 30 June 2026 it bites on a "FATF call for action country", defined by reference to the Financial Action Task Force's list of High-Risk Jurisdictions subject to a Call for Action as it has effect from time to time. At the FATF plenary of 19 June 2026 that list comprised the Democratic People's Republic of Korea, Iran and Myanmar. The separate list of Jurisdictions under Increased Monitoring, informally the grey list, no longer triggers mandatory enhanced due diligence in the United Kingdom, though FATF mutual evaluations remain a geographical risk factor under regulation 33(6)(c) and a firm must still apply enhanced measures wherever it identifies high risk. HM Treasury's Money Laundering Advisory Notice of June 2026 sets out the position.

Politically exposed persons sit at regulation 35. The definition at regulation 35(12)(a) is an individual entrusted with prominent public functions other than as a middle-ranking or more junior official, with a non-exhaustive list at regulation 35(14) running from heads of state and government, ministers and legislators through senior judiciary, central bank boards, ambassadors and high-ranking armed forces officers to the boards of state-owned enterprises and international organisations. Family members include spouses and civil partners, children and their spouses or civil partners, and parents. Known close associates covers joint beneficial ownership or other close business relations, and sole beneficial ownership of an entity known to have been set up for the PEP's benefit, with regulation 35(15) limiting the enquiry to information in the firm's possession or credible publicly available information.

Where a PEP relationship is proposed or continued, regulation 35(5) requires senior management approval, adequate measures to establish source of wealth and source of funds, and enhanced ongoing

monitoring. The two sources are not synonyms, and conflating them is a common finding: source of wealth explains how the person's total assets were accumulated, source of funds where the money in this transaction came from. Regulation 35(9) requires the enhanced measures to continue for at least 12 months after the person ceases to be entrusted with the function, or longer if the firm considers it appropriate; regulation 35(11) provides that they fall away for family members and close associates when the PEP ceases to hold the function, whether or not that period has run.

Domestic PEPs are treated distinctly. Section 77 of the Financial Services and Markets Act 2023 required the Treasury to amend Part 3 of the Money Laundering Regulations so that where the customer is a domestic PEP, or a family member or known close associate of one, the starting point for the regulation 35(3) assessment is that they present a lower level of risk than a non-domestic PEP, and where no enhanced risk factors are present the extent of enhanced measures is less. That is now regulation 35(3A), with enhanced risk factors defined at regulation 35(12)(f) as risk factors other than the domestic PEP status itself. Section 78 required the FCA to review its PEP guidance and publish conclusions within twelve months. Its multi-firm review of 18 July 2024 found firms using definitions of relatives and close associates wider than the Regulations allow, ineffective arrangements for reassessing status after a person leaves office, risk ratings without clear rationale, and poor communication with affected customers. Finalised guidance FG25/3 followed on 7 July 2025, revised on 15 July 2025 to clarify that non-executive board members of United Kingdom civil service departments should not be treated as PEPs.

Simplified due diligence sits at regulation 37 and is permitted only where the firm has determined the relationship or transaction presents a low degree of risk, having taken account of the risk factors in Schedule 3. The 2026 amendment extended regulation 37(2)(a) to cover the new pooled account provisions at regulation 29(11) to (13), which was the point of that reform: making clear when simplified measures may be applied to accounts holding many clients' money.

■ Sanctions screening

Sanctions compliance is legally separate from anti-money-laundering compliance, and the vocabulary should stay separate too. In the United Kingdom, designations are made under regulations issued pursuant to the Sanctions and Anti-Money Laundering Act 2018. At the time of writing, and this is a recent change that will catch out any process built before it, the UK Sanctions List maintained by the Foreign, Commonwealth and Development Office is the single authoritative list of designations: the OFSI Consolidated List of Asset Freeze Targets closed on 28 January 2026 and is no longer updated. Any screening system still pointed at the old file is reading a frozen snapshot.

The Office of Financial Sanctions Implementation enforces financial sanctions. Under section 146 of the Policing and Crime Act 2017 as amended, and as set out in OFSI's *Financial sanctions enforcement and monetary penalties guidance* last updated on 9 February 2026, the maximum civil monetary penalty is the greater of £1,000,000 or 50 per cent of the estimated value of the funds or economic resources involved, or £1,000,000 where no value can be estimated. OFSI has stated an intention to seek to double this to the greater of £2,000,000 or 100 per cent of the breach value; at the time of writing that had been announced but the statutory cap was unchanged. Relevant firms must also notify HM Treasury where they know or have reasonable cause to suspect that a person is designated or has breached a prohibition, and where they hold frozen assets.

The United States position reaches further. The Office of Foreign Assets Control maintains the Specially Designated Nationals and Blocked Persons list and applies its 50 Percent Rule, under which an entity owned in the aggregate, directly or indirectly, 50 per cent or more by one or more blocked

persons is itself blocked whether or not it appears on any list. That is why ownership analysis and sanctions screening cannot be separate workflows: an entity can be blocked without ever being named.

The screening itself is a name-matching problem with unusual properties. Names arriving from customer records are not canonical: the same person may appear with given and family names reversed, with or without patronymics, with diacritics stripped, transliterated from Arabic, Cyrillic, Chinese or Korean by any of several competing conventions, abbreviated, or misspelled. Matching engines therefore combine exact matching with phonetic algorithms, edit-distance measures and token-based comparison, each returning a score, with a configurable threshold above which an alert is raised. Lower the threshold and recall rises while precision falls; raise it and the reverse. No setting gives both, and the threshold is a risk decision a supervisor will ask to see justified.

What reduces alert volumes without reducing detection is secondary identifiers and good data. Date of birth, place of birth, nationality, passport or national identity number and full address, held in structured fields rather than free text, let an engine discount a name match that is impossible on other grounds; list entries carry these to varying degrees, and the poorest-quality entries generate the most noise. Commonly cited industry estimates put the share of alerts closed with no report at between 85 and 95 per cent, but those figures come from vendors and industry bodies rather than regulators and should be treated as indicative. The cost also falls unevenly: common names in large populations, and naming conventions producing many identical full names, concentrate screening burdens on particular ethnic and national groups without anyone having designed that outcome. A firm that does not measure alert rates by customer segment will not know it is happening.

■ Transaction monitoring and suspicious activity reporting

Regulation 28(11) requires ongoing monitoring; how it is done is left to the firm. In practice there are two families of approach and mature firms run both.

Rules-based monitoring encodes typologies as deterministic tests: cash deposits aggregating above a value in a rolling window, rapid movement in and out leaving a near-zero balance, transactions structured just below a reporting threshold, activity inconsistent with the stated business, payments to jurisdictions the firm has rated high risk, or sudden dormancy reversal. Rules are explainable, testable and auditable, which regulators like, and trivially avoidable by anyone who learns the thresholds, which is their weakness. Model-based monitoring scores behaviour against learned baselines, per customer or per peer group, and surfaces the anomalous. It catches patterns nobody wrote a rule for, at the cost of two problems: the firm must explain to a supervisor why a model produced a given output, and a model trained on historically labelled data learns the biases in that labelling. Above-the-line and below-the-line testing and documented calibration are standard supervisory expectations rather than sophistication.

When suspicion arises, the firm's nominated officer, appointed under regulation 21(3) and known in FCA-regulated firms as the money laundering reporting officer, decides whether to report. In the Senior Managers and Certification Regime this is senior management function SMF17, a named individual with personal accountability. The report goes to the UK Financial Intelligence Unit at the National Crime Agency through the SAR Portal. From the SARs Annual Report 2025, covering April 2024 to March 2025:

Measure	2023-24	2024-25
SARs received	872,048	866,616
DAML requests received	57,081	57,666

Measure	2023-24	2024-25
Funds denied through DAML requests	£240.1m	£382.6m
Account freezing orders, forfeitures and restraints obtained through DAMLs	1,785	2,048
Terrorism Act SARs received	1,132	1,101
Defence Against Terrorist Financing requests	406	397
Reported SAR confidentiality breaches	8	7

Two further figures bear on how the system behaves. The average turnaround for a DAML decision was 2.8 days against the statutory seven working day period; and 151 refused cases required one or more moratorium extensions under section 336A, involving 18 law enforcement agencies and denying £120.6 million, of which £103 million came from a single DAML. The headline totals are heavily influenced by a few large cases, which is why the UKFIU publishes a three-year rolling average alongside them: £298.4 million for the period ending 2024-25.

The distribution by sector is equally telling. Of SARs submitted through the SAR Portal in 2024-25, banking accounted for 85.54 per cent and financial services a further 10.34 per cent, with virtual assets at 1.67 per cent, gambling 0.88 per cent, accountancy 0.71 per cent, legal 0.40 per cent and property 0.10 per cent. The UKFIU is careful to say it is for each sector and its supervisor to judge whether that volume is proportionate to risk. It is nonetheless clear that the overwhelming majority of the United Kingdom's financial intelligence comes from banks, which means anything a bank cannot see is largely invisible to the system.

The United States regime under the Bank Secrecy Act is more mechanical. Under 31 CFR 1020.320 a bank must file a suspicious activity report no later than 30 calendar days after initial detection of facts that may constitute a basis for filing, with an additional 30 where no suspect has been identified, and may not disclose the report or anything revealing its existence. Alongside it sits the currency transaction report for cash above \$10,000, a threshold report rather than a suspicion report, with a corresponding offence of structuring to evade it.

■ The risk-based approach

The risk-based approach is FATF Recommendation 1, and it is not permission to do less. It is an instruction to allocate effort in proportion to assessed risk, which requires a firm to have assessed risk in writing. In the United Kingdom that is regulation 18, requiring a written risk assessment that takes account of information from supervisors and considers risk factors relating to customers, countries, products, services, transactions and delivery channels, and regulation 19, requiring policies, controls and procedures proportionate to that assessment and approved by senior management.

The national layer sits above it. HM Treasury and the Home Office published the *National Risk Assessment of Money Laundering and Terrorist Financing 2025* on 17 July 2025, the fourth comprehensive United Kingdom assessment. It rates retail banking, money service businesses, and payment and electronic money institutions as high risk for both money laundering and terrorist financing, and cryptoasset service providers as high for money laundering and medium for terrorist financing. It assesses that trade-based schemes move over £10 billion, and describes a realistic possibility that up to £10 billion could be laundered through the United Kingdom property market annually. A firm's regulation 18 assessment is expected to engage with those findings, not restate them.

Internationally the emphasis has shifted. In February 2025 the FATF revised Recommendation 1 and its Interpretive Note to strengthen the requirement that controls be proportionate and to encourage countries to promote financial inclusion, and later published *Guidance on Financial Inclusion and Anti-Money Laundering and Terrorist Financing Measures*. The FATF is explicit, including in the standing text accompanying its increased-monitoring list, that its standards do not envisage de-risking or cutting off entire classes of customer, and that flows for humanitarian assistance, legitimate non-profit activity and remittances should be neither disrupted nor discouraged. That followed its February 2021 project on the unintended consequences of the standards, which led in 2023 to amendments to Recommendation 8 curbing disproportionate measures against non-profit organisations.

■ Correspondent banking and de-risking

A correspondent relationship is one bank providing services, typically an account and payment execution, to another bank in another jurisdiction. It is the mechanism by which a bank in a small economy reaches the dollar, the euro or sterling at all.

The due diligence is heavier than for an ordinary customer because the correspondent takes on not just the respondent but, indirectly, the respondent's entire customer base. Regulation 34 requires a firm entering a correspondent relationship with a respondent from outside the United Kingdom to gather sufficient information to understand its business fully, determine its reputation and the quality of its supervision from credible public sources, assess its financial crime controls, obtain senior management approval, and document the respective responsibilities. Relationships with shell banks are prohibited outright, as is continuing one with an institution known to allow its accounts to be used by a shell bank. From 1 February 2027 the new regulation 34A extends materially the same requirements to cryptoasset exchange providers and custodian wallet providers dealing with third-country providers of similar services. The standard instrument for gathering this information is the Wolfsberg Group's Correspondent Banking Due Diligence Questionnaire, at the time of writing at version 1.4, with the Group's FAQs recommending a refresh cycle of 12 to 18 months.

The pattern in the data is a long, uneven contraction. The Committee on Payments and Market Infrastructures at the Bank for International Settlements publishes a quantitative review of correspondent banking data; its August 2020 commentary reported that the number of active correspondents worldwide had declined by about 22 per cent between 2011 and 2019 and by about 3 per cent in 2019 alone, with country-level declines ranging from 23 per cent in advanced economies to 41 per cent in small island developing states and dependent territories, and regional declines over the same period from 13 per cent in Northern America to 34 per cent in Latin America. Meanwhile the value and volume of payments running through those networks kept growing, by roughly 5 and 4 per cent respectively in the final year measured. Fewer relationships carrying more traffic is concentration, not shrinkage, and concentration means a single correspondent's exit can disconnect a jurisdiction.

The consequences for smaller markets are structural. Where a country loses its last correspondent for a currency, payments must route through an additional intermediary, raising cost, lengthening settlement and reducing transparency. The National Risk Assessment 2025 records a rise in nesting, where correspondent banks hold accounts for other correspondent banks so the chain becomes longer and more opaque, and attributes this partly to de-risking. That is how de-risking makes the underlying problem worse: the flows do not stop, they move somewhere with less visibility. The same report notes that de-risking of non-profit organisations, combined with the absence of formal banking where they operate, pushes them towards cash and less well-regulated channels. The Financial Stability Board launched a four-point action plan in November 2015 to address the decline, and correspondent

banking remains part of the G20 cross-border payments roadmap, but none of it changes the arithmetic facing a bank weighing bounded revenue against unbounded liability.

