



KEDBYTE

SOFTWARE · AI · AUTOMATION

HOW MONEY MOVES

Volume I — What Money Is

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 I — What Money Is

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 I — What Money Is **1**

1	The Oldest Question	2
	1.0 What this chapter gives you	2
	The plain version	3
	Where the plain version stops being true.	6
	The technical version	8
	1.98 Common wrong ideas	13
	1.99 Chapter summary in 20 lines	14
2	A Balance Is Not a Thing	16
	2.0 What this chapter gives you	16
	The plain version	17
	Where the plain version stops being true.	19
	The technical version	21
	2.98 Common wrong ideas	27
	2.99 Chapter summary in 20 lines	28
3	Double Entry	30
	3.0 What this chapter gives you	30
	The plain version	31
	Where the plain version stops being true.	34
	The technical version	36
	3.98 Common wrong ideas	43
	3.99 Chapter summary in 20 lines	44
4	Ledgers, Journals and Accounts	46
	4.0 What this chapter gives you	46
	The plain version	47
	Where the plain version stops being true.	49
	The technical version	51
	4.98 Common wrong ideas	58
	4.99 Chapter summary in 20 lines	58
5	What a Bank Is	60
	5.0 What this chapter gives you	60
	The plain version	61
	Where the plain version stops being true.	64
	The technical version	66
	5.98 Common wrong ideas	73
	5.99 Chapter summary in 20 lines	74
6	Two Kinds of Money	76
	6.0 What this chapter gives you	76
	The plain version	77
	Where the plain version stops being true.	79

The technical version	81
6.98 Common wrong ideas	87
6.99 Chapter summary in 20 lines	87
7 Clearing Versus Settlement	89
7.0 What this chapter gives you	89
The plain version	90
Where the plain version stops being true.	92
The technical version	94
7.98 Common wrong ideas	101
7.99 Chapter summary in 20 lines	102
8 Netting	104
8.0 What this chapter gives you	104
The plain version	105
Where the plain version stops being true.	107
The technical version	109
8.98 Common wrong ideas	118
8.99 Chapter summary in 20 lines	119
9 Currency, Decimals and the Money Type	121
9.0 What this chapter gives you	121
The plain version	122
Where the plain version stops being true.	123
The technical version	126
The point of all this	133
9.98 Common wrong ideas	133
9.99 Chapter summary in 20 lines	134
10 Reading a Statement	136
10.0 What this chapter gives you	136
The plain version	137
Where the plain version stops being true.	139
The technical version	141
Where Volume I leaves you	149
10.98 Common wrong ideas	150
10.99 Chapter summary in 20 lines	150



VOLUME I

What Money Is

Obligation, records and double entry. What a balance actually is, what a bank actually is, and the difference between clearing and settlement. Everything here is permanent.

The Oldest Question

1.0 What this chapter gives you

1. You will be able to explain why a payment is an edit to a record rather than the movement of an object, using Ada's notebook as the worked example.
2. You will be able to state Caroline Humphrey's claim precisely — that no barter economy, pure and simple, has ever been described — and say what it does not claim, so that you are not caught out by anyone who has read the source.
3. You will be able to describe how a tally stick worked, why splitting a notched stick lengthways made it hard to forge, and how the words *stock* and *foil* reached modern finance from a piece of hazel.
4. You will be able to place coinage correctly in time: administrative writing at Uruk from about 3100 BC, the first coins some two and a half thousand years later beneath the temple at Ephesus.
5. You will be able to explain why the earliest coins undermine the commodity theory of money rather than supporting it, given that Lydian electrum of roughly 54 per cent gold was tarified as though it were 73 per cent.
6. You will be able to distinguish the unit in which an obligation is denominated from the asset in which it is settled, and recognise that distinction as four thousand years old.
7. You will be able to say what actually makes a record money — that an identifiable party is obliged to honour it, and that something makes the obligation stick — rather than repeating the slogan that money is information.
8. You will be able to name the counterparty behind any unit of money you hold: your bank for a deposit, the Bank of England for a banknote, the retailer for a gift card.
9. You will be able to quote the British proportions that matter: about £91.5 billion of notes across roughly 4.98 billion of them, against bank deposits making up 97 per cent of the money in circulation.
10. You will be able to read a hundred-character Bacs Standard 18 record and explain why three of its fields exist purely so that human beings on both sides can later agree what the record meant.

Ask a hundred people what happened when they last paid for something and ninety-nine will describe a movement. Money left my account and went to theirs. Money came out of the machine. Money changed hands.

Almost none of that is true, and the exceptions are rarer than you think. When you tap a card in a shop in Leeds, nothing physical travels between you and the shopkeeper. When your salary arrives on the twenty-fifth, no lorry of banknotes crosses the country. In every one of these cases, what changed was a record. Somebody's obligation to somebody else was written down, cancelled, netted off, or

handed on. The physical object, when there is one, is a token that stands for the record. The record is the money.

This is not a modern discovery about a modern system. It is the oldest fact about money there is, and the evidence for it is older than coins and older than writing. This chapter is about that evidence: the clay of Mesopotamia, the split hazel sticks of the English Exchequer, and the awkward, well-documented fact that the barter economy every schoolchild is taught about has never actually been found.

The reason this matters for a book about payments is simple. If money were a substance, moving it would be a logistics problem, and the interesting questions would be about vaults and armoured vans. Because money is a record, moving it is an information problem, and the interesting questions are about who keeps the record, how two records are reconciled, what happens when they disagree, who bears the loss when one is wrong, and how quickly the thing becomes final. That is what the rest of this book is about. This chapter is about why.

The plain version

Imagine a village with one shop.

The shop is run by a woman called Ada. She sells bread, milk, nails, batteries and the sort of things you need at short notice. Nobody in the village has much cash, because the nearest bank is forty minutes away, so Ada keeps a notebook under the counter. When Tom the carpenter comes in for a loaf, Ada writes: *Tom, 1 loaf, £1.20*. When Tom finishes a set of shelves for Ada's back room, she writes: *Ada owes Tom, £150*.

Nothing has moved. There are no coins in this story yet. But something real has happened. Tom is owed £150 by Ada. Tom owes Ada £1.20. If Tom keeps buying bread every day for four months, the two numbers will meet in the middle and cancel each other out, and the whole relationship will settle without a single coin ever appearing.

Now watch what happens when the village gets a little bigger.

Sam the mechanic fixes Tom's van. The bill is £80. Tom has no cash. But Tom knows that Ada owes him £150, so he says to Sam: "Go to Ada. Tell her to take £80 off what she owes me and put it against your account instead." Sam agrees, because Sam trusts Ada's notebook, and because Sam owes Ada money anyway for a year of milk and nails. Ada opens the book and writes three lines. Tom's credit falls from £150 to £70. Sam's debt to Ada falls by £80. Tom's debt to Sam is gone.

Look at what just occurred. A debt was paid. Nobody handed over anything. No object changed owners. A record was edited, and everyone accepted the edit because everyone accepted the record-keeper. This is a payment. It is the whole of payment, in miniature, and every payment system in the world is a more elaborate version of Ada's notebook with more checks, more speed and more lawyers.

Now here is the part that surprises people. This is not a clever modern shortcut that people invented once coins became inconvenient. As far as anyone can tell from the archaeological and anthropological record, this is the *original* arrangement. Ada's notebook came first. Coins came very much later.

■ The story you were probably told

The version most of us learn goes like this. Long ago, people bartered. If you had shoes and wanted bread, you had to find a baker who happened to want shoes. This was maddening, so somebody hit on the idea of a thing everyone would accept — cattle, salt, shells, eventually lumps of silver — and money was born to fix the problem of barter.

It is a beautiful story. Adam Smith told a version of it in 1776 in *The Wealth of Nations*, using the example of a butcher with more meat than he can eat and a brewer and a baker who want some of it but have nothing the butcher happens to need. Smith's conclusion was that every prudent person would start keeping a stock of "some one commodity or other, such as he imagined few people would be likely to refuse in exchange", and that metals eventually won because they do not rot and can be cut into precise amounts.

The trouble with the story is that when anthropologists went looking for a society that actually worked this way, they could not find one. Not a single one. What they found instead, over and over, in every corner of the world, was some version of Ada's notebook: people who knew each other, kept mental or written track of who owed what, and settled up occasionally or never. Where they found barter, they usually found it happening between strangers, or between people who had no shared record-keeper and no intention of building a relationship, or in places where a money system had recently fallen apart.

The reason is not hard to see once you look at Ada's village. Barter is what you do when you cannot trust a record. If you and I are neighbours, a running tab is easier for both of us than lugging goods around and haggling over exchange rates. It is only when I do not expect to see you again, or cannot enforce anything against you, that we both insist on a simultaneous swap where neither of us has to trust the other for a single day.

■ Two copies of one fact

The obvious weakness of Ada's notebook is Ada. What if she writes the wrong number, or rubs a line out, or dies and leaves the book to a nephew who reads Tom's £150 as £15?

Human beings solved this a very long time ago, and one of the neatest solutions was a stick.

Take a piece of hazel or willow, about the length of your hand, and cut notches across it. A wide notch means a large amount, a narrow notch a small amount. Now — and this is the clever part — split the stick lengthways, right through every notch, so that you end up with two halves that each carry the same notches, along the same grain, with the same knots and flaws.

Give one half to the person who is owed the money. Keep the other half yourself. Neither half can be altered without the alteration being obvious, because if either of you adds or removes a notch the two halves will no longer match. To settle the debt, the two halves are brought back together. If they fit, the record is proved. If they do not, somebody is lying, and the wood says so.

That is a tally stick. Two copies of one fact, made from a single piece of wood so that no forger can produce a matching partner. The English Exchequer — the government department that collected the king's revenue — used them for something like seven hundred years, and the vocabulary has stuck to English so firmly that we no longer notice it. The longer half of the split stick, held by the person owed the money, was called the *stock*. That is why an investor holds *stock*, and why a person who lends to a company is a *stockholder*. The shorter half was the *foil*, held by the debtor, which is why getting a raw deal is getting the *short end of the stick*.

And here is the point that ought to stop you. A tally stick was not a receipt in the modern sense, a dead scrap of paper filed away. It circulated. If the Exchequer owed you £100 and gave you a tally saying so, you could pay your own creditor by handing over the tally. He would accept it because he, in turn, could hand it to the taxman and have it counted against his own bill. The stick moved from hand to hand carrying an obligation with it, and it did this for centuries. It was, in every practical sense, money. It was also, unmistakably, a record. Not gold, not silver, not a substance of any kind. A

notched stick, whose entire value came from the fact that a specific institution had written down that it owed you something.

■ The clay before the wood

Go back much further and you find the same idea in a different material.

In the cities of southern Mesopotamia — modern-day Iraq — around five thousand years ago, people were running large organised economies centred on temples and palaces. Grain came in from the fields. Rations went out to workers. Sheep were counted, beer was issued, labour was allocated. All of this had to be tracked, and the tracking was done in clay, because clay was free and it lasted.

The earliest of these records are not sentences. They are not literature or prayer or law. They are accounts. They say things like: *this many measures of barley, from this store, to this person*. One clay tablet in the British Museum, made in the very last centuries of the fourth millennium BC, is a record of beer, marked with five different kinds of number sign. It is about the size of a postcard. It is not a poem. It is a ledger entry.

This is worth sitting with, because it turns the usual story of civilisation upside down. Writing was not invented so that people could tell stories and later, as an afterthought, applied to business. On the best current evidence, it went the other way. The pressure that produced writing was the pressure of keeping track of who owed what to whom. The alphabet is a side effect of accountancy.

And how long after this did coins arrive? Roughly two and a half thousand years. The first objects most people would recognise as coins — small, standardised, stamped lumps of precious metal — appear in what is now western Turkey around the end of the seventh century BC. Ninety-three of them were dug up by British Museum archaeologists beneath the great temple of Artemis at Ephesus, left there as offerings.

Put that gap in terms you can feel. If the clay ledgers of Uruk had been written on the day of the Battle of Hastings, coins would not have been invented until about four hundred years from now. Writing down debts is not a step on the road to coins. Coins are a late refinement on a system of written obligation that was already ancient when they turned up.

■ What money is, in one sentence

Every form of money you have ever handled is a record of an obligation before it is an object.

The five-pound note in your pocket says so on its face, in words most people have stopped reading: *I promise to pay the bearer on demand the sum of five pounds*. That is not decoration. It is a description of what the note is. The note is the Bank of England's IOU. The bank owes you; the note is the evidence.

The number in your banking app is the same thing one step further along. It is not a store of your pounds sitting in a box with your name on it. It is your bank's written admission that it owes you that amount, payable on demand. If you have ever wondered why bank failures are frightening in a way that losing a wallet is not, this is why. A wallet is a container. A bank balance is a promise.

And the proportions are not close. At the end of February 2026, all the Bank of England banknotes in existence — every fiver, tenner, twenty and fifty in every till, pocket and cash centre in the country — came to about £91.5 billion, spread across roughly 4.98 billion notes. Set that against the money that exists only as entries in bank ledgers and it is a rounding error. The Bank of England's own analysis puts bank deposits at 97 per cent of the money in circulation. Ninety-seven pounds in every hundred is a line in a database. Three are paper.

Now look at what that money does in a year. In 2025, the Bacs system — the one that pays most British salaries and collects most British household bills — processed 6.86 billion payments worth £6.05 trillion. Faster Payments, the system behind the instant transfer you make from your phone, handled 5.55 billion payments worth £4.84 trillion. That is well over twelve billion instructions in a single year, moving nearly eleven trillion pounds, and at no point did a physical object representing value travel from a payer to a payee. What travelled was information: a short, structured message saying *reduce this account by this much, increase that account by this much, and here is the reference so both sides can agree it happened*.

Ada's notebook, scaled to sixty-eight million people.

Where the plain version stops being true

The plain version above is a good map. Like all good maps it lies by simplification, and four of those lies matter enough to correct before we build anything on top of them.

■ Barter did exist; what has never been found is a barter economy

It is fashionable now to say flatly that barter is a myth. That overshoots, and a reader who repeats it at dinner will be caught out by anyone who has actually read the source.

The foundational statement is Caroline Humphrey's, in the journal *Man* in 1985, and it is more careful than its popular version: "No example of a barter economy, pure and simple, has ever been described, let alone the emergence from it of money; all available ethnography suggests that there never has been such a thing." Read the clauses. She is not saying that people never swap goods directly. Of course they do; it happens in playgrounds, at borders, between prisoners, and across the front lines of wars. She is saying two narrower and much stronger things: that no society has ever been observed running its whole economy on direct exchange without money, and that money has never been observed emerging out of such a society.

Humphrey's own fieldwork makes the point sharply. She studied the Lhomi of north-eastern Nepal in 1979 and 1980, a community that did a great deal of barter. They did not do it because they had never heard of money. They did it because currency was scarce and their links to monetary markets had broken down. Her argument is that barter in the observed world is overwhelmingly a *post-monetary* phenomenon: it appears where a money economy has decayed, been cut off, or run short of the stuff. It coexists with money rather than preceding it.

So the correction is this. Barter is real. What is unsupported is the sequence — barter first, then money to fix barter — and the idea that money is fundamentally a solution to the inconvenience of swapping shoes for bread. The evidence points instead to credit and reckoning coming first, with barter as one of several things that happen at the edges: between strangers, across hostile boundaries, and in the wreckage of monetary systems.

There is a second, subtler correction hiding here. Societies without general-purpose money often do not have one undifferentiated market at all. Paul Bohannan's study of the Tiv of central Nigeria, published in *American Anthropologist* in 1955, described distinct spheres of exchange, with goods in one sphere not freely convertible into goods in another, because the culture did not treat them as commensurable. General-purpose money did not merely lubricate exchange in such societies; it dissolved walls that people had wanted kept up. Money is not a neutral efficiency improvement laid over an existing market. It changes what counts as a market.

■ “Credit came first” does not mean neighbours writing IOUs

The Mesopotamian evidence is real, but it is easy to over-read it in a modern direction.

The tablets from Uruk are overwhelmingly *institutional* accounts. They are the internal records of large temple and palace estates disbursing rations and tracking obligations owed to and by the institution. They are much closer to the payroll and stock-control ledgers of a very large organisation than to two neighbours agreeing a private loan. The obligations they record are frequently administrative — quotas, allocations, deliveries due — rather than freely negotiated credit between equals. Calling this “credit” is fair, but you should picture a state accounting department, not a friendly overdraft.

The second over-read concerns the tokens. Denise Schmandt-Besserat’s thesis — that small clay tokens used for counting from around the eighth millennium BC were sealed inside clay envelopes called *bullae*, that the shapes were then impressed on the outside of the envelope so you could see the contents without breaking it, and that the impressions eventually made the tokens redundant and became writing — is enormously influential and genuinely elegant. It is also contested in the detail. Not every specialist accepts that the token system maps cleanly onto the earliest sign inventory, or that a single evolutionary line runs from token to envelope to tablet. Treat the token-to-writing story as the leading hypothesis rather than a settled fact, and note that the broader claim — that the earliest writing was administrative and economic — rests on the tablets themselves and does not depend on the token thesis being right.

■ The tally stick was a receipt that *became* negotiable, and the sources disagree about the details

Two cautions. First, the tally began life as proof of payment into the Exchequer, not as a freely circulating currency. Its negotiability was a consequence of what it was — a hard-to-forge, institution-backed record of a sum owed by the Crown — rather than its design purpose. A tally in a strongbox was a receipt. A tally passed to a creditor was money. Same object, different use. Money is a role played by a record, not a class of object, and the tally is the cleanest illustration in European history.