■ Supervision, enforcement and what will date this chapter

The United Kingdom's supervisory architecture is fragmented. The FCA supervises banks, payment institutions, electronic money institutions and registered cryptoasset businesses. His Majesty's Revenue and Customs supervises money service businesses, trust or company service providers, high value dealers, art market participants, and estate agency and accountancy businesses not otherwise supervised. Professional body supervisors cover the legal and accountancy sectors, overseen by the Office for Professional Body Anti-Money Laundering Supervision within the FCA. The Gambling Commission supervises casinos.

Enforcement inside the FCA's perimeter is generally brought under the Principles for Businesses rather than the Regulations, most often Principle 3 on management and control systems. The Nationwide Building Society case of December 2025 is representative of both mechanics and amounts: £44,078,500, reduced by 30 per cent from £62,969,297 for early resolution, for inadequate anti-financial-crime systems and controls between October 2016 and July 2021, specifically ineffective arrangements for keeping due diligence and risk assessments current across personal current account customers and for monitoring their transactions. The pattern across such cases is instructive: the failings are almost never an absence of policy. They are failures to apply it at scale, to keep customer information current, to tune monitoring to the actual business, and to close the loop between an alert and a decision.

The European Union is moving the opposite way. Regulation (EU) 2024/1624, the single rulebook, applies directly in all Member States from 10 July 2027, with Directive (EU) 2024/1640 to be transposed by the same date and Regulation (EU) 2024/1620 establishing the Authority for Anti-Money Laundering and Countering the Financing of Terrorism, based in Frankfurt and operational since 1 July 2025, expected to begin directly supervising a small group of the highest-risk cross-border institutions from 2028.

Four things here will date, and a reader in 2029 should verify all of them. The United Kingdom's Regulations are mid-reform: the 2026 amendment has commencement dates still ahead of it, on 1 February 2027 and 25 October 2027, and HM Treasury committed in July 2025 to further guidance on the meaning of establishing a business relationship, when source of funds checks are required, and when enhanced due diligence is mandatory. The lists move by design: FATF updates its at each plenary, roughly three times a year, and the United Kingdom's mandatory trigger now tracks the call for action list as it has effect from time to time, while the UK Sanctions List became the single authoritative designation list on 28 January 2026. OFSI's announced doubling of its civil penalty maximum had not taken effect as this was written, and the FCA enforcement figure was current at December 2025. And the European regime has not yet started, so the comparison drawn here between a deregulating United Kingdom and a harmonising European Union is a hypothesis that will by then have been tested.

52.98 Common wrong ideas

Wrong: Laundering is about hiding money. **Right:** It is about explaining it; the car wash owner does not conceal the criminal cash but manufactures a believable answer to where it came from, and pays tax on it because the tax bill is the receipt that proves the story.

Wrong: The anti-money-laundering regime is a net for catching launderers. **Right:** Firms are measured on process rather than outcomes and can comply impeccably while the money still gets away;

the published numbers describe a very large privately funded intelligence-gathering exercise from which comparatively few interventions are extracted.

Wrong: A wrong name match is an error to be engineered away. **Right:** It is the intended output of a system deliberately tuned to over-report, because missing a genuine sanctions target is a strict-liability breach while stopping an innocent customer for ten minutes is an inconvenience.

Wrong: A low-risk customer can be given lighter sanctions screening. **Right:** The risk-based approach governs due diligence only; an asset freeze is an absolute prohibition, and for breaches on or after 15 June 2022 any requirement that the person knew, suspected or believed anything is disregarded.

Wrong: Knowing who the customer is settles the identity question. **Right:** Identity is the easy half; the “more than 25 per cent” bright line is designed around by four holders at 25 per cent each, by chains of companies across jurisdictions, and by nominees, trusts and companies limited by guarantee.

Wrong: Collecting a passport copy verifies identity. **Right:** Verification requires the firm to satisfy itself from a source independent of the customer that the document is genuine and belongs to the person presenting it.

Wrong: Know your customer is something you complete at onboarding. **Right:** Regulation 28(11) makes scrutiny of transactions and keeping information up to date a continuing obligation, and it is the limb enforcement actions most often find missing.

Wrong: Source of wealth and source of funds are two phrases for the same check. **Right:** Source of wealth explains how the person’s total assets were accumulated and source of funds where the money in this transaction came from, and conflating them is a common finding.

Wrong: The bank is being obstructive when it refuses to explain a delay. **Right:** Tipping off is a criminal offence in the United Kingdom, so the staff facing the customer are legally forbidden to explain, which is at the bottom of almost every furious story of this kind.

Wrong: De-risking shows a firm has failed to apply the risk-based approach. **Right:** It is a rational answer to bounded revenue weighed against an unbounded and asymmetric downside, and its effect is that flows move somewhere with less visibility rather than stopping.

52.99 Chapter summary in 20 lines

1. A banknote carries nothing to say where it came from, so what catches a criminal is the story attached to the money rather than the money itself.
2. Money laundering is the manufacture of believable answers to the question “where did that come from”, and anti-money-laundering is the business of asking it properly.
3. A cash business launders by writing down more takings than it had, and its owner will happily pay tax on money never earned, because the tax bill is the receipt.
4. The three moves are getting the cash in, moving it around until the trail is hard to follow, and spending it on something ordinary.
5. Every rule you have been annoyed by exists to interfere with one of those three moves, which is why opening a business account takes eleven days.
6. Each onboarding question does a job: who you are, what the business does, who else benefits, and where the opening money came from.
7. Identity is the easiest of those and the least important; establishing who ultimately benefits is a judgement rather than a fact.

8. Sanctions screening is a legally separate world from money laundering, with absolute prohibition and strict liability where the other has suspicion and reasonableness.
9. Screening deliberately looks for near matches, so a bank taking on four thousand customers a week may see roughly 120 possible matches of which perhaps two are real.
10. That is a property of names rather than a defect in software, and its cost falls unevenly on people whose naming conventions produce many identical full names.
11. Once the account is open, monitoring compares behaviour against what the customer said and against what similar businesses do, and the great majority of alerts are ordinary life.
12. When a firm reports, it must not tell the customer, because tipping off is itself an offence, which explains the unexplained delay.
13. The architecture rests on Part 7 of the Proceeds of Crime Act 2002, whose principal offences sit at sections 327 to 329 and whose day-to-day grip comes from the failure to disclose and tipping-off provisions.
14. The Defence Against Money Laundering process gives the National Crime Agency seven working days to refuse, after which moratorium extensions can lawfully immobilise funds for over half a year on suspicion alone.
15. The Money Laundering Regulations 2017, as amended in 2026, set out what a firm must build: written risk assessment, due diligence, enhanced measures, ongoing monitoring, record-keeping and reporting.
16. The 2026 amendment converted euro thresholds to sterling, narrowed “complex or unusually large”, added rules for pooled client accounts and moved the mandatory enhanced due diligence trigger onto the FATF call for action list.
17. Transaction monitoring runs on rules, which are explainable and avoidable once the thresholds are known, and on models, which find what nobody wrote a rule for and must still be explained to a supervisor.
18. Banking supplies about 85 per cent of the United Kingdom’s suspicious activity reports, so anything a bank cannot see is largely invisible to the system.
19. Enforcement finds failures to apply policy at scale rather than absent policy, and against penalties measured in tens of millions banks exit whole customer categories.
20. De-risking does not stop the flows; it lengthens the chain through nesting and pushes activity towards cash and less regulated channels, which is a policy problem rather than a compliance one.

Sources, all consulted in August 2026: the Proceeds of Crime Act 2002, sections 327 to 333A, 335, 336A, 339A and 342; the Terrorism Act 2000; the Money Laundering, Terrorist Financing and Transfer of Funds (Information on the Payer) Regulations 2017 as amended by S.I. 2026/621; the Sanctions and Anti-Money Laundering Act 2018; sections 77 and 78 of the Financial Services and Markets Act 2023; section 146 of the Policing and Crime Act 2017 as amended; the Economic Crime and Corporate Transparency Act 2023 and Companies House identity verification guidance; 31 CFR 1020.320; the UKFIU SARs Annual Report 2025; the National Risk Assessment of Money Laundering and Terrorist Financing 2025; HM Treasury’s Money Laundering Advisory Notice of June 2026; the FCA, including FG25/3, the 2024 multi-firm review on politically exposed persons and the Nationwide announcement of December 2025; OFSI’s financial sanctions enforcement and monetary penalties guidance; the FCDO UK Sanctions List; the Financial Action Task Force, including the June 2026 plenary lists, the February 2025 revision to Recommendation 1 and the financial inclusion guidance; BIS CPMI correspondent banking commentaries; the Financial Stability Board; the Wolfsberg Group CBDDQ v1.4; Regulations (EU) 2024/1620 and 2024/1624 and Directive (EU) 2024/1640; the United Nations Office on Drugs and Crime; and the Bank Policy Institute, Getting to Effectiveness, 2018.

Regulation

53.0 What this chapter gives you

1. You will be able to explain why a firm holding £900,000 of other people's money on an ordinary Tuesday is required to have only 125,000 euros of its own capital, and why that is not the mismatch it looks like.
2. You will be able to work out which licence a firm needs — authorised payment institution, small payment institution, authorised or small e-money institution, or registered account information service provider — from its volumes and the services it intends to provide.
3. You will be able to explain what safeguarding actually is, why it dominates everything else in this sector, and what changed on 7 May 2026 when CASS 15 came into force.
4. You will be able to tell somebody why safeguarded money is not insured money, and why the average shortfall across twelve failed payments firms was 65 per cent with distributions taking 2.3 years.
5. You will be able to distinguish the £85,000 that appears in scam reimbursement from the £85,000 people wrongly imagine sits behind safeguarded balances.
6. You will be able to explain why a merchant's card costs keep rising even though interchange is capped at 0.2% on consumer debit and 0.3% on consumer credit.
7. You will be able to name the four public bodies and one department that share the United Kingdom's payments perimeter, and say which statute each draws its powers from.
8. You will be able to say why "PSD2 applies in the UK" is false, and what to say instead when quoting a rule.
9. You will be able to describe what an important business service and an impact tolerance are, and why a service is not made unimportant by being well protected.
10. You will be able to separate what is law from what is announced policy — the abolition of the PSR, the PSD3 package, the repeal of assimilated law — and date any statement you make about this field.

The plain version

Imagine a left-luggage office at a railway station. You arrive with a suitcase, hand it over, and get a paper ticket with a number on it. When you come back you hand over the ticket and you expect the same suitcase.

Now ask what could go wrong, and notice that there are two completely different kinds of wrong.

The first kind: the office runs out of money. Rent goes up, custom falls away, and one morning the shutters stay down. That is a business failing. It is sad, it is common, and mostly the world copes.

The second kind: you come back and your suitcase is gone, because the owner sold it. Or because the owner kept all the suitcases in one heap, lost track of whose was whose, and used a few to prop open a door. That is a different sort of failure entirely, and no amount of the owner's own money in the till would have prevented it, because the problem was never the till. The problem was the room out the back.

Almost everything in payments regulation attempts to make the second kind of failure impossible. When a firm holds your money on its way somewhere else, your money is the suitcase and the firm's own money is the till. Regulators care about the till, because a firm that runs out of cash stops working. They care far more about the room out the back, because that is where your money actually is.

So let us follow a firm through it.

Rina runs a small business called Northgate Pay. Her customers are letting agents. When a tenant pays rent, the money comes to Northgate Pay first, sits for a day or two while the agent's software matches it to a property, and then goes out to the landlord. Northgate Pay never owns that money; it holds it in transit. On an average day about £900,000 belonging to other people is sitting in Rina's accounts, and across a month roughly £27 million passes through.

Four things now happen to Rina, and they are the four things this chapter is about.

She needs permission. Moving other people's money as a business is not something you may simply do. You have to be licensed by the Financial Conduct Authority, the FCA, and the licence comes in different sizes. Because Rina's monthly volume is well over the small-firm threshold, she needs the full version: authorisation as what the law calls an authorised payment institution. That is not a form you fill in over a weekend. It is a business plan, a three-year financial forecast, background checks on everyone who runs or owns the firm, and written descriptions of how she will handle complaints, spot money laundering, cope when her systems break, and keep customer money separate. The regulator has three months to decide, but only once the application is genuinely complete, and most first drafts are not.

She needs starting money of her own. As things stood at the time of writing in August 2026, the law required a firm doing what Rina does to hold initial capital of 125,000 euros. Note the currency: British law here is still written in euros, a leftover from where these rules came from. And note the size: 125,000 euros against £900,000 of customer money sitting there on a normal Tuesday. The capital is not there to cover the customer money and could not possibly cover it. It is there so the firm has a cushion of its own, enough to absorb an ordinary bad year without immediately becoming somebody else's problem.

She must keep the customer money somewhere else entirely. This is the room out the back, and in the trade it is called safeguarding. Rina must hold the rent money in a separate account, at a different bank, clearly marked as containing other people's money, with a letter from that bank acknowledging in writing that it has no right to grab the money to settle Rina's overdraft. Anything she is still holding at the end of the working day after the day it arrived has to be in there. Not most of it. All of it. And under rules that came into force on 7 May 2026, she must count it every business day, twice over: once against her own records, once against what the bank says is actually there. Differences are fixed the same day. Once a month she sends the regulator a return. Once a year an outside auditor checks the arrangement, unless she is small enough to be exempt, and the exemption line was drawn at £100,000 of customer money over a fifty-three week look-back.

And she must decide, in advance, how long she is allowed to be broken. The regulator does not ask Rina to promise her systems will never fail. It asks her to name the handful of things she does that

would genuinely hurt somebody if they stopped — for her, getting rent to landlords — and then to write down the longest outage a customer could survive. Four hours. Twelve. A day. Then she must map everything that service depends on, including other people's computers she does not control, and test whether she could really stay inside her own number when things go badly wrong. The regulator's interest is not in the promise. It is in whether she has done the arithmetic honestly.

Now the fifth thing, which is not about Rina at all, but about the price of card payments.

When you buy a £50 pair of shoes with a debit card, the shop does not get £50. A small slice is taken out and passed back to the bank that issued your card. That slice is called interchange. Left alone it would be set by the card networks, who are in the odd position of setting a fee paid by one of their customers to another, with no particular pressure to keep it low.

So the law caps it. As the rules stood at the time of writing in August 2026, interchange on a consumer debit card transaction within the UK was capped at 0.2% of the value, and on a consumer credit card at 0.3%. On your £50 shoes, that is ten pence on debit and fifteen pence on credit. It sounds trivial. Across the country it is billions of pounds a year, which is why it is one of the most bitterly litigated numbers in British commerce.

And finally, the thing that makes this chapter the one most likely to date: the map of who is in charge is itself being redrawn. There were, at the time of writing, two payments regulators in the UK. The FCA licenses and supervises firms. The Payment Systems Regulator, the PSR, regulates the shared plumbing, and it was the PSR that set the interchange rules and the rules on refunding scam victims. In March 2025 the government announced it would abolish the PSR and move its work into the FCA, and the consultation response confirming that decision was published on 21 April 2026. As at the time of writing in August 2026 the PSR still existed and still had its powers, because the merger needs an Act of Parliament that had not yet been passed.

That is the plain version. Permission, capital, the room out the back, an honest number for how long you may be broken, a cap on the card fee, and two regulators becoming one.

Where the plain version stops being true

Safeguarded is not insured, and the £85,000 you have heard of is not this £85,000. The analogy suggests that if the office closes you walk round the back and collect your suitcase. It is not that clean. Money held by a payment institution or an e-money institution is not a bank deposit and is not covered by the Financial Services Compensation Scheme. There is no government-backed £85,000 guarantee behind it. What safeguarding creates is a pool of assets meant to be paid out to customers ahead of the failed firm's ordinary creditors, and whether that pool is full is a question of fact. Historically it very often was not. The FCA stated in its September 2024 consultation CP24/20 that of the twelve payments firms that became insolvent between the first quarter of 2018 and the second quarter of 2023, the average shortfall between funds owed to customers and funds actually safeguarded was 65%, and for e-money institutions alone 80%. Where money was eventually distributed it took on average 2.3 years to reach customers. Sixty-five pence in the pound, two and a bit years late, is a materially different outcome from collecting your suitcase. Separately, the £85,000 that appears in British payments conversation usually refers to something else entirely: the maximum reimbursement per claim under the mandatory refund rules for authorised push payment scams, confirmed by the PSR in policy statement PS24/7 and applying since 7 October 2024. Same number, different mechanism, different money, no connection.

There is no single “payments regulator”, and several of the most binding rulebooks are not law at all. The plain version names two regulators and implies they cover the field. They do not. The Bank of England supervises the payment systems that matter most to financial stability under Part 5 of the Banking Act 2009; the Prudential Regulation Authority supervises banks’ safety and soundness; anti-money-laundering supervision runs on its own track under the Money Laundering Regulations 2017; data protection sits with the Information Commissioner’s Office. And running alongside all of it are rulebooks that bind harder than most legislation while being nobody’s law: the Visa and Mastercard operating rules, private contracts revised at least annually and enforced with fines through the acquiring chain; Pay.UK’s scheme rules for Bacs, Faster Payments and Direct Debit; the PCI Data Security Standard, published by an industry council with no statutory authority whatsoever, which nonetheless determines whether a merchant may accept cards at all. A firm can be perfectly compliant with the Payment Services Regulations and be shut down in a fortnight by a card scheme. Practitioners spend more of their week on the private rulebooks than on the statutes.

PSD2 is not, and never was, the law in the United Kingdom, and the British copies of the European rules are actively drifting. People say “PSD2 applies in the UK” as shorthand and it is false in a way that matters. The Second Payment Services Directive was a directive: an instruction to member states to legislate. The United Kingdom’s legislation was the Payment Services Regulations 2017 and a set of FCA rules, and after departure from the European Union what remains is assimilated law that Parliament and the FCA can change unilaterally, and have. The Financial Services and Markets Act 2023 provides for the wholesale repeal of assimilated financial services law, to be replaced by FCA Handbook rules, on a timetable still running. Divergence already exists: on 19 March 2026 FCA rule changes took effect removing the mandatory £100 single-transaction contactless limit and allowing firms with strong fraud controls to set their own, a change with no European equivalent. Meanwhile the European Union is replacing PSD2 outright. Anyone quoting an article number should say which body of law they are quoting and on what date.

The interchange cap does not cap what a merchant pays, and it does not apply to most of the transactions merchants complain about. Interchange is one component of the merchant service charge. The other two are the scheme fees, charged directly by the card networks and not capped at all, and the acquirer’s own margin. Capping interchange at 0.2% has never stopped the total cost of card acceptance rising. Further, the caps apply to consumer debit and credit cards on transactions within the relevant territory. They do not apply to commercial cards, nor where the issuer is outside the territory, and three-party schemes such as American Express fall outside them except in defined circumstances. This is not a loophole; it is the deliberate scope of the rule. It is also why the most expensive card in a merchant’s mix is usually a foreign-issued commercial credit card, on which no cap bites at all.

The technical version

■ The institutional map, as at August 2026

Four public bodies and one department share the United Kingdom’s payments perimeter, and they derive their powers from four different statutes.

The **Financial Conduct Authority** authorises, registers and supervises payment service providers under the Payment Services Regulations 2017 (SI 2017/752) and electronic money issuers under the Electronic Money Regulations 2011 (SI 2011/99). For banks and building societies it acts as conduct regulator under the Financial Services and Markets Act 2000. It also publishes the finalised guidance *Payment Services and Electronic Money — Our Approach*, not binding but the document practitioners

actually read; version 6 appeared in November 2024 and an amended version was published in draft alongside PS25/12 for May 2026.

The **Payment Systems Regulator** is the economic regulator of designated payment systems under Part 5 of the Financial Services (Banking Reform) Act 2013. It is also the competent authority for the interchange fee caps under the Payment Card Interchange Fee Regulations 2015 (SI 2015/1911), and holds functions under Part 8 of the PSRs 2017 on access to payment systems and under the Payment Accounts Regulations 2015 on the Current Account Switch Service.

The **Bank of England** supervises recognised payment systems and specified service providers under Part 5 of the Banking Act 2009, and operates CHAPS and the real-time gross settlement infrastructure. The **Prudential Regulation Authority** supervises deposit-takers and designated investment firms. **HM Treasury** sets the framework: its November 2024 National Payments Vision and the associated Payments Forward Plan set the direction, and the Payments Vision Delivery Committee coordinates.

On 11 March 2025, as part of its Regulation Action Plan, the government announced its intention to abolish the PSR and consolidate its functions primarily within the FCA. It consulted in September 2025 in *A Streamlined Approach to Payment Systems Regulation* and published the consultation response on 21 April 2026. The response confirms the intention to transfer the PSR's functions entirely to the FCA, integrating them into the FSMA framework so far as practicable; to retain a designation regime for bringing payment systems in and out of scope; to carry over objectives equivalent in substance to the PSR's competition, innovation and service-user objectives; to transfer its functions under the PSRs 2017, the Payment Card Interchange Fee Regulations 2015 and the Payment Accounts Regulations 2015, together with the authorised push payment reimbursement requirements; to create a single access regime by removing the Part 8 PSRs route in favour of the FSBRA one; and to make specific directions appealable to the High Court rather than the Competition Appeal Tribunal. Consolidation requires primary legislation, and the response says only that the government will legislate "as soon as parliamentary time allows". As at the time of writing in August 2026 that legislation had not been introduced and the PSR remained a separate body.

■ Authorisation categories under the PSRs 2017 and EMRs 2011

Five categories matter. The distinctions are about scale, about whether the firm issues e-money, and about which of the payment services in Schedule 1 to the PSRs 2017 it intends to provide.

Category	May do	Scale limit	Initial capital	Ongoing own funds
Authorised payment institution	Any Schedule 1 payment service, including payment initiation and account information	None	EUR 125,000 general; EUR 50,000 payment initiation only; EUR 20,000 money remittance only; none for account information only	Method A, B or C as the FCA directs
Small payment institution	Schedule 1 services other than payment initiation and account information	Monthly average of payment transactions over the preceding 12 months not exceeding EUR 3 million	None	None

Category	May do	Scale limit	Initial capital	Ongoing own funds
Authorised e-money institution	Issue e-money and provide payment services	None	EUR 350,000	2% of average outstanding e-money; Method A, B or C for unrelated payment services
Small e-money institution	Issue e-money and payment services other than payment initiation and account information	Average outstanding e-money not exceeding EUR 5 million; unrelated payment transactions not exceeding EUR 3 million monthly average	Set by Schedule 2 EMRs	2% of average outstanding e-money where required
Registered account information service provider	Account information services only	None	None	None

All figures in that table are as they stood at the time of writing in August 2026, and all are stated in euros in the United Kingdom's own legislation, the amounts having been carried across from the European originals without conversion.

Three points a practitioner will insist on. First, a payment institution may not issue e-money; an e-money institution may do both, which is why firms issuing prepaid cards, wallets or stored balances need the EMI licence and the higher capital. Second, initial capital is a floor applying from day one, whereas own funds is a continuing calculation. Method A is 10% of the previous financial year's fixed overheads. Method B is a tiered percentage of payment volume: 4% of the first EUR 5 million of monthly volume annualised, 2.5% of the next EUR 5 million, 1% of the next EUR 90 million, 0.5% of the next EUR 150 million and 0.25% of the remainder, times a scaling factor of 0.5 for money remittance and 1 otherwise. Method C applies tiered percentages to a relevant indicator built from interest income and expense, commissions and other operating income, with a floor of 80% of the three-year average. The FCA directs which applies. Third, providing payment services in the UK without authorisation or registration is a criminal offence under regulation 138 of the PSRs 2017, not a technical breach.

On process: under regulation 9 of the PSRs 2017 and regulation 9 of the EMRs, the FCA must determine a complete application within three months of receiving it, and an incomplete one within twelve months. The clock runs from completeness, not submission, which is why published median timelines bear little relation to the statutory figure; the Approach Document confirms the point at paragraph 3.196. Agents are separately registered under regulation 34 and the principal remains fully responsible for their acts and omissions, while distributors of e-money are not registered in the same way, a difference that has caused repeated supervisory difficulty.

■ Safeguarding, and why it dominates everything else

The prudential logic here is unusual. Capital requirements in banking exist to absorb losses so the institution continues as a going concern. In payments the capital numbers are deliberately small relative to the customer balances held, because these firms do not take credit risk on customer money as a

bank does. The regulatory weight therefore falls almost entirely on the gone-concern question: when this firm fails, do its customers get their money back, in full, quickly?

Under regulation 23 of the PSRs 2017, an authorised payment institution must safeguard “relevant funds”: broadly, sums received from or for the benefit of a payment service user for the execution of a payment transaction, and sums received from another provider for that purpose. Regulations 20 to 22 of the EMRs impose the equivalent obligation on e-money institutions for funds received in exchange for e-money issued.

There are two permitted methods and a firm may use each for different funds. The **segregation method** requires relevant funds to be kept separate from the firm’s own money and, where still held at the end of the business day following the day of receipt, placed in a separate account with an authorised credit institution or the Bank of England, or invested in secure liquid assets. The **insurance or guarantee method** requires cover by a policy with an authorised insurer, or a comparable guarantee from an authorised insurer or credit institution, with proceeds payable into a separate account on insolvency. In practice the overwhelming majority segregate, because insurers write very little of this cover and price it punitively.

The gap between what the regulations say and what happens in insolvency was exposed by *Re Ipagoo LLP* [2022] EWCA Civ 302, decided on 9 March 2022. Ipagoo was an authorised e-money institution that went into administration having failed to safeguard sufficient relevant funds. The FCA argued that the EMRs impose a statutory trust over relevant funds from the moment of receipt. The Court of Appeal held that they do not. It did hold that the regulations create a statutory asset pool for the benefit of e-money holders, and that where relevant funds should have been safeguarded but were not, the pool falls to be topped up from the general estate. Customers therefore rank ahead of ordinary creditors in respect of the pool, but without the proprietary certainty a trust would give, and insolvency costs can bite in ways they would not under a clean trust.

That case, together with the insolvency data, drove consultation CP24/20, published on 25 September 2024 and closed on 17 December 2024, which proposed two stages: an interim set of rules to raise compliance with the existing statutory requirements, and an end-state regime replacing the statutory safeguarding provisions with a sourcebook modelled on the Client Assets Sourcebook, including an express statutory trust.