Second, be careful about the notch conventions, because sources genuinely conflict. The *Dialogus de Scaccario*, the treatise on Exchequer procedure written around 1179, gives them in body parts: a cut the thickness of the palm for a thousand pounds, the breadth of a thumb for a hundred, the breadth of the little finger for twenty, the width of a swollen barleycorn for a single pound, and a narrower cut for a shilling. The 1911 *Encyclopædia Britannica* gives them in inches, and does not match: one and a half inches for a thousand pounds, one inch for a hundred, three-eighths of an inch for ten, a half-notch of that size for a pound, three-sixteenths for a shilling, the smallest notch for a penny, and small drilled holes for halfpennies. Note that the two schemes do not even agree on which denominations existed — the *Dialogus* has a twenty-pound notch and the *Britannica* has a ten-pound notch. The honest position is that the conventions varied across seven centuries and that no single table describes them all.

Anyone who tells you the tally system was uniform for seven hundred years is telling you something the sources do not support.

■ Information is only money when someone is obliged to act on it

This is the correction that matters most for everything that follows, and it is where the slogan “money is information” starts to mislead.

A record is not money because it is a record. It is money because a specific, identifiable party is obliged to honour it, and because there is a mechanism that makes that obligation stick. Ada’s notebook worked because Ada was solvent, present, honest and answerable to a village that would make her life difficult if she cheated. Take away any one of those and the notebook becomes a list of numbers.

Three consequences follow, and they run through the rest of this book.

The first is that every unit of money in a modern economy has a counterparty. Your bank balance is your bank's liability. A banknote is the central bank's liability. A gift card is the retailer's liability. When people say a payment "settles", what they mean precisely is that the obligation has been discharged and the record can no longer be reversed. Settlement is the moment a promise stops being a promise, and a great deal of the machinery in this book exists to define that moment exactly, because until it arrives somebody is carrying risk.

The second is that cash is genuinely different, and the plain version glosses over it. When I hand you a twenty-pound note, the transaction is complete at the instant of handover: no clearing, no settlement window, no counterparty left owing anything, no possibility of recall. That is a real and unusual property, and it is why cash remains stubbornly useful in exactly the situations where records are unreliable — power cuts, disasters, distrust and crime. Every electronic payment is an attempt to reproduce the finality of cash across distance and time, and none fully manages it.

The third is that Alfred Mitchell-Innes's famous formulation — writing in the *Banking Law Journal* in 1913, "Money, then, is credit and nothing but credit. A's money is B's debt to him, and when B pays his debt, A's money disappears" — is a position in a long-running argument, not a neutral summary of the facts. Innes was demolishing the commodity theory of money, and he was right about the Exchequer in ways that mainstream histories of his day had got badly wrong. He was also making a maximalist claim. Economists still disagree about whether commodity-money origins matter, about how far the state's power to demand taxes explains a currency's acceptance, and about whether "credit" is the right frame for physical cash. This book leans towards the credit view because the historical record supports it and because it is by far the more useful frame for understanding payment systems. You should know that it is a lean and not a law.

The technical version

■ The anthropological critique, stated precisely

The claim under examination is what economists call the *double coincidence of wants* narrative, given its canonical modern statement in Adam Smith, *An Inquiry into the Nature and Causes of the Wealth of Nations* (1776), Book I, Chapter IV, "Of the Origin and Use of Money". Smith's argument is deductive rather than evidential: he reasons from the division of labour to the necessity of a commonly accepted medium, then to the selection of metals on the grounds of durability and divisibility. He offers no ethnographic case, because in 1776 there was none to offer.

The empirical counter-position is set out in Caroline Humphrey, "Barter and Economic Disintegration", *Man* (New Series), Volume 20, Issue 1 (1985), pages 48 to 72, whose ethnographic base is fieldwork among the Lhomi of the Upper Arun valley in north-eastern Nepal in 1979 and 1980. Her thesis is that observed barter is characteristically a phenomenon of economies "de-coupled from monetary markets" — that it coexists with money rather than preceding it, and that under severe currency scarcity money itself may become one of the items bartered rather than the index of value for the rest.

The argument was popularised, and considerably extended, in David Graeber, *Debt: The First 5,000 Years* (Melville House, 2011), which treats the barter narrative as an origin myth internal to economics rather than a historical finding. Graeber's synthesis has been criticised by economists and by some historians on points of detail; his central negative claim — that no one has produced the barter economy the textbooks require — has held up.

A related strand concerns the non-fungibility of value categories. Paul Bohannon, “Some Principles of Exchange and Investment among the Tiv”, *American Anthropologist*, Volume 57, Number 1 (1955), pages 60 to 70, described goods segregated into distinct spheres, with movement within a sphere ordinary and movement between spheres exceptional and morally freighted. Marcel Mauss’s *Essai sur le don* (1925) had earlier established that obligation-creating gift exchange, rather than spot barter, dominates in many non-market societies. The cumulative picture is not that early societies had a primitive form of our market, but that they had systems of obligation whose logic differs from market exchange at the root.

■ The Mesopotamian record

Proto-cuneiform, the earliest attested phase of the cuneiform writing system, appears at Uruk in southern Mesopotamia and is conventionally dated to approximately 3100 to 2900 BC. The Metropolitan Museum of Art’s account of the origins of writing notes that the site “was dominated by large temple estates whose need for accounting and disbursing of revenues led to the recording of economic data on clay tablets”, that roughly 6,000 proto-cuneiform tablets survive with more than 38,000 lines of text, and that among the earliest signs are pictures of “rations that needed to be counted, such as grain, fish, and various types of animals”.

A representative object is British Museum registration number 1989,0130.4, the so-called British Museum Society tablet, acquired at Christie’s in 1989 from the Erlenmeyer collection and displayed in Room 56. It is dated 3100 to 3000 BC, measures 9.4 by 6.87 by 2.3 centimetres, and is catalogued as a record of beer impressed with five different types of numerical symbol. The presence of multiple distinct numerical notations on a single tablet is itself informative: proto-cuneiform used different sign systems for different commodity classes, so that counting sheep and measuring grain did not share a notation.

The prehistory of these tablets is the token thesis associated with Denise Schmandt-Besserat, set out in *Before Writing, Volume I: From Counting to Cuneiform* (University of Texas Press, 1992). The proposed sequence runs: small geometric clay tokens used as counters from the Neolithic; tokens sealed inside hollow clay envelopes (*bullae*) to fix a transaction’s content; tokens impressed on the envelope’s exterior before sealing so the contents could be read without breaking it; and finally the impressions alone, on a flat tablet, the tokens having become redundant. The thesis remains influential and remains disputed in its details.

On units and rates, the Old Babylonian and earlier systems were sexagesimal. The mina was divided into 60 shekels; the shekel was a weight of silver of roughly 8 grams. Michael Hudson’s work on palatial credit records the equivalence of the shekel-weight of silver with a *gur* of grain of 300 *sila*, an equation set administratively by the palace rather than discovered in a market. Hudson’s argument on interest is that the canonical Mesopotamian rate was a product of the number system rather than of economic conditions: one shekel per mina per month, that is one sixtieth per month, giving 12 shekels per mina per year, or 20 per cent per annum in decimal terms, and doubling the principal in 60 months under simple accrual. He notes, importantly for our purposes, that a debt denominated in silver did not require silver in settlement — “What was called a ‘silver’ debt thus did not mean that actual silver had to be paid, but simply that the interest rate was 20 percent” — with barley an accepted means of discharge.

That distinction between the unit in which an obligation is denominated and the asset in which it is settled is not an antiquarian curiosity. It is the same distinction that separates a currency from a settlement asset in a modern payment system, and it is four thousand years old.

The legal architecture around these obligations is visible in the laws of Hammurabi, of the eighteenth century BC. The provisions on merchants and agents are unambiguous about documentation: an agent must give a receipt for money received, and one who “is careless, and does not take a receipt for the money which he gave” bears the consequence. On deposits, the code requires a party leaving goods for safekeeping to “show everything to some witness, draw up a contract, and then hand it over for safe keeping”, and denies remedy to one who deposits “without witness or contract”. A society that legislates in this much detail about receipts and witnesses is a society whose economy runs on records.

■ The English Exchequer and the tally

The Exchequer tally is documented in the *Dialogus de Scaccario*, written about 1179 and attributed to Richard FitzNeal, Treasurer of England. The physical description in the 1911 *Encyclopædia Britannica* is of “a willow or hazel stick about one inch in depth and thickness, and roughly shaped like a thick knife-blade”, notched across its width and then split lengthways through the notches.

The two halves had names. The longer piece, retained by the party to whom the sum was owed, was the *stock*; the shorter, retained by the Exchequer as the debtor’s record, was the *foil*. Details of the sum were also written on the wood in ink before splitting, so that each half carried both a notched and a written record. Verification was by physical fit: if genuine, the halves “fitted so exactly that they appeared evidently to be parts the one of the other”. This is a mechanical integrity check with no cryptographic component and no central register, and it worked for centuries because the grain and flaw pattern of a split stick was effectively unforgeable by the technology of the period.

On notch values, as noted above, the sources conflict. Set them side by side.

Denomination	Dialogus de Scaccario (c. 1179)	1911 Encyclopædia Britannica
£1,000	cut the thickness of the palm	notch 1.5 inches wide
£100	breadth of a thumb	notch 1 inch wide
£20	breadth of the little finger	not listed
£10	not listed	notch 0.375 inch wide
£1	width of a swollen barleycorn	half-notch of the £10 size
1 shilling	a single cut, narrower	notch 0.1875 inch wide
1 penny	not specified in this form	smallest notch
halfpenny	not specified in this form	a small drilled hole

The correct reading of that table is that the convention changed over time and that neither source describes the whole period.

The tally’s monetary character rests on its transferability. Mitchell-Innes, in his 1913 *Banking Law Journal* article “What is Money?”, put it directly: “Tallies were transferable, negotiable instruments, just like bills of exchange, bank-notes or coins.” His accompanying claim about the Exchequer has aged best: “The general belief that the Exchequer was a place where gold or silver was received, stored and paid out is wholly false. Practically the entire business of the English Exchequer consisted in the issuing and receiving of tallies.” From this he drew the conclusion that gives the credit theory its name: “Money, then, is credit and nothing but credit. A’s money is B’s debt to him, and when B pays his debt, A’s money disappears.”

The system’s end is precisely datable and unusually consequential. An Act of 1782 provided for the discontinuance of wooden tallies, but the change took effect only on the death of the last chamberlain, in 1826. The accumulated stock of obsolete tallies — two cart-loads of them — remained in the Palace of Westminster until 16 October 1834, when Richard Weobley, the Clerk of Works, acting on Treasury

instruction, had them burned in the two underfloor stoves in the basement of the House of Lords. Two labourers, Joshua Cross and Patrick Furlong, fed the stoves through the day. The first flames were seen at six o'clock that evening. The fire destroyed the House of Lords, the House of Commons, St Stephen's Chapel, the Painted Chamber, the libraries and committee rooms of both Houses, the Speaker's House, most of the procedural records of the Commons going back to the late fifteenth century, and the British standard yard and pound of 1496. Westminster Hall and the Jewel Tower survived. Charles Dickens, addressing the Administrative Reform Association two decades later, remarked that in the matter of the tallies "all the red tape in the country grew redder at the bare mention of this bold and original conception".

The Palace of Westminster you see today exists because a government disposed carelessly of an obsolete accounting medium. It is the most expensive records-management incident in British history.

■ Coinage as a late and specialised development

The chronology matters, so state it plainly. Administrative writing in Mesopotamia is attested from approximately 3100 BC. Coinage — standardised, stamped, portable metal of certified weight — appears roughly two and a half millennia later.

The evidential anchor is the deposit beneath the Temple of Artemis at Ephesus, where British Museum excavators recovered 93 electrum coins deposited as offerings during the latter part of the seventh century BC. The Lydian royal issues bearing a lion's head run from about 4.7 grams, one third of a Lydian weight stater, down through fractions to pieces of 0.29 grams (one forty-eighth) and 0.15 grams (one ninety-sixth) bearing a lion's paw. Analysis of the metal is instructive: the alloy is around 54 per cent gold, 44 per cent silver and 1 to 2 per cent copper, but the coins were tariffed as if they were natural Lydian electrum of roughly 73 per cent gold. The issuer was, in other words, extracting a seigniorage from the gap between the metal's intrinsic content and its official rating.

That fact deserves emphasis, because it undermines the commodity story at its supposed strongest point. Even the very first coins were not simply weighed metal. They were metal whose accepted value depended on an institution's stamp rather than on its assay. The stamp is a record. Coinage did not replace institutional trust with intrinsic value; it packaged institutional trust into a portable, hard-to-counterfeit form.

■ Obligation on the modern balance sheet

The Bank of England's *Quarterly Bulletin* for 2014 Q1 carried the article "Money creation in the modern economy" by Michael McLeay, Amar Radia and Ryland Thomas. It is the single most useful primary source for a British reader on what money now is, and its central statements are worth quoting exactly.

On composition: "Broad money is made up of bank deposits — which are essentially IOUs from commercial banks to households and companies — and currency — mostly IOUs from the central bank."

On proportion: "Of the two types of broad money, bank deposits make up the vast majority — 97% of the amount currently in circulation."

On creation: "Whenever a bank makes a loan, it simultaneously creates a matching deposit in the borrower's bank account, thereby creating new money." And, correcting the textbook account: "Banks do not act simply as intermediaries, lending out deposits that savers place with them."

Set the physical stock against that. Bank of England banknote statistics record over 4.98 billion notes in circulation with a total value of about £91.5 billion as at 28 February 2026 — three per cent of the money in the economy, held in the form everyone pictures when the word "money" is said.

The consequence for a payments practitioner is structural. If 97 per cent of money is a liability recorded in commercial bank ledgers, then a payment between customers of two different banks is a message that causes two banks to adjust two sets of liabilities, plus a settlement between those banks in a third asset — central bank reserves — that neither customer ever sees. Every layer of the payments stack in this book exists to make that three-way adjustment happen correctly, in the right order, with an unambiguous point of finality.

■ What actually moves: the United Kingdom in numbers

Pay.UK's annual statistics for calendar year 2025 give the retail picture.

System	Volume, 2025	Value, 2025
Bacs (total)	6.86 billion items	£6.05 trillion
of which Direct Debit	5.03 billion items	£1.53 trillion
of which Direct Credit	1.83 billion items	balance of Bacs total
Faster Payments	5.55 billion items	£4.84 trillion
Cheque and credit (image clearing)	78.7 million items	£112.7 billion

The wholesale picture sits above it. The Bank of England reports that in 2024 CHAPS, the United Kingdom's high-value real-time gross settlement system, settled over £87 trillion, averaging over £344 billion each working day across around 208,000 payments a day. Roughly 208,000 instructions a day carry more value than 27 million Bacs items, which tells you most of what you need to know about the difference between retail and wholesale payment.

Now look at the shape of a single Bacs instruction, because it is the moment where the argument of this chapter becomes concrete. The Bacs Standard 18 format, the submission format used across the United Kingdom for Direct Debits and Direct Credits, is a fixed-width record of 100 characters, extended to 106 characters for multi-processing files.

Positions	Length	Field
1–6	6	Destination sorting code
7–14	8	Destination account number
15	1	Destination account type (zero filled)
16–17	2	Transaction code
18–23	6	Originating sorting code
24–31	8	Originating account number
32–35	4	Free format / RTI or PAYE indicator
36–46	11	Amount in pence
47–64	18	User's name
65–82	18	User's reference
83–100	18	Destination account name
101–106	6	Bacs processing day (multi-processing files only)

One hundred characters. Two account identifiers, a two-character verb, an amount in pence, and three eighteen-character strings that exist purely so that human beings on both sides can later agree about what the record meant. That is a payment. That is the whole of it.

The record then moves through the three-day cycle that Pay.UK's system principles set out: day one, input, when payment instructions and Direct Debit Instructions are submitted to the Bacs system; day

two, processing, when payment service providers process the items; day three, entry and settlement, when providers apply the payments to customer accounts and interbank settlement occurs. Input, processing and entry occur only on processing days; Saturdays, Sundays and English bank holidays are excluded. That three-day rhythm is not a technical limitation of any modern computer. It is a convention about when obligations become final, and conventions about finality are what payment systems fundamentally are.

The industry is now migrating that structure to a richer one. ISO 20022 replaces fixed-width, position-defined records with structured XML messages carrying a far larger and more explicitly typed set of fields. For cross-border payments, the coexistence period during which Swift's legacy MT messages and ISO 20022 both ran ended on 22 November 2025, after which MT is no longer supported for payment instructions in the live environment, and in-flow translation became a chargeable contingency from 1 January 2026. The direction of travel is towards more information per payment, not less — remittance data, structured parties, purpose codes, legal entity identifiers. If money were a substance, none of this would matter. Because money is a record, the richness of the record is the product.

■ Why this framing organises the whole book

Every question that follows in this book is a question about records.

Clearing is the process of working out, from two sets of records, what each party owes the other. Settlement is the moment at which an obligation is extinguished and the record becomes final. Reconciliation is the discipline of proving that two independently kept records of the same events agree — which is what the two halves of a tally stick did mechanically, in wood, in the twelfth century. Double-entry bookkeeping is the same insight applied inside a single organisation. A chargeback is a rule about when a settled record may be unwound and who bears the loss. Anti-money-laundering controls are rules about what the record must contain and who must be able to read it. A cryptographic ledger is a mechanism for making a shared record tamper-evident without a trusted keeper, which is to say a tally stick for people who cannot meet to fit the halves together.

The units and institutions change. The problem does not. Somebody must write down who owes what to whom, somebody must be trusted to keep that writing honest, and everybody must agree when a promise has stopped being a promise.

Volume I of this book asks what money is: obligation, denomination, settlement asset, legal tender, currency, and the several different things people mean when they say “cash”. The volumes that follow ask how the information moves: through schemes and rails, files and messages, banks and processors, borders and clearing houses, and the specific British machinery — Bacs, Faster Payments, CHAPS, the Image Clearing System, Open Banking — that most readers will actually be building against.

Start from the wrong premise and every one of those chapters becomes a list of arbitrary rules to memorise. Start from the right one and they become variations on a single, very old idea, already fully formed when somebody in Uruk pressed a number into wet clay to record how much beer was owed to a workman.

Money is information. This book is about how that information moves.

1.98 Common wrong ideas

1. **Wrong:** money was invented to solve the inconvenience of barter. **Right:** no society has ever been observed running its whole economy on direct exchange, and money has never been observed emerging from one; credit and reckoning come first in the evidence we actually have.

2. **Wrong:** barter is a myth. **Right:** barter is real and well documented, but the observed cases are overwhelmingly post-monetary, appearing where currency is scarce or a money economy has broken down, as among the Lhomi of north-eastern Nepal.
3. **Wrong:** the Mesopotamian tablets show neighbours writing IOUs to one another. **Right:** they are overwhelmingly institutional accounts of temple and palace estates, far closer to payroll and stock control than to a friendly overdraft.
4. **Wrong:** the token-to-envelope-to-tablet account of the origin of writing is settled fact. **Right:** it is the leading hypothesis and is contested in its details, and the broader claim that the earliest writing was administrative rests on the tablets themselves rather than on the token thesis.
5. **Wrong:** the tally stick was designed as a circulating currency. **Right:** it began as proof of payment into the Exchequer, and its negotiability was a consequence of being a hard-to-forge, institution-backed record of a Crown debt — a receipt in a strongbox, money when passed to a creditor.
6. **Wrong:** the notch conventions were uniform across the seven centuries the tally was in use. **Right:** the *Dialogus de Scaccario* and the 1911 *Encyclopædia Britannica* conflict, and do not even agree on which denominations existed, so no single table describes the whole period.
7. **Wrong:** money is information. **Right:** a record becomes money only when a specific, identifiable party is obliged to honour it and a mechanism makes that obligation stick; remove Ada's solvency, presence, honesty or answerability and the notebook is a list of numbers.
8. **Wrong:** cash is simply the physical version of a bank balance. **Right:** handing over a note completes the transaction at the instant of handover, with no clearing, no settlement window, no counterparty left owing anything and no possibility of recall, which is a property no electronic payment fully reproduces.
9. **Wrong:** "money is credit and nothing but credit" is the neutral summary of the historical evidence. **Right:** it is Mitchell-Innes's position in a long-running argument, and this book leans towards it because the record supports it and it is the more useful frame, but a lean is not a law.
10. **Wrong:** the three-day Bacs cycle is a technical limitation of the computers involved. **Right:** it is a convention about when obligations become final, and conventions about finality are what payment systems fundamentally are.

1.99 Chapter summary in 20 lines

1. Almost everyone describes a payment as a movement of money, and almost nobody is right.
2. What changes in a payment is a record: somebody's obligation to somebody else is written down, cancelled, netted off or handed on.
3. Ada's village notebook contains the whole of payment in miniature, because Tom's debt to Sam is discharged by three edited lines and nothing changes hands.
4. The textbook story that money was invented to fix the awkwardness of barter comes from Adam Smith in 1776 and is deductive rather than evidential.
5. Caroline Humphrey's careful 1985 statement is that no barter economy, pure and simple, has ever been described, let alone the emergence of money from one.
6. Barter itself is real, but observed barter is overwhelmingly a post-monetary phenomenon that co-exists with money rather than preceding it.
7. Societies without general-purpose money often kept goods in separate spheres of exchange, so money did not merely lubricate a market but changed what counted as one.
8. The earliest writing at Uruk, around 3100 BC, is administrative: rations, grain, sheep and beer, recorded with several distinct numerical notations.
9. Those tablets are the internal accounts of large temple and palace estates, and should be read as a state accounting department rather than as private credit.

10. Coins appear roughly two and a half thousand years after those tablets, in what is now western Turkey at the end of the seventh century BC.
11. Even the first coins were valued by the issuer's stamp rather than by assay, since electrum of about 54 per cent gold was tarified as though it held 73 per cent.
12. The English Exchequer tally makes the same point in wood: one stick split lengthways through every notch, verified by physical fit, unforgeable by the technology of the day.
13. Tallies circulated from hand to hand because the taxman would accept them, which made a notched stick money without making it a substance.
14. The longer half was the stock and the shorter the foil, and the careless disposal of two cart-loads of obsolete tallies burned down the Palace of Westminster in 1834.
15. A record is money only when an identifiable party is obliged to honour it and there is a mechanism that makes the obligation stick.
16. Every unit of money in a modern economy therefore has a counterparty: your bank for a deposit, the central bank for a note, the retailer for a gift card.
17. Cash is the genuine exception, because handover is final and leaves nobody owing anything, which is why it survives where records are unreliable.
18. In Britain the proportions are stark: about £91.5 billion of banknotes exist, against bank deposits making up 97 per cent of the money in circulation.
19. In 2025 Bacs and Faster Payments together carried over twelve billion instructions worth nearly eleven trillion pounds, and not one of them moved an object.
20. Every question in the rest of the book — clearing, settlement, reconciliation, chargebacks, finality — is a question about who keeps the record and when a promise stops being a promise.

Sources used: Caroline Humphrey, "Barter and Economic Disintegration", Man (NS) 20:1 (1985), 48–72; Adam Smith, The Wealth of Nations (1776), Book I Ch. IV; Paul Bohannan, American Anthropologist 57:1 (1955), 60–70; A. Mitchell Innes, "What is Money?", Banking Law Journal (May 1913); The Metropolitan Museum of Art, "The Origins of Writing"; The British Museum, object 1989,0130.4; Denise Schmandt-Besserat, Before Writing (1992); Michael Hudson on palatial credit and Mesopotamian interest; the laws of Hammurabi (Harper translation); Dialogus de Scaccario (c. 1179) and the 1911 Encyclopædia Britannica entry "Tally"; UK Parliament, "The Great Fire of 1834" and "Tally Sticks"; John H. Kroll, "The Coins of Sardis" (Sardis Expedition); Bank of England, "Money creation in the modern economy", Quarterly Bulletin 2014 Q1, and Banknote Statistics (28 February 2026); Bank of England, RTGS and CHAPS introduction (2024 figures); Pay.UK, Annual Summary of Payment Statistics 2025 and Bacs System Principles v18 (January 2026); Bacs Standard 18 Message Implementation Guide; Swift ISO 20022 implementation guidance.

A Balance Is Not a Thing

2.0 What this chapter gives you

1. You will be able to explain why the number in your banking app is a promise rather than a location, and why the envelope-in-a-drawer picture of a bank is the wrong one.
2. You will be able to state, in the Bank of England's own words, why bank lending creates deposits rather than deposits funding loans.
3. You will be able to say what genuinely constrains a bank's lending — capital, liquidity, funding cost and the price of money — and why no reserve multiplier appears on that list in the United Kingdom.
4. You will be able to distinguish a ledger balance from an available balance, and diagnose the customer who says a payment "took the money twice".
5. You will be able to cite *Foley v Hill* and explain why a depositor is an unsecured creditor of the bank rather than the beneficiary of a trust.
6. You will be able to state the FSCS limit precisely — £120,000 per eligible depositor per authorised firm for failures from 1 December 2025 — and explain why it attaches to the banking licence and not to the brand on the card.
7. You will be able to explain the difference between FSCS deposit protection and safeguarding under the Payment Services Regulations and Electronic Money Regulations, including the narrow case in which the FSCS looks through a safeguarding account.
8. You will be able to describe a bank run in terms of maturity transformation, sequential service and multiple equilibria, and say why deposit insurance works mainly by preventing runs rather than by paying for them.
9. You will be able to explain why the queue outside Northern Rock was an aftershock and why about \$42 billion left Silicon Valley Bank in a single day without anyone queuing at all.
10. You will be able to say what legal tender actually means, why it rarely comes up, and why Scottish and Northern Irish banknotes are not legal tender anywhere.

Open your banking app. There is a number at the top of the screen. Most people, if pressed, would say that number represents money that is somewhere — in a vault, in a safe, in a very large drawer with their name on it. Some part of the mind pictures a stack of notes, slightly dusty, waiting.

There is no stack. There is no drawer. There is no vault with your name on it, and there never was.

What there is instead is a record. Your bank has written down that it owes you a certain amount, and the number on your screen is a rendering of that record. It is not a description of where anything is. It is a description of a promise.

This is not a cynical observation and it is not a warning. It is simply the mechanical truth of how deposit banking works, and almost everything else in this book depends on the reader having absorbed it properly. Payments do not move substances around. They rewrite records. Until that is clear, the rest of it — clearing, settlement, netting, why a refund takes five days — will feel arbitrary. Once it is clear, most of it becomes obvious.

The plain version

Imagine you are going away for a fortnight and you ask your neighbour to look after £500 in cash.

There are two entirely different arrangements she might be agreeing to, and it matters enormously which one it is.

In the first, she takes your five hundred pounds, puts it in an envelope, writes your name on the envelope, and puts the envelope in a drawer. She does not touch it. When you come home she gives you back the same envelope with the same notes inside. If you asked to see it at three in the morning, it would be there. This is safekeeping. The money never stopped being yours; she was just holding it.

In the second, she takes your five hundred pounds and puts it in her purse. She spends some of it on shopping. She lends a bit to her son. She uses some to pay a plumber. When you come home, she gives you five hundred pounds — but not *your* five hundred pounds, because those particular notes are long gone, scattered across a supermarket, a plumber's till and her son's petrol tank. What she gives you is five hundred pounds of somebody else's money that she happened to have. What you had for two weeks was not an envelope. It was a promise.

Nearly everyone believes a bank is the first arrangement. Every bank in the world is the second one.

■ Following £500 into a branch

Walk into a high street branch on a Tuesday morning and hand over £500 in twenty-pound notes. Watch what actually happens.

The cashier counts the notes. She puts them in her till. At the end of the day the till is emptied, the notes are bundled, and eventually they go off to a cash centre in a van, where they are mixed in with the notes from every other branch. Nobody writes your name on them. Nobody could. There is no mechanism by which a particular note stays associated with a particular person, and there is no reason anyone would want one.

Two things get written down. The bank's own stock of banknotes goes up by £500 — those notes are now the bank's notes, in exactly the way the shopping money was your neighbour's. And, separately, the number beside your name in the bank's records goes up by £500.

That second number is not a location. It is an obligation. It says: *this institution owes this person five hundred pounds, payable on demand.*

Now the bank does what banks do. It has thousands of customers doing the same thing, and it knows from long experience that on any ordinary day most of them will leave most of their money where it is. So it keeps enough notes and enough instantly available funds to handle the ordinary comings and goings, and it puts the rest to work — lending it to somebody buying a house, buying government bonds, financing a business. That is the entire commercial logic of a bank. It is not a cloakroom. It is an intermediary that stands between people who want their money back at any moment and borrowers who want to keep it for twenty-five years.

■ The arithmetic that shows there is no vault

Here is the number that makes the point unarguable.

Every Bank of England note that exists anywhere in the world — in tills, in pockets, under mattresses, in cash centres, in shoeboxes in Spain — adds up to about **£91.5 billion**, spread across roughly **4.98 billion individual notes**. That is the total. Those are all of them.

And yet the money held in British bank accounts runs to trillions of pounds. The Bank of England's own explanation puts it plainly: about **96% of the money in the UK is bank deposits** and only about **4% is physical cash**.

So if every account holder in the country walked into a branch tomorrow and asked for notes, the notes would not exist. Not because anyone stole them. Not because of any wrongdoing at all. Because notes were never the plan. The overwhelming majority of money in Britain has never had a physical form and never will. It is entries in the records of banks.

■ So who makes sure you get your money back?

If a bank is a promise, the sensible next question is what happens when a bank cannot keep its promise.

Britain's answer is the **Financial Services Compensation Scheme**, the FSCS. It is a statutory scheme, funded by levies on the industry, and it does something very simple: if an authorised UK bank, building society or credit union fails, the FSCS pays your money back up to a limit, without you having to sue anybody or wait for the wreckage to be sorted out.

Since **1 December 2025** that limit has been **£120,000 per eligible depositor per authorised firm**. Before that date it was £85,000, a figure that had stood since 2017. The Prudential Regulation Authority — the arm of the Bank of England that supervises banks — consulted on raising it to £110,000 in March 2025 and then settled on £120,000 in its final rules, having recalculated against inflation up to September 2025.

The FSCS aims to pay most depositors **within seven days** of the failure, automatically, without a claim form. For certain large one-off sums — the proceeds of selling your home, an inheritance, a redundancy payment, an insurance payout — there is extra cover called a **temporary high balance**, currently up to **£1.4 million for six months** from the date the money reaches you.

If you have more than £120,000, the answer is not to worry about it but to spread it. Two banks, two limits. That is the whole strategy, and it works.

■ What a bank run actually is

Now put the two halves together, and the mechanics of a bank run fall out on their own.

A bank has promised everyone their money on demand. It cannot possibly have everyone's money available at once, because most of it is out on loan. This is not a scandal; it is the design, and it works perfectly well as long as people withdraw at the ordinary, boring rate.

A run is what happens when they stop. If enough customers try to take their money out at the same time, the bank has to sell things quickly to raise cash. Selling quickly means selling cheaply. Selling cheaply produces losses. Losses make the bank look weaker. Looking weaker frightens more customers, who withdraw, which forces more selling.

The vicious part is that running is rational. If you think there is a serious chance the bank will run out, the sensible thing to do is get in the queue early, because the money goes to whoever asks first.

Everybody knows that everybody knows this. A bank can therefore be destroyed by a belief, and the belief does not have to be true when it starts.

This is precisely why deposit protection exists. If you are certain the state-backed scheme will make you whole within a week, you have no reason to queue. Deposit insurance does not mainly work by paying people after a failure. It works by making the failure less likely in the first place, because it removes the incentive to run.

■ The exception that proves the rule

There is one form of money that really is a thing you hold: the note in your wallet.

A £20 note is not a record of what anyone owes you that has to be looked up. It is the claim itself, made physical. Hand it over and the payment is complete — instantly, finally, with no institution consulted, no message sent, no possibility of reversal, and no record that you were ever involved. There is nothing to reconcile because nothing was written down.

Everything else in this book is the elaborate machinery required to reproduce, over a network, what a banknote does for free when it changes hands. Seen that way, cash is not the primitive version of digital money. Digital money is a very complicated attempt to catch up with cash on the one dimension where cash is unbeatable: handing it over settles it.

Where the plain version stops being true

That account is good enough to explain at dinner and it is wrong in four specific places. The corrections matter, because each one is a place where confident, articulate people say things that are simply not so.

First: the bank does not lend out your deposit. That story is backwards.

The plain version said the bank keeps some of your £500 and lends the rest. Almost every school textbook says this, and the Bank of England has publicly said it is not how modern banking works. In *Money creation in the modern economy*, published in the Bank's Quarterly Bulletin in 2014, McLeay, Radia and Thomas put it flatly: "Rather than banks receiving deposits when households save and then lending them out, bank lending creates deposits." And: "Whenever a bank makes a loan, it simultaneously creates a matching deposit in the borrower's bank account, thereby creating new money."

Read that carefully, because it inverts the intuition. When a bank grants a £250,000 mortgage, it does not go looking for £250,000 of other people's savings to hand over. It writes two entries: a new asset (the borrower owes it £250,000) and a new liability (the borrower's account now shows £250,000). The deposit did not exist a moment earlier. Commercial banks create the great majority of the money in the economy, in the act of lending.

Which means the accompanying story — that banks are limited by a "reserve ratio", multiplying up a fixed base of central bank money — is also not how the constraint works. The United Kingdom has no mandatory minimum reserve ratio. What used to look like one, the Cash Ratio Deposit scheme under which larger institutions placed non-interest-bearing deposits at the Bank of England, was a funding mechanism for the Bank's own operations, not a lending constraint, and it was replaced by the Bank of England Levy on **1 March 2024**. Banks are constrained by capital requirements, by liquidity requirements, by what they can lend profitably, and by monetary policy acting on the price of money. Not by a multiplier.

Your £500 is still a real liability and the bank still has to be able to pay it. But the sentence "the bank lends out your money" should be retired.

Second: “a row in a database” is itself a simplification, and in a revealing direction.

The chapter’s own framing needs correcting. In a well-built banking system, your balance is usually not a stored number that gets edited. It is a *derived* figure, computed from an append-only sequence of postings — credits and debits, each immutable once written. You do not update a balance; you add an entry and the balance follows. This is not a stylistic preference. It is what makes a ledger auditable, and it is why the accounting principle in the next chapter is the single most load-bearing idea in the book.

There is a second, more practical consequence. There is rarely only one balance. Your **ledger balance** reflects postings that have been made. Your **available balance** reflects the ledger balance adjusted for holds — card authorisations that have been approved but not yet cleared, cheques not yet through, funds earmarked for a standing order due tomorrow. These two numbers are routinely different, and the difference is not an error. When a customer says a payment “took the money twice”, the overwhelmingly common explanation is that they are comparing an available balance against a ledger balance and seeing an authorisation hold that has not yet dropped off. Anyone building a payments product who models “balance” as a single mutable integer will discover this the expensive way.

Third: cash is not outside the system. It is a claim too — just a bearer one.