The outcome was policy statement **PS25/12, published on 7 August 2025**, with the interim rules re-named the **Supplementary Regime** and the end-state rules the **Post-Repeal Regime**. The Supplementary Regime, made by the Payments and Electronic Money (Safeguarding) Instrument 2025 and inserted into the Handbook as **CASS 15, came into force on 7 May 2026**. Its principal requirements, as stated in PS25/12 and accurate at the time of writing in August 2026, are internal and external reconciliations of safeguarded funds on each business day, not required on weekends, public holidays and days when relevant foreign markets are closed, with discrepancies resolved promptly; a new monthly safeguarding return, revised after a pilot exercise; an annual safeguarding audit by a qualified auditor for authorised payment institutions and authorised e-money institutions, subject to an exemption where the firm has not been required to safeguard more than £100,000 of relevant funds at any time over a period of at least fifty-three weeks; a contingency plan at least three months before the expiry of any safeguarding insurance policy or comparable guarantee, that lead period retained despite industry objection; retention of the existing range of permitted secure liquid assets, the FCA explicitly maintaining a low risk tolerance towards broadening it; and resolution pack requirements aligned to CASS practice. The FCA also removed the proposed limited assurance engagement for firms claiming

to hold no relevant funds, replacing it with guidance that a statutory auditor should notify the FCA if a firm has failed to safeguard.

Two significant CP24/20 proposals did **not** proceed at this stage: the requirement to receive relevant funds directly into a designated safeguarding account, and the corresponding treatment of funds received through agents and distributors. Both were held back pending further consideration alongside the statutory trust proposal. The FCA has said it will not proceed with the Post-Repeal Regime without further consultation, and will review implementation once a full audit period has been completed.

The regime applies to authorised payment institutions other than those providing only payment initiation or account information services, authorised e-money institutions, small e-money institutions, and credit unions issuing e-money in the UK. Small payment institutions may opt in.

Underneath all of this sits the **Payment and Electronic Money Institution Insolvency Regulations 2021 (SI 2021/716)**, a bespoke special administration regime for these firms with the return of relevant funds as a statutory objective and a court-approved distribution plan mechanism. HM Treasury launched an independent review of it in 2024 and the report has since been published.

■ Interchange regulation

The caps entered European law through Regulation (EU) 2015/751 on interchange fees for card-based payment transactions, applying from 9 December 2015. That regulation was assimilated into United Kingdom law and amended so that its operative provisions now refer to UK transactions.

As the assimilated regulation stood at the time of writing in August 2026, Article 3 provides that payment service providers must not offer or request a per-transaction interchange fee of more than 0.2% of transaction value for any UK debit card transaction, with a Treasury power to lower that percentage, impose a fixed maximum, or permit an alternative not exceeding the sterling equivalent of EUR 0.05 per transaction, provided the scheme's aggregate annual interchange stays within 0.2% of annual transaction value. Article 4 sets the equivalent cap for any UK credit card transaction at 0.3%, again with a Treasury power to lower but not raise it. Legislation.gov.uk records the regulation as subject to revocation under Schedule 1 to the Financial Services and Markets Act 2023, a change not yet applied at the time of writing.

The regulation is not only about price. Articles 6 to 11 carry the business rules that shape the acquiring market: licensing, prohibiting territorial restrictions; separation of payment card scheme and processing entities; co-badging and the payer's right to choose the brand or application at the point of sale; unblending, which requires acquirers to offer merchant service charges itemised by category and brand of card rather than a single blended rate; the Honour All Cards rule, which prevents schemes obliging merchants to accept every card of a brand irrespective of category; and steering, which permits merchants to steer customers towards cheaper instruments. Unblending made the true cost of card acceptance visible to merchants, and is the reason interchange-plus pricing exists at all. Separately, regulation 166 of the PSRs 2017 prohibits surcharging on most consumer card payments in the UK.

The post-departure history is the interesting part. After the end of the transition period, Mastercard and Visa increased interchange on card-not-present transactions between the UK and the European Economic Area, for consumer debit and credit cards, from 0.2% and 0.3% to 1.15% and 1.5% respectively. The PSR opened market review MR22/2. Its final report, MR22/2.7, found that Mastercard and Visa, which between them account for 99% of UK debit and credit card payments, were not subject to effective competitive constraints; that the increases were costing UK businesses an additional £150 million to £200 million annually; and that neither scheme could demonstrate any specific assessment

underpinning the decision. The PSR consulted on a two-stage remedy: an interim cap followed by a longer-term cap.

Visa, Mastercard and Revolut challenged the PSR's power to impose such a cap by judicial review. **The High Court dismissed all three challenges on 15 January 2026, in [2026] EWHC 64, upholding the PSR's statutory power to impose price caps on cross-border interchange fees.** The PSR then decided, in MR22/2.9, not to proceed with an interim cap, preferring to move directly to a longer-term cap once it had developed a robust methodology, on which it consulted in MR22/2.8. As at the time of writing in August 2026, no cap on UK-EEA cross-border interchange had been imposed. HM Treasury's April 2026 consultation response states that the ability to impose price controls will carry over to the FCA on consolidation.

Running in parallel, and quite separately, is the private litigation. The Competition Appeal Tribunal's judgment of 27 June 2025 in the umbrella interchange proceedings, brought by over two thousand merchant claimants, held that the schemes' default interchange fee arrangements restricted competition in the acquiring market contrary to Article 101(1) of the Treaty on the Functioning of the European Union in respect of inter-regional and commercial card transactions. Permission to appeal was sought in August 2025 and the appellate proceedings were live at the time of writing. Interchange is therefore regulated on two tracks at once — a regulatory cap track and a competition law damages track — and they can reach different answers about the same fee.

■ Operational resilience

The FCA's operational resilience regime was made by policy statement **PS21/3, published in March 2021**, with the rules in SYSC 15A taking effect on **31 March 2022** and a transitional period ending on **31 March 2025**. Its scope expressly includes entities authorised and registered under the PSRs 2017 and the EMRs, alongside banks, building societies, insurers, enhanced-scope Senior Managers and Certification Regime firms and recognised investment exchanges.

The regime rests on four obligations. A firm must identify its **important business services**, defined as services provided to clients which, if disrupted, could cause intolerable levels of harm to one or more clients or pose a risk to the soundness, stability or resilience of the UK financial system. It must set an **impact tolerance** for each, being the maximum tolerable level of disruption, measured as a length of time and any other relevant metric. It must **map** the people, processes, technology, facilities and information supporting each service, **test** its ability to remain within tolerance under severe but plausible scenarios, and produce a board-approved **self-assessment**.

By 31 March 2025 firms were required to have completed mapping and testing to the sophistication necessary to demonstrate that they can remain within impact tolerances. The FCA's supervisory observations emphasise several things firms routinely get wrong: that a service should not be excluded from the list on the strength of a single factor; that identification must be done without reference to the firm's response and recovery capabilities, so a service is not made unimportant by being well protected; that interdependencies and sub-contracting chains must be considered; and, critically, that the firm remains responsible for staying within its impact tolerance even where a third party delivers the service. A firm may rely on a third party's own testing, but must satisfy itself that the methodology and scenarios are adequate.

Two further layers sit on top.

Incident and third-party reporting. Regulation 98 of the PSRs 2017 requires payment service providers to maintain a framework for managing operational and security risks with an annual assessment, and regulation 99 requires notification of major operational or security incidents without

undue delay. Following consultation CP24/28, published 13 December 2024 and closed 13 March 2025, the FCA published **policy statement PS26/2 on 18 March 2026**, with finalised guidance FG26/3 and FG26/4. PS26/2 creates a single regime shared across the FCA, the PRA and the Bank of England, and **its rules come into force on 18 March 2027**. It defines an operational incident, sets reporting thresholds, and introduces a standardised single submission regardless of which regulator the report is for, with a two-tier model of standard and enhanced reporting. The FCA threshold is that the firm reasonably believes the incident poses a risk of intolerable harm to consumers from which they cannot easily recover, or a risk to the safety and soundness of the firm or other market participants, or a risk to market stability, integrity or confidence in the UK financial system. Reportability is not limited to incidents affecting an important business service, and payment service providers will discharge their regulation 99 obligations through this regime. The third-party rules apply to a narrower population expressly including authorised e-money institutions and authorised payment institutions, requiring notification of new or significantly changed material third party arrangements and an annually submitted register.

Critical third parties. The Financial Services and Markets Act 2023 gave HM Treasury power to designate a third party providing services to regulated firms as a critical third party, and gave the regulators rule-making and enforcement powers over designated entities. The final rules, published jointly as FCA PS24/16 and PRA PS16/24 on 12 November 2024, **took effect on 1 January 2025**, but apply to a provider only from the date its designation order comes into force. **On 10 July 2026 HM Treasury designated the first critical third parties, with effect from 13 July 2026: Microsoft Ireland Operations Limited, Google Cloud EMEA Limited, Amazon Web Services EMEA SARL and Oracle Corporation UK Limited.** Designation does not reduce the responsibility of the firms relying on those providers; it adds a supervisory relationship with the provider itself.

■ The European Union position, and where the two regimes now differ

The Second Payment Services Directive, Directive (EU) 2015/2366, applied from 13 January 2018, with the regulatory technical standards on strong customer authentication and common and secure communication, Delegated Regulation (EU) 2018/389, applying from 14 September 2019. The Second Electronic Money Directive, 2009/110/EC, sits alongside it. The European Banking Authority issues guidelines and technical standards; national competent authorities supervise.

On 28 June 2023 the European Commission published two proposals: **PSD3**, COM(2023)366, a directive covering licensing, prudential requirements and supervision, and the **Payment Services Regulation**, COM(2023)367, a directly applicable regulation covering conduct, transparency and rights and obligations. Splitting the regime this way is the point: conduct rules become uniform without national transposition, while authorisation remains a directive.

The legislative position, as at the time of writing in August 2026: Parliament adopted its first-reading position on 23 April 2024; Parliament and Council reached provisional political agreement on 27 November 2025; and the text agreed at early second-reading negotiations was approved in the ECON committee on 5 May 2026. The European Parliament's legislative train summary, updated 20 June 2026, records that the deal still needs formal adoption by Parliament and Council before it can enter into force, and the Parliament's own Legislative Observatory still shows the file as awaiting the Council's first-reading position, with an indicative plenary sitting date of 14 December 2026. Nothing has therefore reached the Official Journal, and the practitioner expectation of publication in the first half of 2026 has already been overtaken. The application clock itself is not speculation: the final compromise texts published in the Council's public document register, 8221/26 for the Regulation and 8222/26 for the Directive, both dated 17 April 2026, provide that the Regulation applies twenty-one months after

its entry into force, that member states must adopt, publish and apply their transposing measures within the same twenty-one months, and that the payee verification provisions in Articles 50 and 57 of the Regulation apply only from twenty-seven months. Entry into force is twenty days after Official Journal publication, so the whole calendar can be counted precisely from the day that publication happens, and not a day before.

The substantive changes that matter most, as reported in the agreed texts: e-money institutions cease to exist as a separate category and become a sub-category of payment institutions, with e-money issuance treated as a payment service in the annex to PSD3 and the Second Electronic Money Directive repealed; payment institutions gain an option, at national discretion, to safeguard customer funds at a central bank; providers must offer a payee verification service and inform the payer of discrepancies; liability extends to failures to implement appropriate fraud prevention mechanisms, and to impersonation fraud where the customer has reported the matter to the police and the provider; online platforms that fail to remove fraudulent financial content after notification become liable to providers that have reimbursed defrauded customers; advertisers of financial services must demonstrate to very large online platforms and search engines that they are authorised or exempt; and users gain a right of access to human customer support.

Separately, and already in force, the **Instant Payments Regulation (EU) 2024/886**, adopted on 13 March 2024, amends the SEPA Regulation to require euro instant credit transfers, charge parity with ordinary transfers, a free verification of payee service, and at-least-daily sanctions screening of customers rather than transaction-by-transaction screening. Its deadlines, as published by the European Central Bank, are staggered:

Requirement	Applies to	Deadline
Receiving instant payments, and equality of charges	Euro area member states	9 January 2025
Sending instant payments, and verification of payee	Euro area member states	9 October 2025
Receiving instant payments, and equality of charges	Non-euro area member states	9 January 2027
Receiving instant payments	EMIs and PIs, euro and non-euro area	9 April 2027
Sending instant payments	EMIs and PIs in euro area member states	9 April 2027
Sending instant payments, and verification of payee	Non-euro area member states, and their EMIs and PIs	9 July 2027
Sending outside business hours from national currency accounts	Non-euro area member states	9 June 2028

The same regulation, through amendments to PSD2 and the Settlement Finality Directive, makes payment institutions and e-money institutions eligible to participate directly in designated payment systems. Since October 2025 non-bank payment service providers meeting the TARGET Guideline requirements have been able to access TARGET services including T2 and TIPS. That is a structural change of the first order: it removes the requirement for a non-bank to hold its settlement position at a commercial bank that is also its competitor.