The plain version implied that a banknote is the one form of money that is not somebody’s promise. Look at the note. It says, in English, *I promise to pay the bearer on demand the sum of twenty pounds*, above the signature of the Bank’s Chief Cashier. That is not decorative. A Bank of England note is a liability of the Bank of England, and it appears as one on the Bank’s balance sheet, alongside the reserve accounts that commercial banks hold there.

What distinguishes a note is not that it is not a claim. It is that it is a claim which transfers by physical possession, against an issuer nobody expects to fail. Nobody has to be told, nobody has to agree, nothing has to be looked up. That is a different property from “not being a promise”, and confusing the two leads people to strange conclusions about what digital cash would have to be.

Scottish and Northern Irish notes make this sharper still. Those are the liabilities of commercial banks, not of the Bank of England, and the issuing banks are required to hold ring-fenced backing assets covering their entire note issue at all times.

Fourth: the run you can photograph is usually not the run that kills the bank.

The image everyone carries is the queue outside the branch. It is a nineteenth-century image and it has become, for the most part, theatre.

Northern Rock is the case in point. The queues that formed in September 2007 are the most reproduced photographs in modern British banking. But by the summer of that year, retail deposits made up only about **23% of Northern Rock’s liabilities**; the other roughly 77% was wholesale funding and securitised notes. The Bank for International Settlements published Hyun Song Shin’s analysis of exactly this, and his conclusion is worth quoting: “The Northern Rock depositor run, although dramatic on television, was an event in the aftermath of the liquidity crisis at Northern Rock, rather than the event that triggered its liquidity crisis.” The real run had already happened in the wholesale markets in August, when short-term lenders declined to roll their funding over.

Silicon Valley Bank in 2023 makes the modern version explicit. There were no queues. There was no branch network to queue at. Depositors withdrew about **\$42 billion in a single day, on Thursday 9 March 2023**, by wire transfer, largely because they had heard the same thing at the same time on the same channels. The bank was closed the following morning. A twenty-first-century run is an API call

executed by a lot of people at once, and it is faster than anything the queue-outside-the-branch mental model prepares you for.

So: the run is real, the mechanism in the plain version is right, and the picture is wrong. Runs happen at the speed of the fastest funding channel, and retail depositors are rarely the fastest channel any more.

The technical version

■ The deposit as a liability

A commercial bank's balance sheet is the cleanest way to see what your balance is.

On the asset side sit the things the bank owns or is owed: loans and mortgages advanced to customers, holdings of government and corporate securities, reserve balances held at the central bank, and physical cash. On the liability side sit the things the bank owes: customer deposits, wholesale borrowing, issued debt securities. The residual — assets minus liabilities — is equity.

Your current account balance is a liability of the bank. Specifically it is an unsecured claim against the bank, generally repayable on demand, ranking in an insolvency according to the statutory creditor hierarchy. It is not property held on trust for you, and it is not segregated.

English law has been unambiguous about this since **Foley v Hill (1848) 2 HLC 28**, decided in the House of Lords, which established that the relationship between banker and customer is one of debtor and creditor rather than trustee and beneficiary. Lord Cottenham LC:

Money, when paid into a bank, ceases altogether to be the money of the principal; it is then the money of the banker, who is bound to return an equivalent by paying a similar sum to that deposited with him when he is asked for it.

That is the legal statement of everything in this chapter, written 178 years ago and never seriously disturbed since. The practical consequence is that when a bank fails, depositors are creditors in an insolvency, which is precisely why a compensation scheme exists.

The same logic applies one tier up. The Bank of England's own consolidated balance sheet at **31 March 2025** showed, on the sterling liabilities side, **reserves balances of £694,512 million** and **notes in circulation of £91,230 million**, against total liabilities of £861,868 million. A banknote is a liability of the issuer, recorded as such.

■ What the record physically is

In a core banking system, an account is an entity with attributes — an identifier, a product type, a currency, a status, an owning customer, an interest and fee configuration — and a stream of postings against it.

A posting carries at minimum a value date, a booking date, an amount with an explicit currency, a direction (debit or credit), a reference to the transaction that produced it, and a link to the contra entry required by double entry. Postings are written once. Corrections are made by writing further postings, not by editing history: a mistaken credit is reversed with a compensating debit, and both remain visible. This is the same discipline that makes any well-built audit trail work, and it is the reason a bank statement is a list of movements rather than a series of snapshots.

The balance is a function of the postings. Real systems typically maintain running totals for performance — recomputing from the beginning of time on every enquiry would be absurd — but those

totals are caches, and reconciling them back to the posting stream is a routine control. If the derived total and the sum of postings ever disagree, the postings are right by definition.

Several distinct balances are maintained or derived, and the vocabulary is not perfectly standardised across the industry:

Balance	Roughly means	Typically differs because of
Ledger or book balance	Sum of all postings made to the account	Nothing; it is the accounting truth
Cleared balance	Postings that have completed clearing	Items in a clearing cycle not yet complete
Available balance	What the customer may actually spend	Authorisation holds, uncleared items, agreed overdraft, earmarks
Value-dated balance	Balance as at a future or past value date	Forward-dated instructions and back-value adjustments

Different institutions use these terms with slightly different boundaries, and international groups frequently use them inconsistently between systems. Where a specification matters — a reconciliation file, an API contract — the definition should be stated rather than assumed.

Two further properties of the record are worth naming now because they recur throughout the book. Amounts are held in the currency's minor units as exact integers, never as binary floating point, for reasons taken up in detail in Volume I's chapter on the money type. And every instruction that moves money needs an idempotency discipline, because a network that can deliver a message twice will eventually deliver a payment instruction twice, and the ledger must be able to recognise the second one as a duplicate rather than post it.

■ Two tiers of money

The phrase “your money” conceals a distinction that governs almost every settlement rule later in this book.

Commercial bank money is a deposit liability of a commercial bank. It is what households and businesses hold and what the number in your app represents. Its value depends on the solvency of that particular bank, which is why deposit protection exists and why the identity of the bank matters.

Central bank money is a liability of the central bank. In the United Kingdom it takes two forms: reserve balances held by eligible institutions in accounts at the Bank of England, and banknotes. Reserves are the settlement asset for the UK's high-value systems: they sit in the Bank's Real-Time Gross Settlement infrastructure, and a CHAPS payment is settled by moving reserves between accounts there. Settlement in central bank money is final in a way that settlement in commercial bank money is not, because the Bank of England does not fail.

That two-tier structure — customers holding claims on banks, banks holding claims on the central bank — is the skeleton of the whole payments system. When you pay someone at another bank, your bank's liability to you falls, their bank's liability to them rises, and the imbalance between the two banks is squared up in central bank money. Chapter 6 takes this apart properly; it is introduced here because “a balance is a liability of a specific institution” is the fact that makes it necessary.

■ Why the bank can lend, and what actually constrains it

Since bank lending creates deposits, the interesting question is not where the money comes from but what stops banks creating an unlimited amount of it. Four things, in practice.

Capital. Under the Basel III framework as implemented in the UK by the PRA, banks must hold regulatory capital against risk-weighted assets, plus buffers. Every new loan consumes capital. Capital is finite and expensive.

Liquidity. The Liquidity Coverage Ratio requires a bank to hold sufficient high-quality liquid assets to survive a 30-calendar-day stress scenario, with a minimum of 100%. The standard was phased in from 60% on 1 January 2015, rising ten percentage points a year to 100% on 1 January 2019. The LCR is the direct regulatory answer to the run mechanism described above: it is a rule about surviving a month of people wanting their money back.

Profitability and competition. A bank creates a deposit when it lends, but it cannot control where that deposit goes. The borrower spends it, and it very often ends up at another bank, which the lending bank must settle in central bank money. Lending therefore has a funding cost, and imprudent lending has a credit cost.

Monetary policy. Bank Rate sets the price at which reserves are remunerated and thereby the reference price of money throughout the system. The central bank influences the quantity of broad money by influencing its price, not by rationing a reserve base.

Structurally, the largest UK banking groups have also been required since **1 January 2019** to place their retail deposit-taking business inside a legally separate ring-fenced bank, insulated from investment banking activity.

■ Deposit protection: what the FSCS actually covers

Precision matters here more than anywhere else in the chapter, because this is the topic on which confident misinformation is most common.

The limit. £120,000 per eligible depositor, per authorised firm, for firm failures occurring on or after **1 December 2025**. This was set by the PRA in policy statement **PS24/25** (November 2025), following consultation paper **CP4/25** (March 2025), which had proposed £110,000. The previous limit of £85,000 had been in place since 2017.

Per licence, not per brand and not per account. The limit attaches to the authorised firm — the banking licence — not to the trading name on the card. The FSCS states the position directly: protection is “up to £120,000 per person per banking licence”. Two of its own worked examples: HSBC trades as HSBC and as first direct on a single licence, so the £120,000 covers the total across both; Nationwide Building Society’s licence also carries the Derbyshire, Cheshire and Dunfermline brands. A depositor holding £120,000 with each of two brands that share a licence is protected for £120,000 in total, not £240,000. The authoritative check is the Financial Services Register maintained by the FCA, not the marketing.

Joint accounts. Each account holder has their own entitlement, so a joint account is generally treated as split equally between the holders for the purposes of the limit.

Temporary high balances. Up to **£1.4 million for six months**, on top of the standard limit, where the money arises from a defined qualifying event. The FSCS lists these as including the sale of a main residence, property purchase or equity release, inheritance, insurance payouts, retirement benefits, redundancy, divorce or dissolution of a civil partnership, and compensation for unfair dismissal or

wrongful conviction. Protection runs from the point the money becomes legally transferable to the depositor or is first credited. Compensation for personal injury, disability or incapacity is protected without limit.

Timing. The FSCS works to a seven-day payout for the bulk of depositors, achieved by requiring failed firms to produce a **Single Customer View** file — a standardised extract of depositor identity, contact details, account details and aggregate balances — within 24 hours of the failure event. Note a genuine inconsistency in the published wording: FSCS industry materials refer to compensation “within seven working days of a deposit taker failing”, while consumer-facing materials say customers “will typically get their money back within seven days”. The operating requirement placed on firms is what matters in practice: the SCV must be ready before anything goes wrong, which is why SCV testing is a standing supervisory expectation and not a crisis activity.

Disclosure lag. Firms were given until **31 May 2026** to update their customer-facing information sheets and exclusion lists to the new limit. Documents published before that date, including some still in circulation, may quote £85,000 while the rule in force is £120,000. If you are checking a bank’s own PDF rather than the FSCS site, check its date.

What is not covered. The FSCS deposit protection scheme covers deposits with PRA-authorised deposit takers. It does not cover funds held with electronic money institutions or payment institutions. The FSCS is unambiguous: “We can’t protect the money you have with e-money institutions and payment providers.”

This is the single most consequential misunderstanding in UK consumer fintech, because e-money accounts look exactly like bank accounts. They have sort codes and account numbers, they receive salaries, they issue cards. What protects the customer’s money there is a different regime: **safeguarding** under the Payment Services Regulations 2017 and the Electronic Money Regulations 2011, which requires the firm either to segregate relevant funds in a designated safeguarding account with a credit institution or to cover them with an insurance policy or comparable guarantee. Safeguarding is a real protection and it is not nothing. But it is a claim on segregated assets resolved through an insolvency process, not a seven-day payout from a statutory scheme, and the FSCS warns that money held this way “could be tied up for a while during the insolvency process”.

There is one important bridge between the two regimes. Following PRA policy statement **PS2/23**, with rule amendments in force from **12 March 2023**, the FSCS can “look through” a safeguarding account and compensate the underlying customers of an e-money or payment institution if **the bank holding the safeguarded funds** fails. This does not extend protection to the case where the e-money institution itself fails and its bank does not. The distinction is narrow and it is exactly the sort of thing that gets flattened in marketing copy.

Covered deposits and bail-in. Deposits protected by the FSCS are legally excluded from the bail-in tool under the UK’s special resolution regime. When the Bank of England writes down or converts a failing bank’s liabilities, protected deposits are not in scope. This is the resolution-side counterpart of deposit insurance and it is worth knowing, because it means the answer to “could they take my savings to save the bank?” is no, within the protected limit.

■ The mechanics of a run, stated properly

The formal treatment is Diamond and Dybvig, “Bank Runs, Deposit Insurance, and Liquidity”, *Journal of Political Economy* volume 91, number 3 (1983), pages 401–419 — work for which Douglas Diamond and Philip Dybvig shared the 2022 Sveriges Riksbank Prize in Economic Sciences with Ben Bernanke.

The model isolates three features that together generate the run.

Maturity transformation. The bank issues liabilities that are short and liquid — demand deposits — against assets that are long and illiquid — mortgages, loans, held-to-maturity securities. That transformation is the economic service the bank provides, and it is inherently fragile.

The sequential service constraint. Depositors are served in the order they arrive, and the bank pays in full until it cannot pay at all. Withdrawing early is therefore strictly better than withdrawing late if you believe others will withdraw.

Multiple equilibria. Given the first two, the model has more than one self-consistent outcome. If nobody expects a run, nobody runs and the bank is fine. If everybody expects a run, running is individually rational and the bank fails — even if its assets, held to maturity, would have been worth more than its liabilities. Deposit insurance works by eliminating the bad equilibrium: if you will be paid regardless, you have no reason to be first.

The critical distinction the model forces is between **illiquidity** and **insolvency**. A bank is insolvent when its assets are worth less than its liabilities. A bank is illiquid when it cannot convert good assets into cash fast enough to meet demands. A run turns the second into the first, because forced sales into a falling market crystallise losses that a patient holder would never have taken.

The three cases already cited illustrate the range.

Case	What happened	Figures
Northern Rock, 2007	Wholesale funding withdrawn first; retail run followed the announcement of support	BoE liquidity support announced 14 September 2007; retail deposits about 23% of liabilities by summer 2007; wholesale liabilities fell from £26.7bn in June 2007 to £11.5bn in December 2007; HM Treasury guarantee announced 17 September, extended 20 September and again 9 October 2007
Silicon Valley Bank, 2023	Digital run at unprecedented speed; uninsured deposit base	About \$175bn total deposits at 31 December 2022, of which roughly \$151bn (about 95%) uninsured; about \$42bn withdrawn on 9 March 2023; closed 10 March 2023
Silicon Valley Bank UK, 2023	Resolved rather than liquidated	About £8.8bn total assets and £6.7bn deposits; sold to HSBC on 13 March 2023 using the Bank of England's private sector purchaser stabilisation power under the Banking Act 2009, with all depositors' money confirmed safe

The SVB UK case is the one practitioners should study, because it shows the modern default. The preferred outcome is not a payout at all. It is a resolution — a transfer of the failing bank's business to a purchaser over a weekend, so that on Monday morning the deposits still exist, still work, and simply have a different owner. Compensation schemes are the backstop for when that cannot be arranged, not the first line.

Note also the SVB deposit composition: about 95% uninsured. Deposit insurance calms retail depositors. A bank funded overwhelmingly by corporate treasurers holding tens of millions each has almost no insured base to calm, which is why concentration and insured share are supervisory concerns and not just balance-sheet trivia.

■ Cash, precisely

Physical currency is the exception, and being precise about it clarifies the rule.

Bank of England notes in circulation, from the Bank's own banknote statistics measured at the end of February each year:

Denomination	Value, £ millions (2026)	Volume, millions of notes (2026)
£5	1,993	399
£10	12,948	1,295
£20	58,633	2,932
£50	17,932	357
Total	91,505	4,984

A separate line in the same statistics, £4,704 million of "other notes", covers higher-denomination notes held as backing for the note issues of banks in Scotland and Northern Ireland — the internal high-value notes long known in the trade as giants and titans.

Legal tender is where public understanding is worst, and the Bank of England's own explainer is blunt that the term "has a narrow technical meaning that will rarely come up in everyday life". It concerns the discharge of a debt in court proceedings. In England and Wales, Royal Mint coins and Bank of England notes are legal tender. In Scotland and Northern Ireland, only Royal Mint coins are — banknotes issued by Scottish and Northern Irish banks are not legal tender anywhere, including Scotland and Northern Ireland, and circulate on acceptance rather than statute. Even the coins carry limits: 1p and 2p coins are legal tender only up to 20p, and 5p and 10p coins up to £5. No retailer is obliged to accept any particular means of payment, and debit cards, credit cards, cheques and contactless payments are not legal tender anywhere in the UK.

The Scottish and Northern Irish issuers operate under the Scottish and Northern Ireland Banknote Regulations 2009, supervised by the Bank of England. All such notes must be fully backed by ring-fenced assets at all times, with a minimum of 60% of those backing assets held in Bank of England notes or UK coin. A Scottish banknote is therefore a commercial bank liability with a statutory asset backing behind it — a genuinely different instrument from a Bank of England note, despite looking like the same kind of object in the same wallet.

What makes cash distinctive as a payment instrument is a short list of properties that the rest of this book spends thousands of words trying to reconstruct electronically. It is a bearer instrument, so possession is title. It settles finally at the moment of handover, with no clearing cycle and no settlement lag. It carries no counterparty risk beyond the issuer. It works with no connectivity, no power and no third party. It is irreversible: there is no chargeback on a twenty-pound note. And it leaves no record, which is simultaneously its greatest privacy virtue and the reason it attracts the attention of anti-money-laundering regimes.

Counterfeiting, incidentally, is not the systemic problem folklore suggests. In 2025 the Bank of England received about 200,000 counterfeit notes with a notional face value of £3.97 million, which it reports as fewer than one in 24,390 notes in circulation.

■ What this means if you are building something

Four practical consequences follow from this chapter, and they show up in real systems constantly.