On operational resilience, the European regime is the **Digital Operational Resilience Act**, Regulation (EU) 2022/2554, which has applied since **17 January 2025**. It covers ICT risk management, incident classification and reporting, resilience testing including threat-led penetration testing for significant

entities, ICT third-party risk management with prescribed contractual terms, and direct oversight of critical ICT third-party providers by the European Supervisory Authorities. Firms operating on both sides will find the UK and EU regimes conceptually similar and mechanically incompatible: impact tolerances and important business services on one side, ICT risk management and register-of-information reporting on the other, with different thresholds and different clocks.

■ The United Kingdom's own moving parts

Three UK changes bear directly on how a payments firm operates, and each has an explicit date.

The **Payment Services (Amendment) Regulations 2024** came into force on **30 October 2024**, permitting a payment service provider to delay execution of an outbound sterling payment within the UK where it has reasonable grounds to suspect fraud or dishonesty and needs to make further enquiries. The delay must be no longer than necessary and in any event no longer than the end of the fourth business day from receipt of the payment order. The provider must tell the customer of the delay and its reasons unless disclosure would be unlawful, and is liable for any charges or interest the customer incurs.

The **authorised push payment reimbursement requirement** has applied since **7 October 2024**, imposed by the PSR through Specific Direction 20 for Faster Payments and mirrored by the Bank of England for CHAPS. The **maximum reimbursement is £85,000 per claim**, confirmed in PSR policy statement PS24/7 in October 2024, which reduced a previously confirmed maximum of £415,000 before go-live. The PSR estimated that £85,000 fully covers 99.8% of Faster Payments authorised push payment scams by volume and 90% by value.

The **contactless rule changes announced by the FCA in December 2025** took effect on **19 March 2026**, removing the mandatory £100 single-transaction limit and the cumulative limit and giving firms with strong fraud controls the flexibility to set their own. Nothing obliges a firm to change anything, and terminal and scheme-rule changes are also needed before higher-value contactless card transactions can be processed. The pattern will recur: a fixed numerical rule replaced by a risk-based standard, with the supervisory question shifting from compliance with a threshold to the demonstrable quality of a firm's own fraud monitoring.

■ What this means for a firm actually operating

The standing obligations of a mid-sized authorised payment institution, as at August 2026, run wider than most people building a payments product expect: own funds on the applicable method; daily safeguarding reconciliation and the monthly return; the annual safeguarding audit; the regulation 98 risk framework and its annual assessment; the SYSC 15A self-assessment, mapping and testing; professional indemnity insurance for payment initiation or account information services; the Money Laundering Regulations 2017; the Consumer Duty in PRIN 2A where it applies; complaints under DISP; strong customer authentication; the reporting cycle under SUP 16.13; registration of agents before they act; and, from 18 March 2027, notification of material third party arrangements and the register.

The single most common cause of serious regulatory difficulty in this sector is not any of those in isolation. It is the gap between what a firm's safeguarding documentation says and what its bank accounts actually contain, discovered late. Everything the FCA has done since CP24/20 shortens the interval at which a discrepancy becomes visible: from an annual audit, to a monthly return, to a daily reconciliation. A firm that treats the daily reconciliation as a compliance chore has misunderstood it. It is the early warning system for the only failure mode that genuinely destroys customers.

■ What is most likely to be wrong in this chapter first

Three things, and it is better to name them than to pretend otherwise.

The abolition of the Payment Systems Regulator is government policy with a published consultation response but no Act of Parliament. When the legislation passes, every reference here to the PSR as a separate body becomes historical, and the statutory basis for interchange regulation and for the reimbursement rules will move.

The European Union's PSD3 and Payment Services Regulation package had reached committee approval of an agreed text but, on the author's checking in August 2026, had not been verified as adopted and published in the Official Journal. Once it applies, e-money institutions cease to be a separate category in the European Union while remaining one in the United Kingdom: the largest structural divergence between the two regimes since departure.

And the whole of the United Kingdom's assimilated payment services law is subject to repeal under the Financial Services and Markets Act 2023, with firm-facing requirements expected to move into FCA rules. When that happens, the regulation numbers used throughout this chapter — regulation 23 for safeguarding, regulation 98 for security risk, regulation 138 for the criminal offence, Schedule 3 for capital — will cease to exist even where the substance survives in renumbered form. Every figure and deadline here should be checked against the source before it is relied on, and the sources are named below precisely so that it can be.

53.98 Common wrong ideas

Wrong: Money held by a payment institution or e-money institution is protected up to £85,000 like a bank deposit. **Right:** It is not a deposit and the Financial Services Compensation Scheme does not cover it; safeguarding creates a pool meant to be paid to customers ahead of ordinary creditors, and whether that pool is full is a question of fact.

Wrong: The £85,000 figure in British payments is the safeguarding guarantee. **Right:** It is the maximum reimbursement per claim under the mandatory authorised push payment scam rules confirmed in PSR policy statement PS24/7, applying since 7 October 2024 — the same number attached to a different mechanism, with no connection between them.

Wrong: Safeguarded funds are held on trust from the moment they are received. **Right:** *Re Ipagoo LLP* held that the Electronic Money Regulations impose no statutory trust; they create a statutory asset pool topped up from the general estate, so customers rank ahead of ordinary creditors but without the proprietary certainty a trust would give.

Wrong: There is a payments regulator, and complying with its rules is what keeps you trading. **Right:** The FCA, the PSR, the Bank of England, the PRA and HM Treasury all hold pieces of the perimeter, and the private rulebooks — the Visa and Mastercard operating rules, Pay.UK's scheme rules, the PCI Data Security Standard — bind harder than much legislation; a firm perfectly compliant with the Payment Services Regulations can be shut down in a fortnight by a card scheme.

Wrong: PSD2 applies in the United Kingdom. **Right:** PSD2 was a directive instructing member states to legislate; the United Kingdom's legislation was the Payment Services Regulations 2017 plus FCA rules, now assimilated law that Parliament and the FCA can and do change unilaterally, as the removal of the mandatory £100 contactless limit on 19 March 2026 demonstrates.

Wrong: Capping interchange caps what a merchant pays to accept cards. **Right:** Interchange is one component of the merchant service charge alongside uncapped scheme fees and the acquirer's own margin, and capping it has never stopped the total cost of acceptance rising.

Wrong: The interchange caps apply to all cards. **Right:** They apply to consumer debit and credit cards on transactions within the relevant territory; commercial cards, foreign-issued cards and three-party schemes such as American Express fall outside them by design, which is why the most expensive card in a merchant's mix is usually one on which no cap bites at all.

Wrong: The FCA has three months to decide an authorisation application, so budget three months. **Right:** The clock runs from completeness rather than submission, an incomplete application attracts a twelve-month deadline, and most first drafts are not complete.

Wrong: Outsourcing an important business service moves the resilience obligation to the supplier. **Right:** The firm remains responsible for staying within its own impact tolerance even where a third party delivers the service, and the designation of a critical third party adds supervision of the provider without reducing anybody else's responsibility.

Wrong: The daily safeguarding reconciliation is a compliance chore. **Right:** It is the early warning system for the only failure mode that genuinely destroys customers, which is why the FCA has shortened the interval at which a discrepancy becomes visible from an annual audit to a monthly return to a daily count.

53.99 Chapter summary in 20 lines

1. Payments regulation is built around one failure in particular: a firm losing the customer money it was only holding in transit, which no amount of the firm's own capital would have prevented.
2. Moving other people's money as a business requires a licence from the Financial Conduct Authority, and the licence comes in sizes.
3. Five categories matter, distinguished by scale, by whether the firm issues e-money, and by which of the Schedule 1 payment services it provides.
4. Initial capital ranges from nothing for an account information service provider to EUR 350,000 for an authorised e-money institution, and the amounts are still stated in euros in British legislation because they were carried across from the European originals without conversion.
5. Initial capital is a floor from day one, whereas ongoing own funds is a continuing calculation under Method A, B or C as the FCA directs.
6. Providing payment services without authorisation or registration is a criminal offence under regulation 138 of the Payment Services Regulations 2017, not a technical breach.
7. The statutory three-month determination period runs from the point an application is complete, which is why published timelines bear so little relation to it.
8. Safeguarding carries far more regulatory weight than capital, because the question these rules ask is a gone-concern question: when this firm fails, do its customers get their money back, in full, quickly?
9. The honest answer historically was no — an average shortfall of 65 per cent across twelve insolvencies, 80 per cent for e-money institutions alone, and 2.3 years to distribution.
10. *Re Ipago LLP* established that the regulations create a statutory asset pool rather than a trust, which is why customers rank ahead of ordinary creditors without proprietary certainty.
11. Consultation CP24/20 and policy statement PS25/12 followed, and the Supplementary Regime in CASS 15 came into force on 7 May 2026 with daily internal and external reconciliations, a monthly return and an annual audit subject to a £100,000 exemption.

12. Two proposals were held back for further work: receiving relevant funds directly into a designated safeguarding account, and the express statutory trust of the Post-Repeal Regime.
13. Operational resilience under SYSC 15A asks a firm to identify its important business services, set an impact tolerance in time for each, map the dependencies, test against severe but plausible scenarios and self-assess at board level.
14. A service is not made unimportant by being well protected, and the firm stays responsible for its tolerance even where a third party delivers the service.
15. A single incident and third-party reporting regime under PS26/2 comes into force on 18 March 2027, and HM Treasury designated the first critical third parties with effect from 13 July 2026.
16. Interchange on UK consumer debit is capped at 0.2% of value and on consumer credit at 0.3% under the assimilated Regulation (EU) 2015/751, with a Treasury power to lower but not to raise.
17. That cap reaches only one component of one part of the market, which is why total card acceptance costs have kept rising and why the priciest cards escape it entirely.
18. Interchange is contested on two tracks at once — the PSR's cross-border market review, whose capping power the High Court upheld on 15 January 2026, and the competition damages litigation in the Competition Appeal Tribunal — and the two can reach different answers about the same fee.
19. The institutional map is being redrawn: the government confirmed on 21 April 2026 that the PSR's functions will transfer to the FCA, but consolidation needs an Act of Parliament that had not been introduced by August 2026.
20. The whole body of assimilated payment services law is subject to repeal under the Financial Services and Markets Act 2023, so the regulation numbers used here may survive only in renumbered form, and every figure and date should be checked against the named source before it is relied on.

Sources: the Payment Services Regulations 2017, the Electronic Money Regulations 2011, the Payment and Electronic Money Institution Insolvency Regulations 2021, the Payment Card Interchange Fee Regulations 2015 and assimilated Regulation (EU) 2015/751, all via legislation.gov.uk; FCA CP24/20, PS25/12, CP24/28, PS26/2, PS24/16, PS21/3, the December 2025 contactless statement and the Approach Document Payment Services and Electronic Money — Our Approach; PRA and Bank of England PS16/24; PSR MR22/2.7, MR22/2.8, MR22/2.9 and PS24/7; HM Treasury's National Payments Vision, its April 2026 consultation response A Streamlined Approach to Payment Systems Regulation and its July 2026 critical third party designations; Re Ipagoo LLP [2022] EWCA Civ 302 and the High Court judgment of 15 January 2026; the European Parliament legislative train file on the payment services regulation; and the European Central Bank's Instant Payments Regulation implementation table.

What Breaks and What It Costs

54.0 What this chapter gives you

1. You will be able to explain why a forty-minute terminal outage produces three weeks of consequences, and why the two numbers are not related to each other.
2. You will be able to do the arithmetic that turns 12,400 uncertain items into roughly 413 hours of manual research, and give senior management a recovery estimate they will not like but cannot argue with.
3. You will be able to explain why a partial batch is more dangerous than a total failure, and why resubmitting one converts a partial failure into a duplicate.
4. You will be able to name the five phases of an incident and say which of them firms are habitually bad at.
5. You will be able to list the clocks that start at once in a United Kingdom incident — four hours from first detection, seventy-two hours for a personal data breach, the end of the next business day for an unauthorised transaction, five business days for a scam claim — and say what triggers each.
6. You will be able to explain why the regulators fined TSB for the governance around the decision to go live rather than for the fault itself.
7. You will be able to say which rails can be pulled back and which cannot, and explain that settlement finality is a deliberately created legal state rather than a property of a message.
8. You will be able to explain why batch file submission has no idempotency key, and why duplicate protection is therefore a control the submitter must build rather than a service the scheme provides.
9. You will be able to distinguish operational resilience, the critical third parties regime and incident reporting, and explain why the three are one policy.
10. You will be able to specify a control that alerts on the absence of an expected file, which is the cheapest fix for the most expensive silent failure in the industry.

Every chapter of this book so far has described something working. Cards authorising, files clearing, settlement landing, mandates honoured, disputes adjudicated. That is the right way round, because you cannot understand a failure in a system you do not understand working. But it leaves an impression that needs correcting before we finish, which is that payments is a set of mechanisms that occasionally break.

It is not. Payments is a set of mechanisms that break continually at a low rate, and the entire apparatus around them — the reconciliation, the exception queues, the indemnities, the redress rules, the regulators, the two o'clock in the morning phone trees — exists because that is the normal condition rather than the exceptional one. A payment system is not a machine that works and sometimes fails. It is a machine that fails at a known rate, wrapped in enough correction to make the failures survivable.

The published record is unambiguous about the rate. In March 2025 the House of Commons Treasury Committee published correspondence from nine of the largest banks and building societies operating in the United Kingdom, and the aggregate was at least 803 hours of unplanned technology and systems outage — more than thirty-three days — across at least 158 separate incidents between January 2023 and February 2025. Nine well-capitalised, heavily supervised institutions, in a mature market, with mature controls. That number is not a scandal. It is the baseline.

This chapter is about what happens inside those hours and, much more importantly, in the weeks after them. The thing practitioners learn early and outsiders almost never see is that the outage is the small part. What outlasts it is the reconciliation, and the reconciliation is where the money actually is.

The plain version

Start with a water main.

A pipe bursts under a road at seven in the morning. The water company gets a crew there by nine and the leak is stopped by half past ten. Three and a half hours, and the problem in the technical sense is over. Nobody thinks the story ends there. The road is dug up for a month. The trench has to be backfilled, the tarmac replaced, the two shops whose cellars flooded dealt with, the claims from people whose cars were damaged assessed one at a time, and somebody has to work out whether the pipe that burst was the one that was supposed to have been replaced in 2019.

Payments incidents are exactly like this, and people consistently misjudge them for exactly the same reason. The visible event — the app is down, the card is declined, the wages have not arrived — is the leak. Fixing the leak is a few hours of engineering. The road is the ledger, and the ledger takes weeks.

Here is a real-shaped example with real arithmetic.

Riverside Court is a care home. It employs eighty-four people and pays them monthly, and the total that leaves its bank account on payday is £176,400, an average of £2,100 each. The way it pays them is old and reliable: on the Wednesday, the payroll clerk sends a list to the bank containing eighty-four lines, each line saying who to pay, how much, and which account. Three working days later the money lands.

One month, the confirmation screen does not load. The clerk waits, refreshes, sees nothing, and sensibly assumes the list did not go. So she sends it again. Both lists were correctly signed, both were validly formatted, and the bank has no way of knowing that the second one is not a genuine second run of wages. On the Friday, £352,800 leaves the account. Every single member of staff is paid twice.

Now watch how long each part takes.

The cause takes about four seconds to explain and about four minutes to fix: the software gains a check that warns you when you are about to send a list identical to one sent in the last twenty-four hours. That is the whole technical remedy. It could be written on a Monday morning.

Detection takes three days, and this is the first surprise. Nobody notices on Friday. The staff certainly do not ring up to complain about being paid twice. The bank does not stop it, because nothing about the second list was invalid. The alarm is not an alarm at all: on Monday the finance manager opens the bank statement for an unrelated reason, sees a balance £176,400 lower than she expected, and has to work backwards from a number to a cause. In payments, this is the ordinary way things are found. Not a red light. A figure that is wrong.

Then comes the part that takes eleven weeks.

The money has gone into eighty-four accounts belonging to eighty-four adults, and it is now legally theirs to argue about. Riverside writes to all of them on the Tuesday. Sixty send it back within a fortnight: £126,000 recovered. Fifteen more come back over the following six weeks, some awkwardly, because they had already spent part of it — one had cleared a credit card, one had paid a deposit on a holiday. That is another £31,500, some of it in instalments. Six agree to repay over four months. Three never pay at all: one has left the country, one disputes being overpaid, one simply stops answering. £6,300 is written off.

Meanwhile the books do not balance, in a specific and irritating way. The payroll system recorded one payroll. The bank recorded two payments to each person. So for eleven weeks there is an unexplained hole that starts at £176,400 and shrinks unevenly as money dribbles back, and every returned payment has to be matched to the right person by hand, because people send money back with references like “wages” and “sorry”. The cash flow forecast was wrong for two months. The overdraft was used, costing £940 in interest and fees. The auditor asks about it the following March, so somebody writes the whole thing up again seven months later.

Add it up. Four minutes of code. Roughly sixty-five hours of finance and payroll time at Riverside. £6,300 written off, £940 of borrowing costs, one very uncomfortable staff meeting, and a company that will be nervous about payday for a year.

The same shape appears when nothing goes out at all. A café’s card terminal stops taking payments at half past twelve on a Saturday and is fixed in forty minutes, and the café loses perhaps £600 of lunchtime trade that never comes back. But some of the failed taps were not clean failures. A handful of customers were charged, saw a decline on the terminal, and paid again with a different card. Sorting out fourteen double charges — finding them, matching them to receipts, refunding them, answering the two customers who go to their banks instead and start a formal dispute — takes the owner three weeks of intermittent effort and one furious online review.

Forty minutes of outage, three weeks of consequences. If you remember one sentence from this chapter, make it that one: the outage is measured in hours, the recovery is measured in weeks, and the two numbers are not related.

One last thing the plain version needs to establish, because it is why all of this is so laborious. When money moves and settles, it is genuinely gone. It is not sitting somewhere temporary waiting to be recalled. Somebody else now owns it, and getting it back requires either their agreement or a court. That is not a design flaw. It is the entire point of a payment: to end the argument about who owes what. A payment you could quietly reverse would not be a payment, it would be a loan with extra steps. The price of finality — the thing that lets the shopkeeper hand over the goods — is that a mistake is expensive to undo.

Where the plain version stops being true

The clock that matters does not start when you understand the problem. The Riverside story has detection on the Monday and everything flowing from there. In a regulated firm, the regulatory clock is deliberately not tied to your understanding. At the time of writing in August 2026, a UK payment service provider that becomes aware of a major operational or security incident must notify the Financial Conduct Authority without undue delay under regulation 99 of the Payment Services Regulations 2017, with the timing currently specified through the European Banking Authority’s Guidelines on incident reporting under PSD2, EBA/GL/2017/10 as issued on 27 July 2017, applied by SUP 15.14 of the FCA Handbook. That guidance requires an initial report within four hours of first detection. Not from root cause, not from confirmed impact, not from the incident bridge agreeing what to call it.

Firms must therefore be able to file a report saying, in effect, something is wrong, we do not yet know what, here is what we do know. The instinct to wait until you can give a clean answer is the instinct that produces a breach.

Duplicate is easier to explain than partial, and partial is far more dangerous. The Riverside example is a duplicate, which is tidy: you know exactly what happened, twice, to exactly whom. Real incidents are more often partial, and partial is worse in a way that is not intuitive. If a batch of 40,000 payments fails completely, you know the state of the world: nothing happened, resend. If 31,600 succeeded, 2,900 failed cleanly, and 5,500 returned no answer at all, you do not know the state of the world and cannot find out from your own records. You have to reconcile against somebody else's. Until you have, you cannot safely resubmit, because resubmitting the ambiguous 5,500 might duplicate them and not resubmitting them might leave 5,500 people unpaid. A clean total failure is a good day. Ambiguity is the expensive one.

The cost is almost never the outage. Riverside's direct loss was small. The pattern in published cases is that the operational loss is a rounding error against the redress, the remediation, the enforcement and the supervisory tail. TSB's platform migration in April 2018 gives the clearest published ratio: the Financial Conduct Authority and Prudential Regulation Authority stated when fining the bank a combined £48,650,000 on 20 December 2022 that all of TSB's branches and a significant proportion of its then 5.2 million customers were affected by the initial issues, that the problems were not fully resolved for eight months, and that TSB had by then paid £32.7 million in redress. Against that, the engineering hours to restore service were trivial. The asymmetry runs the other way too: the Bank of England's Real-Time Gross Settlement system was down for roughly nine hours on 20 October 2014, delaying 142,759 CHAPS payments worth £289.3 billion, and the Bank's published response of 25 March 2015 records that it settled nine compensation claims totalling £4,056.89. A national settlement outage produced about four thousand pounds of claims; a retail migration produced tens of millions. Where the harm lands, not how big the system is, determines the bill.

"The technical fix is the short part" has a real exception, and it is the migration case. When the failure is a design decision rather than a fault — a new platform that does not behave as the old one did, a data model that lost information in transit, a capacity assumption wrong by an order of magnitude — there is no short fix, because there is nothing to restore to. TSB's disruption ran for eight months. Deloitte's review found the root cause of the 2014 RTGS outage to be defects introduced in functionality changes made in April 2013 and May 2014, meaning the fault had been latent for around eighteen months before it expressed itself. Incidents caused by a change you deliberately made are a different species from incidents caused by a component that failed, and they behave differently in every dimension: longer, more expensive, more likely to attract enforcement, and much more likely to be attributed to governance rather than technology.

The technical version

■ On dating

Every figure, deadline, threshold and rule in this section is stated as accurate at the time of writing, in August 2026. Operational resilience is one of the fastest-moving areas of financial regulation in the United Kingdom and the European Union, and three specific things were in motion as this was written: the FCA's and PRA's new operational incident and third-party reporting regime had been finalised but did not come into force until 18 March 2027; HM Treasury had made its first designations under the critical third parties regime only weeks earlier; and the EU's Digital Operational Resilience Act had been applicable for less than two years, with its supervisory practice still forming. Verify every number below against the primary source before relying on it.

■ The anatomy of an incident

An incident has five phases, and firms tend to be good at the second and third and bad at the first, fourth and fifth.

Detection. The commonest detection mechanism in payments is not an error alert. It is a volume anomaly: throughput that has departed from its expected shape for the time of day, the day of the week and the point in the month. This matters because the failure modes that hurt most produce no errors at all. A file that is never picked up throws nothing. A downstream system that silently truncates a field returns success. A duplicate submission is, by construction, two valid events. Error-rate monitoring finds the loud failures; throughput and reconciliation monitoring find the quiet ones, and the quiet ones cost more. Mature operations therefore alert on the absence of expected traffic — no Bacs input report by 23:00, no settlement file by 06:00, credit volumes 40 per cent below the same weekday last week — rather than only on the presence of errors.

Classification. Classification is a regulatory act, not an engineering one, because it starts clocks. The problem is that it demands an impact estimate at the moment when impact is least knowable. The discipline that works is to classify on plausible worst case and revise downwards later, because every reporting regime tolerates an overstated report far better than a late one.

Containment. A genuine tension runs through every serious incident: stopping the bleeding usually destroys evidence. Failing over destroys the state of the failed node. Restarting a stuck process discards the queue you needed to count. Reprocessing a file overwrites the log that told you what the first attempt did. Good runbooks resolve this in advance by specifying what must be captured before each containment action, so the decision is not being made by a tired engineer at four in the morning who will later be asked, by a skilled person under section 166 of the Financial Services and Markets Act 2000, why there is no record of what happened.

Recovery. Restoring service is not the same as restoring correctness, and conflating them is the classic mistake. Service is restored when payments flow again. Correctness is restored when every payment that should have happened has happened exactly once, and every payment that should not have happened has been reversed or paid for. The gap between those two moments is where the reconciliation lives: typically two to eight weeks for a material incident.

Aftermath. Redress, reporting, root cause analysis, remediation, and the supervisory conversation, which continues long after the engineering is closed.

Command structure is where competent firms differ visibly from incompetent ones. There is one incident commander, who does not fix anything. There is a separate communications lead, because the person diagnosing must not also be the person drafting the customer statement. There is a scribe

keeping a timestamped log, because the log is the evidence and cannot be reconstructed afterwards. And there is a named accountable senior manager: in the United Kingdom, operational resilience typically sits with the Chief Operations function, SMF24, under the Senior Managers and Certification Regime, which means the failure has a person's name attached to it in a way that predates the incident.

■ The clocks

The single hardest operational fact about a modern incident is that several unrelated deadlines start at once, they are triggered by different events, and they are owed to different people.

Obligation	Trigger	Deadline	Basis, as at August 2026
UK major payments incident, initial	First detection	4 hours	PSRs 2017 reg. 99, SUP 15.14, EBA/GL/2017/10
UK personal data breach	Awareness	72 hours	UK GDPR Art. 33
EU major ICT incident, initial	Classification as major	4 hours, and no later than 24 hours from becoming aware	DORA Art. 19(4)(a)
EU major ICT incident, intermediate	Initial notification	72 hours	DORA Art. 19(4)(b)
EU major ICT incident, final	Intermediate report	1 month	DORA Art. 19(4)(c)
Refund of unauthorised transaction	Notification by payer	End of the next business day	PSRs 2017 reg. 76
APP scam reimbursement	Claim	5 business days, subject to stop the clock	PSR reimbursement requirement, from 7 Oct 2024
UK incident report, from 18 Mar 2027	Threshold determination	Ordinarily 24 hours, but 4 hours from first detection for PSPs	FCA PS26/2, FG26/3; PRA PS7/26, SS1/26

Two of those deserve elaboration.

DORA — Regulation (EU) 2022/2554, applicable since 17 January 2025 — sets its initial four-hour clock from classification rather than detection, but caps it by requiring notification no later than twenty-four hours from becoming aware of the incident. Classification itself must be done without undue delay, against the criteria in the Commission Delegated Regulation on classification of major ICT-related incidents. The design intent is to stop firms deferring the clock by deferring the decision.

The United Kingdom is replacing its arrangements. On 18 March 2026 the FCA published PS26/2, with finalised guidance FG26/3 on operational incident reporting and FG26/4 on material third-party reporting, alongside the PRA's PS7/26 and supervisory statement SS1/26. The rules come into force on 18 March 2027. They create a single definition of an operational incident, a single template and a single submission route through FCA Connect for firms that today report separately to two or three authorities. An operational incident is a single event or a series of linked events which disrupt the firm's operations such that they disrupt delivery of a service to an end user external to the firm, or affect the availability, integrity, authenticity or confidentiality of that end user's data. Reporting is split into standard and enhanced tiers, with payment service providers in the enhanced tier and, crucially, retaining the accelerated four-hour deadline from first detection while other enhanced firms work to twenty-four hours from threshold determination. The EBA guidelines will be disapplied for UK PSPs,

who will discharge regulation 99 solely through the new regime. Final reports are due within thirty working days of resolution, or no later than sixty with reasons for the delay.