Model balances as derived from an immutable posting stream, not as a mutable field. Any design in which a balance is updated in place will eventually produce a number that cannot be explained, and an unexplainable balance in a financial system is a regulatory problem as well as an engineering one.

Be explicit about which balance you mean, everywhere — in APIs, in reports, in support scripts, in the screen label the customer reads. “Balance” alone is not a specification.

Know exactly what protects your customers’ money and never let the language drift. If you are an e-money institution, your customers’ funds are safeguarded, not FSCS-protected, and saying otherwise in a marketing email is a regulatory failure and not a wording preference. If you are a bank, know your licence, know which brands share it, and know that the limit is £120,000 per depositor per licence.

And carry the central idea into everything that follows. A payment is never the transportation of a substance. It is the coordinated amendment of at least two records, held by at least two institutions, each of which must end up agreeing about what happened. The next chapter is about the discipline that makes that agreement possible, and it is seven hundred years old.

2.98 Common wrong ideas

1. **Wrong:** the bank keeps some of your deposit and lends the rest of it out. **Right:** bank lending creates deposits, because a new loan writes a new asset and a matching new liability, and the deposit did not exist a moment earlier.
2. **Wrong:** banks are limited by a reserve ratio that multiplies up a fixed base of central bank money. **Right:** the United Kingdom has no mandatory minimum reserve ratio, and the real constraints are capital, liquidity, profitability and monetary policy acting on the price of money.
3. **Wrong:** a shop has to take cash because cash is legal tender. **Right:** legal tender has a narrow technical meaning about discharging a debt in court proceedings, and no retailer is obliged to accept any particular means of payment.
4. **Wrong:** a balance is a number stored in a field and edited when money moves. **Right:** in a well-built system it is derived from an append-only sequence of immutable postings, which is what makes a ledger auditable and a statement a list of movements.
5. **Wrong:** the app showed two different figures, so the payment was taken twice. **Right:** the overwhelmingly common explanation is an authorisation hold that has not yet dropped off, seen by comparing an available balance against a ledger balance.
6. **Wrong:** a banknote is the one form of money that is nobody’s promise. **Right:** a Bank of England note is a liability of the Bank of England and is recorded as one; what distinguishes it is that the claim transfers by physical possession against an issuer nobody expects to fail.
7. **Wrong:** FSCS cover is per account, or per brand, so two accounts mean two limits. **Right:** it is £120,000 per eligible depositor per authorised firm, so brands sharing a licence — HSBC and first direct, for instance — share a single limit.
8. **Wrong:** money in an e-money or payment institution account is FSCS-protected, because it has a sort code, receives salaries and issues cards. **Right:** those funds are safeguarded under the Payment Services Regulations 2017 and the Electronic Money Regulations 2011, which is a claim on segregated assets resolved through insolvency rather than a seven-day statutory payout.
9. **Wrong:** a bank run looks like a queue outside a branch. **Right:** runs happen at the speed of the fastest funding channel, which at Northern Rock was the wholesale market in August 2007 and at Silicon Valley Bank was about \$42 billion of wire transfers on 9 March 2023.

10. **Wrong:** when a bank fails, depositors get compensated. **Right:** the preferred outcome is resolution — Silicon Valley Bank UK was sold to HSBC on 13 March 2023 — and a compensation payout is the backstop for when a transfer cannot be arranged.

2.99 Chapter summary in 20 lines

1. The number at the top of your banking app is not a description of where anything is; it is a description of a promise.
2. Handing £500 to a neighbour who seals it in a labelled envelope is safekeeping, and handing it to one who spends it and repays you later is deposit banking.
3. Every bank in the world is the second arrangement, and English law has said so since *Foley v Hill* in 1848.
4. Notes paid into a branch become the bank's notes, and the only thing written beside your name is an obligation payable on demand.
5. Bank deposits are about 96 to 97 per cent of the money in the United Kingdom, against roughly £91.5 billion of Bank of England notes in existence anywhere in the world.
6. If every account holder asked for notes tomorrow the notes would not exist, not because of any wrongdoing but because notes were never the plan.
7. The textbook claim that a bank lends out your deposit is backwards, because the Bank of England's own bulletin states that bank lending creates deposits.
8. Nothing multiplies up a fixed reserve base, and the binding constraints are capital, liquidity, funding cost and Bank Rate.
9. A balance is best modelled as a figure derived from an immutable posting stream rather than as a mutable field.
10. There is rarely only one balance: ledger, cleared, available and value-dated balances differ routinely, and the differences are not errors.
11. A banknote is a claim too, and what makes it special is that it transfers by possession rather than that it is not a promise.
12. Scottish and Northern Irish notes are commercial bank liabilities with statutory ring-fenced backing, and are not legal tender anywhere.
13. Legal tender is a narrow point about discharging debts in court and obliges nobody to accept anything in a shop.
14. A run follows from three features: liabilities that are short against assets that are long, service in order of arrival, and the rationality of going early.
15. Diamond and Dybvig showed that this produces more than one self-consistent outcome, and deposit insurance works by eliminating the bad one.
16. Northern Rock's real run happened in the wholesale market in August 2007, and the photographed queues came after the announcement of support.
17. Silicon Valley Bank lost about \$42 billion by wire in a day, with roughly 95 per cent of its deposits uninsured, which is what a run looks like when the branch is an API.
18. The FSCS covers £120,000 per eligible depositor per authorised firm from 1 December 2025, per licence rather than per brand, with temporary high balance cover of up to £1.4 million for six months.
19. E-money and payment institutions are outside the deposit scheme entirely, and their customers rely on safeguarding, which is a real protection but a slower and different one.
20. A payment is never the transportation of a substance; it is the coordinated amendment of at least two records held by at least two institutions that must end up agreeing.

Chapter sources: Bank of England — *Quarterly Bulletin 2014 Q1*, McLeay, Radia and Thomas, "Money creation

*in the modern economy”; “What is money” and “What is legal tender” explainers; banknote statistics (page updated 23 April 2026); consolidated balance sheet at 31 March 2025; Bank of England Levy pages; statement on Silicon Valley Bank, 13 March 2023; Scottish and Northern Ireland banknote regime documentation; Freedom of Information response on bail-in and customer deposits. Prudential Regulation Authority — CP4/25 and PS24/25, Depositor protection; PS2/23. Financial Services Compensation Scheme — deposit limit increase, banking licences, temporary high balances, Single Customer View, and e-money protection pages. Financial Conduct Authority — safeguarding requirements for payment and e-money institutions. Bank for International Settlements — Hyun Song Shin, “Reflections on Northern Rock”. Basel Committee on Banking Supervision — Liquidity Coverage Ratio executive summary. *Foley v Hill* (1848) 2 HLC 28. *Diamond and Dybvig*, *Journal of Political Economy* 91(3), 1983. US House Financial Services Committee background materials on the failure of Silicon Valley Bank. UK Parliament research briefing SN04478 on Northern Rock.*

Double Entry

3.0 What this chapter gives you

1. You will be able to state the rule of double entry in one sentence — never write down one end of a journey — and explain why a record-keeping method that checks itself survived seven hundred years without replacement.
2. You will be able to define debit and credit correctly as left and right, and explain why the same £3.20 is a debit in Lloyds' ledger and a credit in the cardholder's.
3. You will be able to post a VAT-inclusive £3.20 coffee sale as a three-legged entry, and say why the shop's cash rose by £3.20 while its equity rose by only £2.26.
4. You will be able to explain why “double” refers to the two sides rather than to a count of lines, and why a model hard-coded to one debit and one credit cannot express a partial refund.
5. You will be able to name the errors a trial balance cannot catch — omission, commission, principle, original entry, reversal and compensation — and explain why fraud generally balances beautifully.
6. You will be able to trace one card payment through the issuer's, acquirer's, merchant's and cardholder's books, and show the same interchange appearing as income in one and cost in another.
7. You will be able to explain why a card-accepting merchant books a receivable rather than cash, and why that gap has to be funded.
8. You will be able to explain why an authorisation posts nothing at all, why “pending” exists on the app, and why a refund is a fresh transaction rather than a cancellation.
9. You will be able to describe what happens at the apex of the system, where an interbank payment merely relabels which reserves account the Bank of England owes.
10. You will be able to list the invariants a ledger implementation must enforce: balance at write time, immutable append-only postings, signed integers in minor units, idempotency keys, multi-leg entries and named accounts for in-flight money.

There is a sentence that people say without noticing what it commits them to. They say: *I spent three pounds twenty on a coffee.*

Notice how many facts that sentence contains. It sounds like one. It is at least two, and if you are the coffee shop it is four or five. Three pounds twenty left something. Three pounds twenty arrived somewhere. The shop is three pounds twenty richer in one respect and slightly poorer in another, because a cup of coffee that cost it forty-one pence to make no longer exists. Some of that three pounds twenty was never the shop's money at all; fifty-three pence of it belongs to His Majesty's Revenue and Customs and the shop is merely holding it. None of this is philosophy. It is all recorded, in books, by people whose job it is, and the discipline they use to record it is seven hundred years old and has not needed replacing.

The discipline is called double entry, and it is the single most load-bearing idea in this book.

Everything that follows — clearing, settlement, netting, interchange, chargebacks, the reason a refund takes five days, the reason your available balance is not your ledger balance, the reason a bank can fail on a Friday and reopen on a Monday with every customer's number intact — is double entry operating at a scale and speed that would have astonished the Venetians who invented it. Not an analogy for double entry. Not something like double entry. It is double entry, running on faster machinery.

The previous chapter established that your balance is not a substance but a record. This chapter is about the rule that governs how such records may legally and sensibly be changed. Get this chapter and the rest of the book becomes a set of variations. Skip it and every later chapter will feel like a list of arbitrary industry customs.

The plain version

Every fact about money is a journey, and a journey has two ends.

If your friend hands you five pounds, that is not one thing that happened. It is two things that happened at the same instant: you have five pounds more, and your friend has five pounds less. If you write down only your half, your notebook is claiming that five pounds appeared out of nowhere. Notebooks that claim money appears out of nowhere are, sooner or later, wrong about everything.

So the rule is this: **never write down one end of a journey. Always write down both.**

That is the whole idea. Everything else is bookkeeping for how to do it tidily.

■ The two columns

Draw a line down the middle of a page. On the left you write where the money went *to*. On the right you write where it came *from*.

Your friend hands you five pounds:

Left: where it went to	Right: where it came from
My pocket, £5	My friend, £5

Two ends. Same amount on both sides. If the two sides do not match, you have made a mistake, and you know it immediately, without anybody having to tell you.

That last part is the reason the idea survived. It is not merely a tidy way to keep records. It is a record-keeping method that **checks itself**. If you add up everything in the left column and everything in the right column, the two totals must be identical. If they are not, something is wrong and you go and find it. A shopkeeper in 1350 with no calculator, no computer and no auditor could catch his own errors before they compounded. That is an enormous thing to be able to do.

■ Aisha's cake stall

Aisha is thirteen. She is running a cake stall at the school summer fair and she has decided to keep proper books, because her uncle is an accountant and has been unbearable on the subject.

She starts by putting in £20 of her own money, so there is a float. Then she spends £12 on flour, eggs, butter and icing sugar. She spends £3 on greaseproof paper and little labels. The school charges £5 to hire a table. On the day, she sells £34 of cakes.

Here is her book. Every row has a left and a right, and every row balances by itself.

What happened	Left (where it went to)	Right (where it came from)
Aisha puts in her own £20	Cash tin, £20	Owed back to Aisha, £20
Buys ingredients	Ingredients, £12	Cash tin, £12
Buys paper and labels	Packaging, £3	Cash tin, £3
Pays for the table	Table hire, £5	Cash tin, £5
Sells cakes	Cash tin, £34	Sales, £34

At the end of the fair, how much is in the tin? Add up everything that went into it and take away everything that came out: £20 in, then £12 out, £3 out, £5 out, then £34 in. That is £34.

Now check it from the other direction, which is the point of the whole exercise. The stall took £34 of sales and spent £20 getting there (£12 plus £3 plus £5), so it made £14 of profit. It also still owes Aisha the £20 float she put in. So the total claim on the stall is £20 plus £14, which is £34 — exactly the cash in the tin.

Two completely separate routes to the same number. That is not a coincidence and it is not luck. It is a consequence of never writing down half a journey. Aisha now knows, without checking anything else, that her book is internally consistent. If she had counted £31 in the tin, she would know instantly that either £3 had gone missing or she had forgotten to write something down, and she would go looking. Before double entry, she would simply have had £31 and a vague feeling.

■ The two words that confuse everybody

Accountants do not say “left” and “right”. They say **debit** and **credit**. Debit means the left-hand column. Credit means the right-hand column. That is all they mean.

Read that again, because almost everyone in Britain believes something else. Almost everyone believes that “debit” means money going out and “credit” means money coming in. That belief comes from bank statements, and bank statements are not lying to you, but they are written from a point of view you have probably never considered.

Look at Aisha’s book. When she sells £34 of cakes and the money goes *into* the tin, that is a **debit** to the cash tin. Money coming in is a debit. When she pays £5 for the table and money goes *out* of the tin, that is a **credit** to the cash tin. Money going out is a credit.

That is the exact opposite of what your bank statement says, and both are correct. We will resolve this properly in a moment. For now, hold on to the only definition that never changes: **debit is the left side, credit is the right side, and every entry needs both.**

■ The coffee, over the counter

Ama buys a flat white from Kalypso Coffee for £3.20 and pays with a £5 note.

In the shop’s book, that single sale is two rows, not one, because two separate journeys took place.

The first journey is the money. Three pounds twenty arrived in the till, and it arrived from a customer buying something. Left: cash, £3.20. Right: sales, £3.20.

The second journey is the coffee. A cup of coffee left the shop’s stock and went into a customer’s hand. The beans, milk, cup, lid and sleeve cost the shop 41 pence. Left: cost of what was sold, £0.41. Right: stock, £0.41.

Both rows balance. Together they say something a single row could never say: the shop is £3.20 richer in cash, £0.41 poorer in stock, and therefore £2.79 better off overall. A shop that only recorded “sales £3.20” would think it was doing brilliantly right up until the moment it noticed it had no beans.

■ The same coffee, paid by card

Now Ama taps her debit card instead.

No money moves. Nothing physical happens at all. What happens is that four different organisations each write down two things.

Ama’s bank writes down that it owes Ama £3.20 less than it did a moment ago, and that it now owes £3.20 to somebody else in the payments chain. Two facts. The shop’s card company writes down that it is owed £3.20 by that chain, and that it owes £3.20 to the shop. Two facts. The shop writes down that it has made a £3.20 sale, and that it is owed £3.20 by its card company. Two facts. And if Ama keeps a household budget, she writes down that she spent £3.20 on coffee and that her bank balance is £3.20 lower. Two facts.

Nobody moved anything. Everybody wrote two things down. That is what a payment is.

And here is the part that makes the whole system hang together: those separate books have to agree with each other. Ama’s bank thinks it now owes £3.20 into the system. Somebody else in the system thinks they are owed £3.20. Those two numbers must match, and at the end of the day somebody checks that they do. When they do not match, that is not a rounding annoyance. That is how fraud is caught, how outages are detected, and how a payments company discovers it has lost real money. There is a whole chapter about it in Volume V and it is titled, without affection, *Reconciliation*.

■ Where this came from

Nobody invented double entry in a single stroke. It grew up among Italian merchants in the thirteenth and fourteenth centuries, in exactly the places you would expect: Florence, Genoa, Venice, and the trading colonies those cities planted around the Mediterranean. The oldest surviving books that unambiguously use the full method belong to a Florentine firm, Giovanni Farolfi & Company, and were written by a partner named Amatino Manucci at the firm’s branch in Salon, in Provence, covering 1299 and 1300. They were identified and analysed by the accounting historian Geoffrey Lee in a 1977 paper in the *Accounting Historians Journal*. Manucci did not invent the technique. He was simply using it, competently, the way a clerk uses a spreadsheet.

The method spread by apprenticeship for two hundred years before anyone thought to write it down properly for strangers. The man who did was **Luca Pacioli**, a Franciscan friar and mathematician born around 1447 in Sansepolcro in Tuscany, who died on 19 June 1517. In November 1494 the Venetian printer Paganino de Paganini published Pacioli’s enormous *Summa de arithmetica, geometria, proportioni et proportionalità*, a general survey of everything then known about arithmetic, algebra and geometry, running to over six hundred pages of dense type.

Buried inside it, as the eleventh treatise of the ninth part, is a section called *Particularis de computis et scripturis* — “particulars of reckonings and writings”. In it Pacioli sets out what he calls the method of Venice: keep three books, a rough memorandum for jotting things down as they happen, a journal for writing them up in order, and a ledger where each account has its own page with debits on the left and credits on the right. Check the ledger periodically by listing every balance and proving that the two columns agree. He is widely quoted as advising that a merchant should not go to sleep at night until his debits equal his credits.

Pacioli did not invent double entry and never claimed to. What he did was put it on a printing press, at the moment printing presses existed, in the commercial capital of Europe. Three hundred years later Goethe had a character in *Wilhelm Meister's Apprenticeship* say of book-keeping by double entry that "it is among the finest inventions of the human mind". The method has been continuously in use ever since, which makes it older than the Bank of England, older than the pound sterling in its modern form, and older than every institution named anywhere in this book.

Where the plain version stops being true

The plain version above is a good foundation and it is not a lie. But four things in it will mislead you if you carry them forward unamended, and the fourth one is where fintechs lose money.

■ Correction one: "debit" does not mean money out

This is the most common misunderstanding in personal finance and it is entirely the fault of the banking industry, which chose to print the words on statements without explaining whose books they refer to.