■ Failure modes, and why some are worse than others

Duplicate submission and duplicate settlement. The commonest serious payments incident, and the one with the worst recovery profile, because a duplicate is by definition a valid instruction. The largest published UK example is Santander's Christmas Day 2021 incident, in which approximately £130 million was paid out in around 75,000 duplicate transactions originating from about 2,000 business accounts. Recovering a duplicate credit means recovering funds from third parties who have no contractual relationship with the party that made the mistake, which is why the money often does not come back.

Stuck file. The instruction was created but never delivered, or delivered and never picked up. Zero errors, zero alarms, zero payments. Only a positive control catches this: a reconciliation that expects a specific file, of roughly a specific size, by a specific time, and alerts on its absence.

Partial batch. Some items processed, some did not, and the boundary is unknown. This is the failure mode that most often produces the wrong recovery action, because the temptation to resubmit the whole file is enormous and doing so converts a partial failure into a duplicate.

Silent corruption. Field-level truncation and encoding faults are especially nasty in the file-based world because the legacy formats are fixed-width and unforgiving. In the Bacs Standard 18 format, the fields carrying the counterparty's name and the payment reference are eighteen characters. Anything that overruns is cut, and a cut reference is a payment that arrives but cannot be applied. The sender sees success. The receiver sees an unidentifiable credit and a reconciliation break.

Idempotency, and why batch payments lack it. Modern payment APIs are built around an idempotency key: a client-generated identifier that lets the server recognise a retry and return the original result rather than performing the action twice. That solves the Riverside problem completely. Batch file submission has no equivalent primitive. A Bacs submission is authenticated and non-repudiable, and is therefore trusted; it is not deduplicated, because two genuine payroll runs to the same eighty-four people for the same amounts on consecutive days is a legitimate business event no scheme can distinguish from an error. Duplicate protection in the file world is a control submitters must build themselves.

■ Reversal asymmetry

Whether a payment can be pulled back is not a question about technology. It is a question about which rail it went through, and the answers differ sharply.

Rail	Reversal mechanism	Practical character, as at August 2026
Bacs Direct Credit	No unilateral reversal; recall by request to the receiving PSP, requiring the beneficiary's consent	Weak. Success depends on goodwill and on the money still being there
Bacs Direct Debit	Direct Debit Indemnity Claim under the Direct Debit Guarantee	Strong for the payer. The paying PSP refunds first and recovers from the originator afterwards, with a challenge process available to the service user

Rail	Reversal mechanism	Practical character, as at August 2026
Faster Payments	Recall request; no guaranteed return	Weak, and fast money is fast in both directions
CHAPS	Irrevocable between direct participants once settled	None, absent the beneficiary's agreement
Card	Chargeback under scheme rules, within scheme-defined windows	Strong for the cardholder, structured, and adversarial

Underneath all of this sits settlement finality. The Financial Markets and Insolvency (Settlement Finality) Regulations 1999, S.I. 1999/2979, allow payment systems to be designated so that transfer orders effected through them are protected from being unwound by insolvency law. CHAPS was designated in 2000 and is a recognised interbank payment system under Part 5 of the Banking Act 2009. Finality is not a technical property of a message; it is a legal state, deliberately created, that exists precisely so that a settled payment cannot be reopened by later events.

The consequences of that design are best shown by the largest published misdirected-payment case of recent years. On 11 August 2020, Citibank, acting as administrative agent on a Revlon term loan, sent lenders the full outstanding principal when it had intended to send only an interest payment. Some returned the money. Others did not. The sum in dispute — principal and accrued interest owed to the lenders who declined to return it — was \$558,558,375.74. The United States District Court for the Southern District of New York held in February 2021 that the discharge-for-value rule gave those lenders a defence against restitution, because they were owed exactly what they received and had no notice of the mistake. The Second Circuit vacated that ruling on 8 September 2022 in *In re Citibank August 11, 2020 Wire Transfers*, No. 21-487, holding among other things that the debt was not then due and that the recipients had constructive notice. The point for a practitioner is not the eventual outcome. It is that recovering an erroneous payment of over half a billion dollars from sophisticated counterparties took two years and two courts, and at the halfway mark the money was gone. Payment systems are built to make payments stick. They succeed.

■ The reconciliation aftermath

This is the part of an incident that no press release describes and every practitioner remembers.

Reconciliation after an incident is a three-way problem, not a two-way one. There is the scheme's or acquirer's view, expressed in settlement and clearing files; the bank's view, expressed as movements on the settlement or client account; and the firm's own ledger, its record of what it believed it did. In normal operation these three agree to within a small and stable break rate. After an incident they do not, and the work is to reduce the disagreement to zero, item by item, with each item explained.

Breaks fall into recognisable categories, and the category determines the remedy:

Break type	Meaning	Typical remedy
In the ledger, not at the bank	We think we paid; no money moved	Reissue after confirming non-payment
At the bank, not in the ledger	Money moved; we have no record of instructing it	Investigate as potential duplicate or unauthorised
Matched, wrong amount	Both agree it happened, disagree on value	Adjust, then find the truncation or rounding fault

Break type	Meaning	Typical remedy
Matched, wrong date	Value-dating disagreement	Usually interest and reporting impact only
Unidentified receipt	Money arrived, cannot be applied	Suspense, then manual research, then escalating attempts to contact
Duplicate	Two records, one intended event	Recall attempt, then recovery, then write-off

Unmatched items go to a suspense account, which is not a solution but a holding pen with a clock on it. Two disciplines matter: ageing, because a two-day break is operations and a ninety-day break is a provision, and named ownership, because unowned breaks are immortal.

The reason this outlasts the outage is arithmetic. Suppose an incident leaves 12,400 items in an uncertain state. Automated matching on amount, date, counterparty and reference will typically resolve 85 to 95 per cent of them, so call it 1,240 items left. Manual research on a payments break — pulling the original instruction, the scheme record and the bank movement, and forming a view — averages fifteen to twenty-five minutes when straightforward and runs to hours when it involves contacting a beneficiary. At twenty minutes each, 1,240 items is roughly 413 hours: eleven weeks for one full-time person, or three weeks for a team of four who also have their normal work to do. Then there is the tail, the items requiring a third party's cooperation, which arrive back over months, and the small residue that never resolves and becomes a write-off with an audit trail.

That is why a forty-minute outage produces a six-week recovery, and why the honest answer to “when will this be finished?” on the day of an incident is a number that senior management will not like.

■ Customer impact, redress and the legal floor

Redress in the United Kingdom is not discretionary generosity. Several distinct regimes apply at once, and they attach to different failures.

For unauthorised transactions, regulation 76 of the Payment Services Regulations 2017 requires the payer's payment service provider to refund the amount as soon as practicable and in any event no later than the end of the business day following the day on which it becomes aware of, or is notified of, the transaction, and to restore the account to the state it would have been in. There are exceptions where the provider has reasonable grounds to suspect fraud and reports them appropriately. Regulations 91 to 95 deal with non-execution, defective execution and late execution, allocating liability between the payer's and payee's providers and requiring the payment to be traced.

For authorised push payment scams, the Payment Systems Regulator's reimbursement requirement has applied to eligible Faster Payments claims since 7 October 2024, with the Bank of England setting the equivalent for CHAPS. Reimbursement is due within five business days, subject to a stop-the-clock provision where the sending provider needs more information, and the maximum mandatory level is £85,000 per claim, a figure the PSR settled on having previously proposed £415,000. Firms may reimburse more, and losses above it can be taken to the Financial Ombudsman Service.

For everything else — the distress and inconvenience of an outage, the consequential losses, the missed mortgage completion, the declined card at the till — there is complaint handling under DISP in the FCA Handbook, the Financial Ombudsman Service, and since July 2023 the Consumer Duty, which requires firms to act to deliver good outcomes for retail customers and specifically covers consumer

support. At the time of writing, the Ombudsman's award limits rose on 1 April 2026 to £455,000 for complaints about acts or omissions on or after 1 April 2019, and £205,000 for earlier ones.

The scale of voluntary distress payments is now public record because the Treasury Committee asked. Barclays confirmed that during its outage from 31 January to 2 February 2025, 56 per cent of online payments failed as a result of severe degradation of mainframe processing performance, and that it expected to pay between £5 million and £7.5 million for inconvenience or distress, taking its total potential payout over the two-year period to as much as £12.5 million. The second-highest figure among the nine firms was £350,000, from Bank of Ireland. One very large number, then a cliff: these payments are driven by the severity and timing of individual events, not by aggregate downtime. The Barclays outage landed on payday.

■ Four published incidents, and what each one teaches

Bank of England RTGS, 20 October 2014. Roughly nine hours of outage in the system that settles CHAPS. Operating hours were extended to 20:00 and all submitted payments, 142,759 of them worth £289.3 billion, settled the same day. Deloitte's independent review, published with the Bank's response on 25 March 2015, found the root cause to be defects introduced as part of functionality changes made in April 2013 and May 2014. The Bank aims for 99.95 per cent availability of RTGS settlement services to CHAPS, and had achieved 100 per cent from 2008 to 2013. Nine compensation claims were settled, totalling £4,056.89. The lessons the Bank accepted were almost entirely about governance, testing and crisis management rather than software: reconstituting the RTGS board under a Deputy Governor, deferring non-routine changes, separating test and pre-production environments, and reducing the barriers to invoking the MIRS contingency solution so that it could switch and fix in parallel rather than treating failover as a last resort. That last point is the most transferable. A contingency you are afraid to invoke is not a contingency.

Visa Europe, 1 June 2018. Charlotte Hogg, then Visa's European chief executive, wrote to the Treasury Committee explaining that the disruption ran from 14:35 on 1 June to 00:45 the following day, that 51.2 million transactions were initiated across Europe in that window and 5.2 million failed, and that 2.4 million UK transactions failed to process. The cause was a very rare partial failure of a component within a switch at the primary of two data centres, either of which was designed to handle 100 per cent of European volume alone. The partial nature of the failure is the whole lesson: because the primary had not cleanly failed, the backup switch did not activate, and the malfunctioning primary kept trying to synchronise messages with the secondary site, creating a backlog that degraded the secondary's own processing. It took nearly five hours to fully deactivate the failing system. Redundancy protects against a component that stops. It does not automatically protect against a component that continues, badly.

TSB, April 2018. A platform migration in which, as the regulators put it, the data itself migrated successfully but the platform immediately experienced technical failures. On 20 December 2022 the FCA fined TSB £29,750,000 and the PRA £18,900,000, a combined £48,650,000 after a 30 per cent early-resolution discount, for operational risk management and governance failures including management of outsourcing risk. Note what was penalised. Not the fault. The governance around the decision to go live.

CrowdStrike, 19 July 2024. At 04:09 UTC a faulty content configuration update to the Falcon sensor caused an out-of-bounds memory read in the Windows sensor client, crashing machines into a boot loop. The update was reverted at 05:27 UTC, seventy-eight minutes later, and machines that booted after the reversion were unaffected. Microsoft estimated roughly 8.5 million Windows devices were affected, less than one per cent worldwide. Recovery took days, because many machines needed in-

dividual physical intervention and could not be fixed remotely by the very software that had broken them. This is the chapter's central ratio at global scale: the vendor's fix took seventy-eight minutes and the world took a week. It is also why the concentration of financial firms on a handful of shared suppliers became a regulatory priority rather than a procurement one.

■ The regulatory frame around resilience

Three regimes now sit on top of the failure modes described above, and they are worth distinguishing because firms routinely conflate them.

Operational resilience is about outcomes, not systems. Under the FCA's PS21/3 and the corresponding PRA rules, firms must identify their important business services, set an impact tolerance for each — the maximum tolerable level of disruption, expressed as a duration alongside other relevant metrics, beyond which further disruption would cause intolerable harm to consumers or risk to market integrity — and map and test to the point where they can remain within those tolerances in severe but plausible scenarios. Services had to be identified and tolerances set by 31 March 2022, and the transition period for being able to stay within them ended on 31 March 2025. The conceptual shift is from availability targets, which are about systems, to impact tolerances, which are about customers and which assume disruption will happen rather than trying to prevent it.

Critical third parties. Section 312L of the Financial Services and Markets Act 2000, as amended by the Financial Services and Markets Act 2023, allows HM Treasury to designate a third-party service provider as a critical third party where, in its opinion, failure or disruption of its services to firms could threaten the stability of, or confidence in, the UK financial system. The joint rules of the Bank of England, PRA and FCA, published on 12 November 2024 as PRA PS16/24 and FCA PS24/16, took effect on 1 January 2025 but bite on a provider only when a designation order comes into force. HM Treasury announced its first designations on 10 July 2026, effective 13 July 2026: Microsoft Ireland Operations Limited, Google Cloud EMEA Limited, Amazon Web Services EMEA SARL and Oracle Corporation UK Limited. Two features are commonly misunderstood. Oversight extends only to the systemic services provided to the financial sector, not to the provider's wider business. And designation transfers nothing: firms remain fully accountable for the risks in their own third-party arrangements, and a designated provider is not thereby safer than an undesignated one.