Your bank statement says "debit" next to the £3.20 coffee. Aisha's cake stall book says a debit is money *arriving*. Both are correct because they are two different books written from two different points of view.

Recall the previous chapter: your deposit is not your money sitting in the bank's vault. It is the bank's **promise to pay you**. In the bank's own ledger, your current account is a *liability* — a debt the bank owes you. When you spend £3.20, the bank's debt to you gets smaller. Reducing a liability is, by the universal convention, a debit. So the bank debits your account, and prints the word on your statement.

Meanwhile, in your own books, your bank account is an *asset* — something you own. Spending reduces an asset. Reducing an asset is a credit. So in your books, the very same event is a credit.

The same £3.20, at the same instant, is a debit in Lloyds' ledger and a credit in yours. Neither is wrong. They are mirror images, and they must be mirror images, because the thing that grew smaller for one party is the thing that grew smaller for the other in the opposite direction.

The practical consequence is that **"debit" and "credit" are meaningless without two further pieces of information: whose book, and what kind of account**. Any sentence of the form "we'll credit your account" is ambiguous until you know who "we" are. When a merchant says "we will credit your card", they mean your card balance will move in your favour. When a bank's core system says "credit the customer account", it means increase the bank's liability, which is the same thing. When a payments engineer says "credit the merchant's payable", they mean the acquirer now owes the merchant more. Same word, three books, three referents. Systems that blur this produce reversed postings, which are the single most tedious class of bug in financial software.

■ Correction two: "double" does not mean two lines

The name is a historical accident and it misleads people who take it literally. "Double" refers to the two *sides*, not to a count of lines. A single journal entry may have three legs, or twelve, or two hundred. The rule is not "one debit and one credit". The rule is:

The sum of the debits equals the sum of the credits.

Return to the flat white. The shop is registered for VAT, and hot drinks in the United Kingdom are standard-rated at 20 per cent. Of the £3.20 the customer paid, £2.67 is the shop's revenue and 53 pence

is output tax that the shop is collecting on behalf of HMRC and will hand over later. That is one economic event and it needs three legs:

Account	Debit	Credit
Cash	£3.20	
Sales		£2.67
VAT payable		£0.53

Debits £3.20, credits £3.20. Balanced, with three lines.

Correspondingly, one real-world event may require several *separate* journal entries. The coffee sale generates one entry for the money and a second for the movement of stock into cost of sales. The card payment in the technical section below generates one entry when the sale is recorded, another when the acquirer settles, and a third when the monthly fee invoice is raised. Do not expect a one-to-one relationship between things that happen in the world and entries in a ledger. There isn't one, and building software on the assumption that there is will produce a system that cannot represent a partial refund.

■ **Correction three: balancing proves consistency, not truth**

This is the correction that matters most and that is stated least often.

A trial balance that agrees tells you exactly one thing: that for every debit posted, an equal credit was posted. It tells you nothing whatsoever about whether the entries were *correct*.

Post a £4,000 laptop purchase to “office refreshments” instead of “computer equipment”, and the trial balance agrees. Post an entire month's rent to nowhere at all — leave it out of the books completely — and the trial balance agrees. Post a sale as £810 when the invoice said £108, on both sides, and the trial balance agrees. Swap the debit and the credit so a payment out looks like a payment in, and the trial balance agrees. Make two independent mistakes of £20 that happen to fall on opposite sides, and the trial balance agrees.

Every one of those has a name in British accountancy training, and the whole class of them — errors of principle, omission, commission, original entry, reversal and compensation — is invisible to the balancing check. The technical section tabulates them.

This has a direct implication for anyone building a payment system. A balanced ledger is a *necessary* condition of correctness and nowhere near a sufficient one. Balance is cheap: any halfway competent implementation enforces it at write time and it will never fail again. The expensive question is whether the entries describe what actually happened, and that question is answered not by the ledger but by comparing the ledger against an independent record produced by somebody else. That comparison is reconciliation, and it is where the real work is.

Fraud, it should be said plainly, generally balances beautifully. A dishonest employee does not create unbalanced entries. They create balanced entries to the wrong account.

■ **Correction four: there is no world ledger, and the two halves are not simultaneous**

The plain version quietly implies a single book in which both ends of a journey get written. For a cake stall, that is true. For payments, it is not true and never has been.

There is no global ledger of money. There are millions of separate ledgers — one per bank, one per payment institution, one per scheme, one per merchant, one at the central bank — each of which

balances *within itself* and each of which is entirely capable of disagreeing with its neighbours. When Ama's bank debits her account and credits an amount payable into the card scheme, that is complete and balanced double entry *inside Ama's bank*. It says nothing about whether the acquirer at the other end has recorded the matching amount receivable, or recorded it for the right value, or recorded it at all.

The two halves are also frequently separated in time, sometimes by days. Ama's card was approved in about a second at 07:42 on a Tuesday. No posting was made at 07:42. Her bank placed a *hold*, which reduced her available balance but did not touch her ledger balance, and the actual double entry was made when the transaction cleared — typically the next business day for a UK card payment. That gap is not a defect. It is a designed consequence of authorisation and clearing being separate events, and it is the reason “pending” exists on your app.

Between the two halves, the money lives in accounts that exist purely to hold incomplete journeys: suspense accounts, clearing accounts, in-transit accounts. A large payment institution's suspense accounts are where its operational health can be read most honestly, because anything sitting in suspense is, by definition, a journey with only one end written down so far. A suspense balance that is growing rather than clearing is an incident in progress.

The technical version

■ The accounting equation

Everything in double entry descends from one identity. In the form used in the United Kingdom and internationally:

$$\text{Assets} = \text{Liabilities} + \text{Equity}$$

An entity's resources must equal the claims on those resources. The claims come from two sources: outsiders (liabilities) and owners (equity). This is not an empirical finding about businesses. It is true by construction, because equity is *defined* as the residual.

The IASB's *Conceptual Framework for Financial Reporting*, revised in March 2018, gives the definitions that UK and international reporting rests on. An asset is “a present economic resource controlled by the entity as a result of past events”. A liability is “a present obligation of the entity to transfer an economic resource as a result of past events”. Equity is “the residual interest in the assets of the entity after deducting all its liabilities”. Income is “increases in assets, or decreases in liabilities, that result in increases in equity, other than those relating to contributions from holders of equity claims”, and expenses are the mirror.

Because income and expenses are themselves changes in equity, the equation expands to cover trading:

$$\text{Assets} + \text{Expenses} = \text{Liabilities} + \text{Equity} + \text{Income}$$

Left-hand-side accounts have natural debit balances. Right-hand-side accounts have natural credit balances. Every rule you will ever be asked to memorise about debits and credits is a restatement of which side of that equation an account sits on.

■ Debits and credits, defined properly

A debit is an entry on the left of an account. A credit is an entry on the right. The words carry no directional meaning about cash. Their effect depends entirely on the account type:

Account type	Natural balance	A debit	A credit
Asset	Debit	Increases	Decreases
Expense	Debit	Increases	Decreases
Liability	Credit	Decreases	Increases
Equity / capital	Credit	Decreases	Increases
Income	Credit	Decreases	Increases

British bookkeeping students learn this as **DEAD CLIC**: Debits increase Expenses, Assets and Drawings; Credits increase Liabilities, Income and Capital.

The etymology is worth knowing because it explains the counter-intuitive part. “Debit” comes through Old French from Latin *debitum*, the past participle of *debere*, “to owe” — recorded in English from the mid-fifteenth century, and in its bookkeeping sense of an entry in an account by 1776. “Credit” arrives in the 1540s from Italian *credito*, from Latin *creditum*, “a loan, a thing entrusted”, from *credere*, “to trust”. The Italian merchant’s ledger was originally a list of people: those who owed him (*debitori*, on the left) and those he owed or who had entrusted him (*creditori*, on the right). Every other account type was later fitted into that same two-column frame. That is why the words feel wrong when applied to a cash account: they were never about cash in the first place.

■ Books of prime entry, ledgers and posting

The structure Pacioli described in 1494 is, with different names and vastly better tooling, the structure in use today.

A transaction is first captured in a **book of prime entry** — the day book, the sales day book, the cash book, or in a modern system the transaction log of the till, the gateway or the API. Nothing is classified yet; it is a chronological record that something happened.

It is then written as a **journal entry**: a dated, balanced set of debits and credits with a narrative and a reference. The journal entry is the atomic unit of accounting. It either exists in full or not at all.

The journal entry is **posted** to the **general ledger**, which is organised by account rather than by date. The list of available accounts is the **chart of accounts**, and its design is one of the more consequential decisions in a finance system: too coarse and you cannot answer questions, too fine and nobody classifies consistently.

High-volume accounts are broken out into **subsidiary ledgers** — one account per customer, per merchant, per cardholder — summarised into a single **control account** in the general ledger. The sub-ledger must sum to the control account. That check, run daily, is the first line of defence in every core banking system on earth, and its failure is a severity-one incident.

Finally, the **trial balance** lists every ledger account balance in two columns and proves the totals agree.

■ The trial balance and what it does not catch

Error	Description	Trial balance detects?
Single-sided entry	Only the debit or only the credit posted	Yes
Same-sided entry	Two debits or two credits posted	Yes
Unequal amounts	Debit and credit posted at different figures	Yes

Error	Description	Trial balance detects?
Casting error	An account balance added up wrongly	Yes
Extraction error	A balance transcribed wrongly into the trial balance	Yes
Omitted balance	An account left off the trial balance	Yes
Error of omission	Transaction never entered at all	No
Error of commission	Right class of account, wrong account within it	No
Error of principle	Wrong class of account entirely, e.g. capital expenditure posted as revenue expenditure	No
Error of original entry	Both sides posted at the same incorrect figure	No
Error of reversal	Debit and credit entered the wrong way round	No
Compensating error	Two independent errors that cancel out	No

Where a difference is found and cannot be resolved immediately, the difference is parked in a **suspense account** so that reporting can proceed, and cleared as the errors are identified. A suspense account is an admission written into the books; a suspense balance that persists across period ends is a finding, not a feature.

■ Worked example one: £3.20 over the counter

Kalypso Coffee, VAT-registered, sells a flat white for £3.20 including VAT at the UK standard rate of 20 per cent, paid in cash. Ingredient and packaging cost is 41 pence.

The VAT-exclusive value is $£3.20 \div 1.20 = £2.6667$, rounded to £2.67, leaving output tax of £0.53.

Journal 1 — the sale:

Account	Dr	Cr
Cash in hand (asset)	£3.20	
Sales (income)		£2.67
VAT payable (liability)		£0.53

Journal 2 — the cost of what was sold:

Account	Dr	Cr
Cost of sales (expense)	£0.41	
Inventory (asset)		£0.41

Gross margin on the cup is £2.26. Note that the shop's cash rose by £3.20 but its equity rose by only £2.26, because 53 pence was never its money and 41 pence of stock was consumed. A single-entry record showing “+£3.20” would overstate the shop's performance by 42 per cent.

■ Worked example two: the same £3.20 across two banks

Ama taps a contactless consumer debit card issued by Lloyds Bank. Kalypso Coffee's acquirer is Worldpay; the merchant's bank account is with NatWest; the card scheme is Visa. Ama's bank and the acquirer's settlement bank both hold accounts at the Bank of England.

For clarity the fee figures below are illustrative of a small UK merchant's contract rather than quoted from a published rate card. The one fee that is not illustrative is the interchange cap: under the Interchange Fee Regulation as applied in the United Kingdom, interchange on a consumer debit card transaction is capped at **0.2 per cent** of transaction value, and on a consumer credit card at **0.3 per cent**, for transactions where the merchant, the acquirer and the issuer are all in the UK. Those caps have applied since 9 December 2015 and remain in force as of August 2026. Suppose Worldpay charges Kalypso a blended merchant service charge of 1.5 per cent.

Interchange at 0.2 per cent of £3.20 is £0.0064 — sixty-four hundredths of a penny. Merchant service charge at 1.5 per cent is £0.048. Both are computed to sub-penny precision and rounded once, at the level of the daily batch, not per transaction. This is not fussiness; it is the reason Volume I closes with a chapter insisting that money must never be represented in floating point.

Issuer (Lloyds), on clearing:

Account	Dr	Cr
Ama — current account (liability to customer)	£3.2000	
Interchange income		£0.0064
Scheme settlement payable		£3.1936

Lloyds' liability to Ama falls by £3.20 — a debit, because reducing a liability is a debit, and this is the entry that prints the word "debit" on her statement. Lloyds retains interchange and owes the remainder into the scheme.

Acquirer (Worldpay), on clearing:

Account	Dr	Cr
Scheme settlement receivable	£3.1936	
Interchange expense	£0.0064	
Merchant payable — Kalypso Coffee		£3.2000

and then, for its own charge:

Account	Dr	Cr
Merchant payable — Kalypso Coffee	£0.0480	
Merchant service charge income		£0.0480

leaving £3.152 owed to the merchant. Note that interchange is *revenue* to the issuer and *cost* to the acquirer: the same 0.64 pence appears as a credit in one institution's books and a debit in another's. That is the four-party model in miniature, and Volume III takes it apart properly.

Merchant (Kalypso Coffee), at the till:

Account	Dr	Cr
Card receivable – Worldpay (asset)	£3.20	
Sales (income)		£2.67
VAT payable (liability)		£0.53

The shop has made a sale but holds no cash. It holds a receivable. This is the single most under-appreciated fact about card acceptance and the reason card-heavy businesses can be profitable and still fail: revenue is recognised at the till, cash arrives days later, and the gap must be funded.

Merchant, on funding the following business day:

Account	Dr	Cr
Bank – NatWest current account (asset)	£3.15	
Merchant service charge (expense)	£0.05	
Card receivable – Worldpay		£3.20

Cardholder (Ama), if she keeps personal books:

Account	Dr	Cr
Coffee and eating out (expense)	£3.20	
Lloyds current account (asset)		£3.20

Set Ama's entry beside Lloyds' entry and the whole chapter is visible in two rows. The identical event is a **debit** to the account in the bank's ledger and a **credit** to the account in hers. Neither party is confused. They are keeping mirror books, and mirrors reverse.

■ Where the central bank's books balance too

The entries above leave one obligation outstanding: Lloyds owes money into the scheme, and the acquirer's side is owed money out of it. That obligation is not discharged by the card network. Card networks calculate; they do not settle. Settlement happens when the participating banks move balances at a central bank.

In sterling, that means the Bank of England's Real-Time Gross Settlement system. Institutions hold either a **reserves account** — available to banks, building societies and certain investment firms, remunerated at Bank Rate as set by the Monetary Policy Committee — or a **settlement account**, available to institutions that do not qualify for reserves but meet the Bank's settlement criteria, including payment system operators and non-bank payment service providers.

The Bank supports several settlement models. CHAPS uses real-time gross settlement: each payment settles individually and irrevocably, debiting one participant's account and crediting another's, throughout the operating day. CHAPS is usually open from 06:00 to 18:00 Monday to Friday excluding England and Wales bank holidays; participants must be open to receive by 08:00 and to send by 10:00, customer payments must be submitted by 17:40 and the system closes to bank-to-bank payments at 18:00. There are over 35 direct participants, with several thousand further institutions reaching CHAPS indirectly. Retail systems, including card schemes, generally use deferred or prefunded net settlement instead, in which a day's obligations are netted down to a handful of positions and settled in cycles — the subject of the netting chapter later in this volume.

Whichever model is used, the entry at the Bank of England has the same shape:

Account	Dr	Cr
Reserves account — paying participant	net amount	
Reserves account — receiving participant		net amount

And here is the observation that ties the two-tier system together, which the chapter on central bank money will develop. Both of those accounts are **liabilities of the Bank of England**. When one commercial bank pays another, the Bank of England's total liabilities do not change by a penny. Nothing is created and nothing is destroyed. The Bank simply records that it owes a different set of people. An interbank payment, at the apex of the system, is a relabelling exercise — which is precisely what the previous chapter said money was.

■ The gap between authorisation and posting

Ama's card was approved in roughly a second. Nothing was posted to her account in that second.

Authorisation is a request for, and a promise of, funds. In double-entry terms it is not a transaction at all: no debit, no credit, no change to the ledger balance. What the issuer records is a **hold** or **shadow** entry, which reduces the *available* balance — the number the app shows — while the *ledger* balance, the accounting figure, is untouched. Only at clearing, when the acquirer submits the transaction for settlement, does a balanced journal entry come into existence.

This produces effects that customers experience as inexplicable and that are in fact perfectly determinate. A pending transaction can disappear without ever posting, because a hold expired and no clearing record arrived. A refund appears slowly because it is a fresh transaction travelling the same route in reverse, requiring its own clearing and settlement, rather than a cancellation of the original. Two identical taps a second apart can both be approved because authorisation is checked against available balance, not against a queue of intentions.

Every one of these is the same underlying fact: authorisation and posting are separate events, and only one of them is double entry.

■ Mirror accounts, and why reconciliation exists at all

Because each institution keeps its own books, the linkage between institutions is always a **pair of accounts that are supposed to be equal and opposite**. Lloyds' "scheme settlement payable" should correspond to a scheme receivable; Worldpay's "merchant payable — Kalypso" should correspond, to the penny, to Kalypso's "card receivable — Worldpay".

In correspondent banking the pattern has its own vocabulary. An account a bank holds *with* another bank, in that bank's books, is a **nostro** — from the Italian for "ours", as in "our account with you". The same account viewed from the other side is a **vostro** — "yours", as in "your account with us". Banks additionally maintain an internal **mirror nostro**: their own running record of what the correspondent's statement ought to say. The daily job of proving that the mirror matches the correspondent's statement is *nostro reconciliation*, and Volume IV returns to it.