Incident and third-party reporting is the data layer that makes the other two supervisable. PS26/2 and PS7/26 also require in-scope firms to notify material third-party arrangements before commitments are finalised and to submit an annual register of them, explicitly so that the regulators can identify concentration risk and inform future critical third party designations. The three regimes are one policy.

■ What good operational practice actually looks like

Strip away the frameworks and the practices that distinguish firms which recover well are unglamorous and specific.

Design for reconciliation before designing for throughput. Every payment instruction should carry an identifier that survives every hop and appears in every record on both sides, so that three-way matching is a lookup rather than an investigation. Retrofitting this after an incident is close to impossible, which is why it has to be a build-time decision.

Make duplicate prevention a control, not a hope. Hash every outbound file and compare it against submissions from the preceding days. Maintain a submission register that operations staff, not just engineers, can read. Require a second authoriser for any submission matching a recent one in value and item count. Set value limits with the sponsoring bank so that an anomalous total cannot be sub-

mitted at all. Riverside's four minutes of code is the entire fix, and almost nobody writes it until after the first time.

Alert on absence, not only on errors. The expected file that never arrives is the most expensive silent failure in the industry.

Rehearse the failover you are afraid of. The Bank of England's own accepted recommendation after 2014 was to reduce the barriers to invoking its contingency solution so that it could switch and fix in parallel. A failover that has never been executed in anger, or that requires a decision nobody wants to make, is a document rather than a capability. This is what the impact-tolerance regime is really testing.

Pre-agree the recovery paperwork: recall templates, beneficiary letters, the customer notice, the redress matrix saying what is paid for what category of harm. Drafting these during an incident guarantees they will be slow, inconsistent and later criticised.

Separate the restore-service team from the restore-correctness team on day one. They have different skills, timescales and definitions of done, and if they are the same people the reconciliation starts two weeks late because everyone is exhausted.

Write a blameless post-incident review with a real timeline, publish the actions with owners and dates, and then check them. The commonest finding in enforcement notices is not that a firm failed to identify a weakness. It is that the firm identified it, wrote it down, and did not fix it in time. Keep the incident log as if a regulator will read it, because one may: contemporaneous, timestamped, factual, including the decisions not taken and why.

■ Returning to the tap

This book opened in a shop in Leeds, with a card touched against a terminal, and with the observation that almost nobody's description of what happens next is correct. Money did not move. No object travelled. A record was edited, and everyone accepted the edit because everyone accepted the record-keeper.

Fifty-three chapters later, you know what the edit involves. That the terminal runs an application which must satisfy a kernel specification, that the token in the phone is not the card number, that the authorisation travels through an acquirer to a scheme to an issuer and back inside a second or two, and that it is a promise rather than a payment. That the movement happens hours later in a settlement file, and finality later still across accounts at the Bank of England. You know who is liable if the card was stolen, who is liable if the shopkeeper never delivers, what the interchange fee is and why it is capped, why the merchant's money arrives on Wednesday, and which regulator would care if it did not.

You also now know the thing the tap is designed to conceal. Behind that gesture sit several dozen organisations, any of which can fail; a legal framework whose entire purpose is to make the result stick even when somebody has made a mistake; a reconciliation apparatus that quietly resolves millions of small disagreements every day; and people whose job, when the disagreements stop being small, is to spend six weeks matching numbers to numbers until the two sides of a ledger agree again.

In 2024, according to UK Finance's UK Payment Markets report, nearly 49 billion payment transactions were made in the United Kingdom, of which 18.9 billion were contactless. Pay.UK's figures for 2025 show approximately 5.03 billion Direct Debits and 1.83 billion Bacs Direct Credits, including around 378 million payroll payments and 267 million state pension payments, running through a system whose cycle is still three working days and whose input files are still eighteen characters wide in

the fields that matter. The overwhelming majority of those payments worked. The ones that did not were caught, in most cases, by somebody looking at a number that was wrong.

That is what money moving actually is. Not a substance travelling, but a very large number of records being written, checked against each other, and corrected. Ada's notebook, with more lawyers.

The tap takes half a second. Everything in this book is what makes the half-second safe to trust, and everything in this chapter is what happens on the days when it is not. Both halves are the system. A reader who understands only the first half understands payments the way somebody who has only watched aircraft take off understands aviation.

Now you have seen the landings too.

54.98 Common wrong ideas

Wrong: The regulatory clock starts when you understand what has gone wrong. **Right:** The initial report is due within four hours of first detection, not from root cause or confirmed impact, so a firm must be able to file a report that says something is wrong, we do not yet know what, here is what we do know.

Wrong: A complete failure is the worst outcome. **Right:** A clean total failure tells you the state of the world and can be resent; ambiguity is the expensive one, because items that returned no answer at all can be neither safely resubmitted nor safely abandoned.

Wrong: The cost of an incident is the outage. **Right:** The operational loss is usually a rounding error against redress, remediation, enforcement and the supervisory tail — TSB's combined fine was £48,650,000 against £32.7 million already paid in redress, while nine hours of national settlement outage in 2014 produced nine claims totalling £4,056.89.

Wrong: The bigger the system that failed, the bigger the bill. **Right:** Where the harm lands determines the bill, which is why a retail migration cost tens of millions and an RTGS outage cost four thousand pounds.

Wrong: The technical fix is always the short part. **Right:** Where the failure is a design decision rather than a fault there is nothing to restore to, which is why TSB's disruption ran for eight months and why the 2014 RTGS defects had been latent for around eighteen months.

Wrong: The incident is over when service is restored. **Right:** Service is restored when payments flow again; correctness is restored when every payment that should have happened has happened exactly once, and the gap between those two moments is typically two to eight weeks.

Wrong: Redundancy protects against the failure of a component. **Right:** It protects against a component that stops, not one that continues badly — Visa's primary switch in June 2018 failed partially, so the backup never activated and the malfunctioning primary degraded the secondary site as well.

Wrong: Error-rate monitoring will find the incident. **Right:** The failure modes that cost most produce no errors at all, so mature operations alert on the absence of expected traffic rather than only on the presence of errors.

Wrong: A settled payment sent by mistake can be pulled back. **Right:** Finality exists precisely to stop that, which is why recovering Citibank's \$558,558,375.74 of mis-sent Revlon principal took two years and two courts, with the money gone at the halfway mark.

Wrong: Designating a cloud provider as a critical third party makes the firms relying on it safer. **Right:** Oversight extends only to the systemic services provided to the financial sector, designation transfers nothing, and firms remain fully accountable for the risks in their own third-party arrangements.

54.99 Chapter summary in 20 lines

1. A payment system is not a machine that works and sometimes fails, but a machine that fails at a known rate wrapped in enough correction to make the failures survivable.
2. Nine of the largest United Kingdom banks reported at least 803 hours of unplanned outage across at least 158 incidents between January 2023 and February 2025, which is the baseline rather than a scandal.
3. The visible event is the leak and the ledger is the road, and the road takes weeks after the leak is stopped.
4. Riverside Court's duplicated payroll of £176,400 needed four minutes of code to prevent, took three days to detect and eleven weeks to reconcile, and ended in £6,300 written off plus £940 of borrowing costs.
5. Detection in payments is ordinarily a wrong number noticed for an unrelated reason, not a red light.
6. Once money has moved and settled it is genuinely gone, and getting it back requires the recipient's agreement or a court, because ending the argument about who owes what is the whole purpose of a payment.
7. Several unrelated regulatory clocks start at once, triggered by different events and owed to different people, and the United Kingdom four-hour clock runs from first detection.
8. Classification is a regulatory act rather than an engineering one, because it starts those clocks at the moment impact is least knowable.
9. Containment usually destroys evidence, so good runbooks specify in advance what must be captured before each containment action.
10. Restoring service is not restoring correctness, and conflating them is the classic mistake.
11. Partial failures are worse than total ones because the state of the world cannot be established from the firm's own records, and the temptation to resubmit the whole file converts a partial failure into a duplicate.
12. Batch file rails have no idempotency key, because two identical payroll runs on consecutive days are a legitimate business event no scheme can distinguish from an error, so duplicate prevention is a control submitters must build themselves.
13. Whether a payment can be reversed depends on the rail: Direct Debit and cards give strong structured remedies, Bacs credits and Faster Payments give only a request, and CHAPS gives none once settled.
14. Post-incident reconciliation is a three-way problem between the scheme's view, the bank's view and the firm's own ledger, and every break must be categorised, aged and given a named owner.
15. The arithmetic is unforgiving: 12,400 uncertain items, ninety per cent matched automatically, twenty minutes each on the remainder, is about 413 hours before the tail and the write-offs even begin.
16. Redress is not discretionary generosity but a stack of regimes — regulation 76 for unauthorised transactions, regulations 91 to 95 for defective execution, the £85,000 scam reimbursement requirement, DISP, the Ombudsman and the Consumer Duty.
17. Barclays expected to pay between £5 million and £7.5 million in distress payments for one outage that landed on payday, while the next highest figure among nine firms was £350,000, because severity and timing drive these payments rather than aggregate downtime.

18. The four published incidents teach four distinct lessons: a contingency you are afraid to invoke is not a contingency, redundancy does not protect against a component that continues badly, regulators penalise the governance of the decision rather than the fault, and a seventy-eight-minute vendor fix can take the world a week.
19. Operational resilience, the critical third parties regime and the new incident and third-party reporting rules coming into force on 18 March 2027 are three parts of one policy, the last being the data layer that makes the other two supervisable.
20. Money moving is not a substance travelling but a very large number of records being written, checked against each other and corrected, and this chapter is what happens on the days when the correcting becomes visible.

Sources, all consulted in August 2026: the Payment Services Regulations 2017, regulations 76, 91 to 95 and 99, and FCA Handbook SUP 15.14; EBA Guidelines EBA/GL/2017/10; FCA PS26/2 with FG26/3 and FG26/4, and PRA PS7/26 with SS1/26, published 18 March 2026 and in force 18 March 2027; FCA PS21/3 and the FCA's operational resilience insights pages; PRA PS16/24 and FCA PS24/16 on critical third parties, with HM Treasury's designation announcement of 10 July 2026; section 312L of the Financial Services and Markets Act 2000 as amended by FSMA 2023; Regulation (EU) 2022/2554 (DORA), Article 19; the Financial Markets and Insolvency (Settlement Finality) Regulations 1999, S.I. 1999/2979; the Bank of England's response to the independent review of the RTGS outage of 20 October 2014, published 25 March 2015; the Bank of England and FCA announcement of the TSB fine, 20 December 2022; Visa's letter to the Treasury Committee of 15 June 2018; the Treasury Committee's publication of bank IT failure correspondence, 6 March 2025; Payment Systems Regulator PS24/7 and its APP reimbursement pages; Financial Ombudsman Service award limits for 1 April 2026; Pay.UK's Bacs glossary and Bacs annual processing statistics 2025; UK Finance UK Payment Markets 2025; Microsoft's statement of 20 July 2024 and CISA's alert of 19 July 2024; and In re Citibank August 11, 2020 Wire Transfers, No. 21-487 (2d Cir., 8 September 2022).

ABOUT THE AUTHOR

Shikhar Singh

Founder & Chairman

KedByte Technologies Private Limited

Shikhar Singh is the Founder & Chairman of KedByte Technologies Private Limited.

He wrote this book from the position of someone building payment software rather than someone commenting on it, and the difference shows in what it chooses to explain. The chapters that go deepest are the ones where he has watched capable engineers lose days: the gap between authorisation and settlement, the several identifiers all called "the transaction number", and reconciliation, which is unglamorous, unavoidable, and where more young financial technology companies fail than for any more interesting reason.

His conviction is that payments has no genuinely hard idea in it, only a very long chain of simple ones, most of which are documented for people who already understand them. The barrier is not difficulty. It is that the industry writes for itself. That is a solvable problem, and this book is an attempt at solving it.

Under his direction, KedByte Technologies Private Limited treats technical education as infrastructure rather than as a product feature: knowledge that should be transferable, verifiable, and free of the vocabulary barriers that keep capable people out of the field.



ABOUT THE PUBLISHER

KedByte Technologies Private Limited

KedByte Technologies Private Limited publishes technical work under the KedByte Press imprint.

The editorial standard applied to this volume is simple to state and difficult to meet. Every explanation must work twice: once for a reader who has never opened a terminal, and once for an engineer who needs the exact figure. Every claim that can be checked must be checked. Every figure that will go stale must carry a date. Where the field genuinely disagrees, both positions are given rather than one being quietly chosen.

Where this book uses real evidence, it uses it unaltered. The network trace in Part F, the failed pushes in Part G and the error messages quoted throughout are reproduced exactly as they were observed. Diagnosis is taught against ambiguous real data, because that is the only kind there is.



KEDBYTE

TECHNOLOGIES PRIVATE LIMITED

COLOPHON

This first edition of *How Money Moves* runs to five volumes and fifty-four chapters, and was set entirely from plain text.

The body text is set in TeX Gyre Pagella, a digital revival of Hermann Zapf's Palatino, chosen for long-form reading. Headings, tables and the KedByte identity are set in Poppins. Code, command output and diagrams are set in DejaVu Sans Mono, and every diagram in this book is drawn in plain ASCII so that it survives being copied anywhere.

The book was composed with pandoc and typeset with XeLaTeX on A4 stock. The single accent colour used on rules, chapter numbers and section headings is KedByte navy.

Shikhar Singh

Founder & Chairman, KedByte Technologies Private Limited



First Edition • KedByte Press