The general principle holds everywhere. Two balanced ledgers can be individually perfect and jointly wrong. Only comparison detects that, and comparison is a separate discipline with a separate cost. A payment institution that has invested in a beautiful ledger and not in reconciliation has built half a system.

■ The balancing invariant in software

Several practical rules follow for anyone implementing this, and they are not stylistic preferences.

Enforce the invariant at write time. A ledger service should refuse to persist a journal entry whose debits do not equal its credits. The check belongs in the same transaction as the write, not in a nightly job. Once enforced, it never fails again, which is exactly the point: correctness that costs nothing to maintain.

Make postings immutable and append-only. Never update a posted entry. To correct an error, post a reversing entry and then post the correct one, leaving all three visible. This preserves the audit trail, makes any historical balance reproducible by replaying entries up to a point in time, and is the reason regulators can ask what a balance was on a given afternoon eighteen months ago and receive an answer.

Use signed integers in minor units. Amounts are held in pence, cents or paise as integers, with the currency and its exponent carried alongside per ISO 4217. Floating point must never touch a monetary amount. The final chapter of this volume is about nothing else.

Make posting idempotent. Every entry carries a client-supplied idempotency key. A retried API call after a timeout must not create a second, equally balanced, entirely spurious journal entry. Double-posting is the most expensive bug in payments precisely because the ledger will accept it without complaint: it balances.

Support multi-leg entries natively. A model that hard-codes exactly one debit account and one credit account cannot express VAT, fees, foreign exchange spreads or partial refunds without contortion. Purpose-built ledger databases — TigerBeetle is the best-known open-source example — take the opposite approach and expose only strict two-account transfers, composing complex movements from several of them. Either design is defensible; what is not defensible is a design in which balance is optional.

Model the incomplete journeys explicitly. Give in-flight money real accounts: authorisation holds, clearing accounts, unsettled scheme positions, unallocated receipts. Every penny in the business should be in a named account at all times, including the pennies whose ownership is currently undetermined. Systems that leave in-flight money implicit discover it during an incident, which is the worst possible moment.

■ What the law requires

Double entry is not merely good practice in the United Kingdom; the outcome it produces is a statutory duty.

Section 386 of the Companies Act 2006 provides that “every company must keep adequate accounting records”, and defines adequate as sufficient “to show and explain the company’s transactions” and “to disclose with reasonable accuracy, at any time, the financial position of the company at that time”. Subsection (3) requires “entries from day to day of all sums of money received and expended by the company” together with “a record of the assets and liabilities of the company”. Companies dealing in goods must additionally keep year-end stock statements and stocktaking records.

Section 388 requires those records to be kept at the registered office or such other place as the directors think fit, open to inspection by the company’s officers, and preserved for three years from the date they were made in the case of a private company and six years in the case of a public company.

The statute does not mandate double entry by name. It does not need to. A record that shows and explains transactions and discloses the financial position at any time is a double-entry record; no other method produces it. For regulated firms the requirements go considerably further — the FCA’s safeguarding rules for payment and e-money institutions impose specific internal and external reconciliation duties on top of the Companies Act baseline — and Volume V takes those up.

■ Why this is the load-bearing chapter

Every mechanism in the rest of this book is a special case of what is above.

Clearing is the exchange of the information needed to make matching entries in two institutions’ books. Settlement is the pair of entries at the central bank that discharges the resulting obligation. Netting is the arithmetic that reduces a million pairs of entries to one pair. Interchange is a credit in the issuer’s ledger and a debit in the acquirer’s. A chargeback is a reversing entry travelling the original route backwards. Reconciliation is the comparison of one institution’s balanced books against another’s. Safeguarding is the requirement that a specific pool of assets equal a specific pool of liabilities at all times. An outage is a period during which one side of a set of journeys was written and the other was not; the incident is the outage, and the aftermath is finding the unmatched halves.

There is no point in the entire apparatus at which anything is transported. There is only a very large number of books, each of which must balance internally, all of which must agree with each other, and none of which is allowed to record half a journey.

Pacioli’s advice was that a merchant should not sleep until the debits equal the credits. Modern payments run continuously and nobody sleeps at all, which is why the check has been moved into the machinery. But it is the same check, and it is asking the same question it asked in Venice in 1494: is there anything here that came from nowhere?

3.98 Common wrong ideas

1. **Wrong:** debit means money going out and credit means money coming in. **Right:** debit is the left-hand side and credit is the right-hand side, so money arriving in a cash tin is a debit to it and money leaving is a credit.
2. **Wrong:** your bank statement is written from your point of view. **Right:** it is written from the bank’s, where your account is a liability, so the bank reducing what it owes you is a debit and that is the word it prints.
3. **Wrong:** “we will credit your account” is an unambiguous instruction. **Right:** it means nothing until you know whose book and what class of account, which is why reversed postings are the most tedious class of bug in financial software.
4. **Wrong:** double entry means every transaction has exactly two lines. **Right:** it means the sum of the debits equals the sum of the credits, and one entry may have three legs, or twelve, or two hundred.
5. **Wrong:** one event in the world produces one entry in the ledger. **Right:** a single coffee sale produces one entry for the money and another for the stock, with further entries at settlement and at invoicing, and software built on a one-to-one assumption cannot represent a partial refund.
6. **Wrong:** a trial balance that agrees means the books are correct. **Right:** it means only that for every debit posted an equal credit was posted, and a whole month’s rent left out entirely leaves the totals untouched.
7. **Wrong:** fraud shows up as an imbalance. **Right:** a dishonest employee does not create unbalanced entries but balanced entries to the wrong account, which is why reconciliation exists as a discipline separate from bookkeeping.

8. **Wrong:** there is a single global ledger into which both halves of a payment are written. **Right:** there are millions of separate ledgers, each balancing within itself and each entirely capable of disagreeing with its neighbours.
9. **Wrong:** the two halves of a payment happen at the same moment. **Right:** the card was approved in about a second and the double entry was made at clearing, typically the next business day, with the money sitting in suspense and clearing accounts in between.
10. **Wrong:** a ledger that enforces balance is a correct ledger. **Right:** balance is cheap and never fails again once it is enforced at write time; the expensive question is whether the entries describe what actually happened, and only comparison against somebody else's record answers it.

3.99 Chapter summary in 20 lines

1. "I spent three pounds twenty on a coffee" sounds like one fact and is at least two, and for the shop it is four or five.
2. The rule of double entry is that you never write down one end of a journey, only both ends together.
3. Debit means the left-hand column and credit means the right-hand column, and the words mean nothing more than that.
4. Because the two column totals must agree, the method checks itself, which let a shopkeeper in 1350 catch his own errors before they compounded.
5. Aisha's cake stall reaches £34 by two entirely separate routes, and that is a consequence of the rule rather than luck.
6. Almost everyone in Britain believes debit means money out, because bank statements are written from the bank's point of view and never explain whose books they are.
7. In the bank's ledger your account is a liability, so spending reduces it, and reducing a liability is a debit.
8. In your own books that account is an asset, so the identical event is a credit: the two parties keep mirror books, and mirrors reverse.
9. "Double" refers to the two sides rather than to two lines, and the real rule is that the sum of the debits equals the sum of the credits.
10. A VAT-registered shop's £3.20 flat white needs three legs, because 53 pence of it was never the shop's money.
11. Some real-world events need several separate journal entries, so there is no one-to-one relationship between happenings and entries.
12. A trial balance that agrees proves internal consistency and says nothing whatever about whether the entries are true.
13. Errors of omission, commission, principle, original entry, reversal and compensation are all invisible to the balancing check.
14. Fraud generally balances beautifully, because the dishonest entry is balanced and merely posted to the wrong account.
15. There is no world ledger: millions of separate books each balance internally and can still disagree with one another.
16. Following one £3.20 card payment through issuer, acquirer, merchant and cardholder shows the same interchange as a credit in one institution's books and a debit in another's.
17. The merchant records a receivable rather than cash, which is why card-heavy businesses can be profitable and still fail on the funding gap.
18. Authorisation posts nothing and touches only the available balance, so pending items can vanish, refunds travel the whole route again, and two identical taps can both be approved.
19. At the apex, an interbank payment moves balances between two accounts that are both liabilities

of the Bank of England, so nothing is created and nothing is destroyed.

20. Clearing, settlement, netting, interchange, chargebacks, reconciliation and safeguarding are all special cases of the same refusal to record half a journey.

Chapter sources: Luca Pacioli, Summa de arithmetica, geometria, proportioni et proportionalità (Venice: Paganino de Paganini, November 1494), tractatus XI of distinctio IX, "Particularis de computis et scripturis"; Mathematical Association of America, Convergence, "Mathematical Treasure: Luca Pacioli's Summa"; Institute of Chartered Accountants in England and Wales, historical accounting literature collection, entries on Pacioli and the Summa; G. A. Lee, "The Coming of Age of Double Entry: The Giovanni Farolfi Ledger of 1299–1300", Accounting Historians Journal 4(2), 1977; Johann Wolfgang von Goethe, Wilhelm Meister's Apprenticeship, Book I, Carlyle translation. IFRS Foundation, Conceptual Framework for Financial Reporting, revised March 2018, project summary. Companies Act 2006, sections 386 and 388 (legislation.gov.uk). Bank of England, "A brief introduction to the Real-Time Gross Settlement system and CHAPS", and the CHAPS operating-hours pages. Payment Systems Regulator, guidance on the Interchange Fee Regulation. Association of Accounting Technicians, study guidance on identifying and correcting errors. Online Etymology Dictionary, entries for "debit" and "credit".

Ledgers, Journals and Accounts

4.0 What this chapter gives you

1. You will be able to explain why a bank's general ledger contains no account with your name on it, and where your own balance actually lives.
2. You will be able to describe the difference between the journal and the ledger — what happened against where we stand — and say why a business of any size needs both shapes of the same information.
3. You will be able to say what a chart of accounts is for, why it is designed backwards from the returns it must produce, and what its segments typically carry.
4. You will be able to separate recording, posting, settlement and valuation as four events with four timings, and use booking date and value date correctly.
5. You will be able to explain a card memo post: why the available balance drops, why the ledger balance does not, and why the amount can change before the real entry lands.
6. You will be able to answer “we missed the cut-off” with the only useful follow-up question, given that CHAPS, Bacs, the card schemes and the bank's own close are four different moments.
7. You will be able to give four independent reasons a bank must have an end of day, and describe the close-of-business sequence from stopping postings to rolling the business date.
8. You will be able to explain why immutability is a set of controls rather than a property of the technology, and what to ask a vendor who sells you an “immutable ledger”.
9. You will be able to distinguish a reversal from an adjustment, say when each is appropriate, and explain why amendment in place is never legitimate.
10. You will be able to explain why a payment often cannot be cancelled, using camt.056, camt.029 and pacs.004, or the ISO 8583 reversal messages, to show that not one of them is an edit.

The previous chapter established that every payment is two facts rather than one, and that the two facts must agree. This chapter is about where those facts are written down, in what order, under what names, and according to what rules — because a principle that is never implemented is just a slogan, and double entry is implemented with almost fanatical specificity.

The distinction matters more than it sounds. It is possible to understand double entry perfectly and still have no idea what a bank's books actually look like, in the same way that it is possible to understand arithmetic and have no idea how a payroll system is organised. The discipline is the easy part. The filing is the hard part, and the filing is where the interesting failures live.

There is also a piece of misinformation to clear away first. Most people, if they think about it at all, imagine that somewhere inside their bank there is a very large table with one row per customer, and that their balance is the number in that row. There is no such table. There is a system that holds customer balances, certainly, but it is not the bank's ledger — it is one of a dozen feeder systems

whose totals are summarised into the ledger. The ledger itself has no idea who you are. It knows only that the bank owes something in the order of fourteen billion pounds to a category of people called personal current account holders, and it knows that figure to the penny.

The rest of this chapter explains how that arrangement works, why it is arranged that way, what “end of day” means to an institution that never actually closes, and why the single most important rule in the whole apparatus is that you never rub anything out.

The plain version

Go back to Ada’s village shop from the first chapter, and imagine that it has done rather well.

Ada now has four shops, a delivery van, eleven staff, an accountant who visits on Thursdays, and about six hundred customers who run accounts with her. The notebook under the counter stopped working around the time she opened the second shop, for a reason that is entirely practical: a single notebook, written in the order things happen, is wonderful for answering the question *what happened today* and hopeless for answering the question *how much does Tom owe me*. To answer the second question from a notebook you would have to read every page since the beginning of time and add up all the Toms.

So Ada does what every business in history has eventually done. She splits her record-keeping into two shapes of the same information.

■ The diary and the folders

The first shape is a diary. Every time anything happens, someone writes it down in the order it happened, with the date, the amount, and what it was for. Nothing is ever skipped and nothing is ever written out of order. This is the record of events, and it is complete. In bookkeeping this is called the **journal**, and older British businesses called it the day book.

The second shape is a set of folders. There is a folder for Tom, a folder for Sam, a folder for the till, a folder for the van, a folder for wages, a folder for rent, a folder for everything Ada cares about separately. Whenever something goes into the diary, it also gets copied into the folders it affects — and because of double entry, it always affects exactly two of them at minimum, one on each side. This set of folders is called the **ledger**, each folder is called an **account**, and the act of copying an event from the diary into the folders is called **posting**.

The diary answers *what happened*. The folders answer *where do we stand*. They contain identical information arranged for two different questions, and a business needs both.

There is a third document, less glamorous and quietly essential: the list of which folders exist. Ada cannot have her Tuesday assistant inventing a folder called “bits and bobs” whenever something puzzling arrives, because in a year nobody will know what is in it. So there is a fixed, numbered, agreed list of every folder that may be used, and if something does not fit any of them, it goes into a holding folder and somebody senior sorts it out. That list is the **chart of accounts**, and the holding folder is the **suspense account**.

■ A day at Ada’s Bank

Now make Ada a banker, because the arithmetic is the same and the example is more useful.

Ada’s Bank opens on Monday morning with the following position. It has £40,000 of banknotes in its vault, £45,000 sitting in its account at the central bank, and £430,000 of loans it has made to people. Against that, it owes £500,000 to its depositors, and the £15,000 left over belongs to Ada. Notice that

the two sides already agree: £40,000 plus £45,000 plus £430,000 is £515,000, and £500,000 plus £15,000 is £515,000.

Four things happen during the day, and one thing happens because the day happened at all.

Tom brings in £300 in cash. The vault goes up by £300 and what the bank owes Tom goes up by £300. Two facts.

Priya pays Sam £120, both of them being Ada's customers. What the bank owes Priya falls by £120; what it owes Sam rises by £120. Two facts, and notice that from the bank's point of view nothing at all has changed in total. It owed £500,000 before and it owes £500,000 after. The money did not go anywhere. The label on it changed.

Sam takes £200 out of the cash machine. The vault falls by £200 and what the bank owes Sam falls by £200.

Ada charges Priya a £5 monthly account fee. What the bank owes Priya falls by £5, and the bank's income rises by £5. This is the first entry in the day that makes Ada better off rather than simply rearranging her obligations.

And then the day itself does something. Ada's £430,000 of loans earn interest at 5% a year, and one day has passed. Five per cent of £430,000 is £21,500 a year, which is £58.90 for the day. Nobody has paid it yet, and nobody will until the borrowers' payment dates, but Ada has earned it, so it goes in the books: the bank is owed another £58.90, and the bank's income rises by £58.90.

That last one is worth pausing on, because it is the entry that people find strangest and it explains a great deal about why banks have an "end of day" at all. Nothing happened. Nobody did anything. No customer touched anything. But time passed, and time passing is itself an economic event when you are in the business of lending money. Somebody or something has to notice that the day has ended and write down what the ending of the day earned. That is a job, it happens once per day, and it is one of the reasons the concept of a banking day exists.

■ Proving it adds up

At the end of Monday, Ada's accounts stand like this. On the left-hand side are the things the bank has and is owed. On the right-hand side are the things the bank owes and has earned.

Account	Left (debit)	Right (credit)
Cash in vault	40,100.00	
Balance at central bank	45,000.00	
Loans to customers	430,000.00	
Interest earned but not yet received	58.90	
Owed to depositors		500,095.00
Ada's own capital		15,000.00
Fee income		5.00
Interest income		58.90
Total	515,158.90	515,158.90

The two columns agree, to the penny. This list is called the **trial balance**, and producing it is the oldest sanity check in commerce. If the columns do not agree, something has gone wrong and you know it immediately, before it turns into a mystery.

Work through the arithmetic once and it stops feeling like magic. The vault is £40,000 plus Tom's £300 minus Sam's £200. Deposits are £500,000 plus Tom's £300 minus Sam's £200 minus Priya's £5 fee — Priya's payment to Sam nets to nothing because both are Ada's customers. Every single line came from an event that was written down twice, once on each side, and so the totals cannot help but match.

■ The rule about the eraser

Now the most important part of this chapter, and the part that governs everything that comes later in this book.

On Tuesday, Ada's assistant discovers that £250 which arrived on Monday for a customer called Priya Shah was credited to a different customer, also called P. Shah. It is an honest mistake and it needs fixing.

The wrong way to fix it — the way that feels natural, and the way every spreadsheet in the world invites you to — is to find Monday's entry, cross out the wrong name, and write the right one. Two seconds, problem solved, books still balance.

Ada's books do not permit this, and no proper set of books anywhere has permitted it for five hundred years. What actually happens is that three separate things are written down, all dated Tuesday, none of them touching Monday's page.

First, an entry that is the exact mirror image of Monday's mistake: take £250 back off the wrong P. Shah and put it back where it came from. This is called a **reversing entry**, and it does not make Monday disappear. Monday's entry is still there, still visible, still wrong. It has simply been cancelled out by an equal and opposite entry the following day.

Second, the correct entry: take the £250 out of the holding account and credit it to Priya Shah, where it should have gone in the first place.

Third — and this is the part that separates a real ledger from a tidy one — a note, attached to both of the above, saying who found the error, who authorised the correction, when, and why.

Six lines of writing where an eraser would have taken none. This looks like bureaucratic self-punishment until you consider what the eraser costs you. Once erasing is permitted, no statement of the books is ever evidence of anything. If Monday's page can be quietly changed on Tuesday, then Monday's page proves nothing about Monday, and neither does any other page. The auditor cannot check the books, because the books are whatever the last person to touch them says they are. The customer cannot dispute a figure, because there is no fixed figure to dispute. And someone stealing money can hide it perfectly by simply amending history.

The reversing entry is therefore not an accounting formality. It is the mechanism by which a set of books becomes evidence rather than opinion. And once you have seen it, you will see it everywhere in payments, because the same logic applies at every level of the system. A refund is not the deletion of a purchase; it is a second, opposite transaction. A cancelled bank transfer is not an unsent message; it is a new message asking the other bank to send the money back. Nothing is ever undone. Things are only ever answered.

Where the plain version stops being true

The shop is a good model and it will get you through a dinner party. Four things about it are wrong in ways that matter.

First: there is no folder with your name on it in the general ledger. There never was.

The plain version implied a single ledger containing an account for each customer. Real institutions do not work this way and could not. A bank with eight million current accounts does not maintain eight million accounts in its general ledger; the general ledger might contain a few thousand accounts in total, and one of them will be something like “personal current accounts, sterling, retail division” holding a single aggregate balance of several billion pounds.

Your individual balance lives in a **sub-ledger** — in this case the deposit system, which is part of what the industry calls the core banking platform. The sub-ledger holds the detail: your account, your balance, your transactions, your overdraft limit. The general ledger holds the summary, in a **control account**, and the two are reconciled, usually daily. If the deposit sub-ledger says the bank owes £14,208,331,402.71 across all personal current accounts and the general ledger control account says £14,208,331,398.44, there is a problem, it is a serious one, and somebody’s afternoon is now about finding £4.27.

This is not merely a technical arrangement; it changes what questions each system can answer. The general ledger cannot tell you anything about you. It can tell you what the bank owes in total, by product, by currency, by legal entity, by branch. That is what it is for, because that is what regulatory returns, statutory accounts and capital calculations are built from. Confusing the two leads people to make wrong assumptions about what a bank can look up and how fast.

Second: recording, posting, settling and valuing are four different events, and they routinely happen on different days.

The plain version used “write it down” and “post it” almost interchangeably. In a real system they are distinct stages with distinct timestamps and distinct meanings, and the gaps between them explain a large fraction of the questions customers ask.

An entry is **recorded** when it is captured into the system. At that stage it may exist in an unposted state — validated or not, awaiting authorisation, awaiting a matching leg. It is **posted** when it is committed to the account and actually changes the balance. It carries a **booking date**, being the day the bank’s books recognise it, and a **value date**, being the day from which it counts for interest and availability. Those two dates are frequently different. A cheque paid in on Friday may be booked on Friday and valued on the following Wednesday. A payment made on a Saturday may be booked with Monday’s date because Saturday is not a business day in the bank’s accounting calendar.

And **settlement** — the moment the two banks are actually square with each other, usually across accounts at the central bank — is a fourth event with its own timing, which frequently happens after the customer has seen the money. Chapter seven is entirely about the difference between clearing and settlement, and the reason it needs a chapter is that the two are conflated constantly, including inside banks.

The most visible consequence of this is the card authorisation. When you tap for £43.20 in a supermarket, no accounting entry is posted anywhere. What happens is a **memo post**: your available balance drops by £43.20 while your ledger balance is untouched, because no double entry has occurred. The real posting happens when the transaction clears, typically one to three business days later, at which point the memo hold is released and a genuine entry lands. This is why the amount can change between the two — a restaurant tip, a fuel pump, a hotel — and why the date on your statement is sometimes not the date you were in the shop.

Third: “end of day” is not a moment when the bank stops, and there is more than one of them.

The shop closes at six. A bank does not close at all — Faster Payments has run around the clock since 2008, cash machines never sleep, and card authorisations arrive at four in the morning from customers in other time zones.

What a bank has instead is a **business date**, which is a label rather than a wall clock. Postings are stamped with a business date, and at some point the system performs a set of end-of-day procedures and rolls the business date forward. Transactions arriving after the roll get tomorrow's date even though it is still today outside. The batch is a logical boundary, not a physical shutdown, and for a bank operating in several countries there may be several business dates alive at once across different books.

Worse, there is no single cut-off, because the bank sits inside other people's calendars as well as its own. In the United Kingdom, CHAPS opens at 6am and closes at 6pm on working days, with customer payments required by 5.40pm. Bacs input files must reach the scheme between 07.00 and 22.30 on input day. Card schemes have their own daily cut-offs which do not align with either. Each of those is somebody's end of day, none of them is the same moment, and the bank's own accounting close is yet another. When somebody in payments says "we missed the cut-off", the useful follow-up question is always *whose*.

Fourth: immutability is a set of controls, not a property of the technology.

The plain version said you never rub anything out. That is the rule. It is not, however, a law of physics.

Underneath almost every general ledger in the world sits a perfectly ordinary relational database in which any sufficiently privileged person can execute an update statement against a posted journal line. What prevents it is not that the operation is impossible but that it is prohibited, monitored, and detectable — application designs that only ever append, database permissions that deny direct write access to production financial tables, audit logging that records every access, hash-chaining or sequence numbering that makes gaps visible, and separation of duties so that the person who creates a journal entry is never the person who approves it.

This matters practically, because software is sold on the claim of being an "immutable ledger" when what is meant is that the vendor's own screens have no delete button. That is a real and useful control, but it is a control, not a guarantee, and anyone assessing a system should ask what stops a database administrator, what stops a support engineer with an emergency fix script, and what would make either visible afterwards. The answer is usually a combination of restricted access, mandatory approval workflow, and independent reconciliation — and if the answer is "nobody would do that", the system does not have immutability, it has trust.

There is also a narrower correction. Reversal is not the only permitted repair. A **reversing entry** restores the position exactly and then the correct entry is made afresh, which is preferred because it preserves the gross record of what was originally done. An **adjusting entry** posts only the net difference, which is shorter but leaves the original figure standing uncorrected in the account. Both are legitimate, both are used, and the choice depends on materiality, on whether the accounting period has closed, and on the institution's own policy. What is never legitimate is amendment in place.

The technical version

■ Accounts and their classes

An account is a classification bucket with a running balance. Every account belongs to one of five classes, and the class determines which side of the entry increases it.

Class	Increased by	Example in a bank
Asset	Debit	Loans and advances to customers; balances with central banks; nostro accounts
Liability	Credit	Customer deposits; debt securities issued; accrued charges
Equity	Credit	Called-up share capital; retained earnings; revaluation reserves
Income	Credit	Interest income; fee and commission income; interchange income
Expense	Debit	Interest expense; staff costs; scheme fees; impairment charges

The perpetual confusion for non-accountants is that a customer deposit is a liability, so the bank *credits* your account when it owes you more. Your statement is written from the bank's point of view, which is why a credit is money in and a debit is money out, and why the words appear reversed relative to the bank's own internal usage in ordinary conversation. Chapter ten returns to this, at length, because it is the single most common source of misreading a statement.

The fundamental identity that all of this enforces is that assets equal liabilities plus equity, extended within a period to include income and expense, which are closed to retained earnings at period end.

■ The chart of accounts

The chart of accounts is the controlled list of every account that may be posted to, together with its class, its posting rules and its reporting mappings. In an institution of any size it is not a flat list of numbers but a structured code, typically composed of independent segments.

A realistic account code in a bank might carry a legal entity identifier, a natural account, a currency, a product or portfolio code, a cost or profit centre, and an intercompany counterparty segment. The natural account alone might be six or seven digits; the full posting string might be thirty characters or more. This structure exists because the same economic fact needs to be sliced several ways: interest income needs to be reportable by product for management, by entity for statutory accounts, by currency for treasury, and by counterparty type for regulatory returns, and it is far cheaper to encode all of that on the entry than to reconstruct it later.

The chart is not designed in isolation. It is designed backwards from the reports it must produce. In the United Kingdom those include the company's statutory accounts under the Companies Act 2006 and the applicable accounting framework, the prudential and financial reporting returns required by the Prudential Regulation Authority, and the statistical returns required by the Bank of England. In the European Union the equivalent frameworks are the European Banking Authority's COREP, covering own funds and prudential metrics, and FINREP, covering financial information anchored to the accounting figures. A chart of accounts that cannot be mapped cleanly to the required returns is not a stylistic problem; it is an annual manual reconciliation exercise for the rest of the institution's life.

Multi-currency handling adds a further layer. Accounts are held in a transaction currency and translated into the entity's functional currency for reporting, with periodic revaluation of monetary balances and the resulting gains or losses posted to defined revaluation accounts. The revaluation run is another of the jobs that must happen at a defined point in the day, using a defined rate, from a defined source.

■ Journals, journal entries and journal lines

A journal is the chronological record. A journal entry is the atomic unit of posting: a header describing the event and two or more lines describing its effects, whose debits and credits must sum to zero before the entry is accepted. The system rejects an unbalanced entry outright, which is the single most valuable structural property of double entry as implemented in software.

A minimally serious journal entry carries, at header level, a unique identifier, the source system, the journal category or type, the business date, the currency, the description or narrative, the creating user or process, the approving user, the approval timestamp, and a reference to the originating business object such as a payment identifier or a card transaction identifier. At line level it carries the full account code string, the debit or credit indicator, the transaction amount, the functional-currency amount and the rate used, the value date if it differs from the business date, and a line-level narrative.

Journals divide broadly into automatic and manual. Automatic journals are generated by feeder systems — deposits, lending, cards, treasury, payments — from business events, according to fixed accounting rules held in an accounting rules engine that maps event types to debit and credit account pairs. Manual journals are typed in by the finance function to record accruals, provisions, reclassifications and corrections. Manual journals are the ones auditors care about, because they are the ones a human chose. International Standard on Auditing 240, which governs the auditor's responsibilities relating to fraud, requires auditors to test the appropriateness of journal entries recorded in the general ledger — with particular attention to entries made outside the normal course of business, entries posted at period end, and entries made by people who do not usually post journals.

Recurring journals — accruals, prepayments, depreciation — are templated and generated on a schedule. Reversing journals are a special case worth naming: an accrual posted at month end is very often flagged to reverse automatically on the first day of the following month, so that when the actual invoice arrives it can be posted gross without double counting. This is an entirely routine use of reversal that has nothing to do with error correction.

■ Sub-ledgers and control accounts

The general ledger is fed by sub-ledgers, each of which holds transaction-level detail for one domain and reports a summary into a control account.

Sub-ledger	Holds	Typical control account
Deposits (core banking)	Individual customer current and savings accounts	Customer deposits, by product and currency
Lending	Individual loan accounts, balances, accrued interest, arrears	Loans and advances to customers
Cards	Cardholder balances, authorisations, presentments, fees	Card receivables; merchant payables
Nostro	The bank's own accounts held at correspondent banks	Balances with other banks

Sub-ledger	Holds	Typical control account
Fixed assets	Individual assets, cost, depreciation, disposals	Property, plant and equipment
Payables and receivables	Supplier invoices and customer invoices	Trade creditors; trade debtors

British practice historically named these the sales ledger and the purchase ledger, with the general ledger itself called the nominal ledger — terminology still common in UK accounting software and worth recognising when it appears.

The **sub-ledger to general ledger reconciliation** is the control that binds the two together, and in a bank it is performed daily rather than monthly. Its logic is simple: the sum of all detail balances in the sub-ledger must equal the balance of the corresponding control account. Any difference is a break, and breaks are aged, categorised and escalated. Where an item cannot be resolved immediately it is parked in a **suspense account** — a temporary general ledger account for items that are known to have occurred but are not yet correctly classified. A suspense account that is cleared daily is housekeeping. A suspense account with a growing balance and items aged over ninety days is one of the clearest early indicators that an institution's operations are failing, and it is precisely what supervisors and auditors look for.

■ Posting versus recording, precisely

The lifecycle of an entry runs: capture, validate, authorise, post, report. Only the fourth step changes a balance.

The distinction is visible on the wire in ISO 20022 cash management reporting. The message camt.053, formally BankToCustomerStatement, is the end-of-day statement covering a completed business day. Each entry within it carries a status; the value BOOK indicates that the entry has been booked to the account, as opposed to being merely pending. Each entry also carries BookgDt, the booking date, and ValDt, the value date, as separate elements — the standard treats them as two different facts about the same entry precisely because they routinely differ. Statement-level balances are typed by code, with OPBD for the opening booked balance and CLBD for the closing booked balance, and separate codes exist for available balances, which is the standard's way of making explicit the ledger-versus-available distinction introduced in chapter two.

The companion messages complete the picture. Camt.052, BankToCustomerAccountReport, carries intraday reporting, including items that are not yet booked. Camt.054, BankToCustomerDebitCreditNotification, notifies the account owner of individual debit or credit entries as they occur. In the older SWIFT FIN world the equivalents are MT940 for the end-of-day customer statement and MT942 for the interim transaction report, in which the same information is encoded positionally inside fixed-format tags — a Tag 61 statement line packs the value date and entry date into the first ten characters, which is compact, unforgiving, and the reason the industry is migrating.

■ The trial balance, and what it does not prove

The trial balance is the list of every account's closing balance with debits in one column and credits in the other. It is produced as a matter of routine, at least monthly and in banks daily, and its agreement is a necessary condition for the books being correct.

It is emphatically not a sufficient one. A trial balance will balance perfectly in the presence of at least four distinct classes of error. An entry omitted entirely, both legs missing, leaves the totals untouched.

An entry posted to the wrong account of the same class — Priya's £250 credited to the wrong customer — leaves the totals untouched, which is exactly the error the plain version used. An entry posted twice in full leaves the totals untouched. And compensating errors, where two mistakes happen to offset, leave the totals untouched. A balanced trial balance means the arithmetic is consistent. It says nothing whatever about whether the entries describe reality.

This is why reconciliation exists as a separate discipline from bookkeeping, and why chapter six of volume five is devoted to it. Balancing proves internal consistency. Reconciliation proves agreement with an external record — the counterparty bank's statement, the scheme's settlement file, the central bank's account. The two controls catch different things and neither substitutes for the other.

■ End of day, and why batch processing exists at all

Modern engineers, raised on event streams and continuous deployment, tend to treat batch processing as a relic. It is not. There are at least four independent reasons a bank must have a defined daily boundary, and none of them is about the age of the software.

The first is that some financial facts are defined per day and cannot exist without a day. Interest accrual is the obvious case: interest accrues on a balance for a period, and the period has to end for the accrual to be computed. So do daily fee applications, daily limit resets, daily average balance calculations, and the ageing of arrears. A system with no end of day cannot compute these because the input to the calculation does not exist until the day is closed.

The second is that the bank sits inside external cycles that are batch by design. The Bacs three-day cycle, run by Pay.UK, is the clearest example in Britain: files are submitted on input day, with transmission to Bacs permitted between 07.00 and 22.30; they are processed and distributed to the receiving payment service providers on processing day; and on entry day, day three, the payments are simultaneously credited to the recipients' accounts and debited from the originator's. That structure is not negotiable by an individual bank, and a bank's own processing has to align to it.

The third is settlement. Obligations between banks are settled at the Bank of England across settlement accounts in the RTGS service, and the settlement arrangements are themselves periodic. CHAPS settles individually and in real time on a gross basis throughout the day, between 6am and 6pm. Bacs settles on a net basis once each business day. The Faster Payments Service, despite running around the clock for customers, settles net three times each business day. The Image Clearing System for cheques runs a two-day cycle and settles net once each business day. LINK, Visa Europe and Mastercard Europe settle 24-hour cycles, with cycles falling on weekends and public holidays settling on the following business day. Every one of those is a boundary the bank's books must respect.

The fourth is that reporting requires a frozen point. A regulatory return, a statutory account, a customer statement and a management report are all statements about a defined moment. If the underlying figures never stop moving, no two reports of the same period will ever agree, and the institution loses the ability to say anything definite about itself.

In practice this is implemented as a sequence of scheduled jobs, known in core banking platforms as close of business or COB. Temenos, whose Transact platform is one of the most widely deployed core banking systems in the world, structures this as pre-COB, COB and post-COB stages. The typical sequence within it is: stop accepting new postings for the current business date; complete or fail any in-flight transactions; execute standing orders and scheduled payments due today; calculate and post interest accruals; apply fees and charges; revalue foreign currency positions; run limit and arrears checks; generate customer statements; extract summarised postings into the general ledger; produce

the trial balance; run the sub-ledger reconciliations; produce regulatory and statistical extracts; take the backup; and finally roll the business date.

The reason this sequence is treated with something close to reverence in banking operations is that it is both mandatory and time-boxed. It must complete before the next business day's processing begins, which in a 24/7 institution means it must complete inside a window measured in hours while customers are still transacting. If COB does not finish, statements do not generate, interest does not post, files do not go to the schemes, and the bank enters the next day without knowing its own position. Delayed or failed end-of-day runs are among the most serious incidents a bank can have, and they are the reason batch scheduling and workload automation remain specialist and well-paid disciplines.

The consequences of getting this wrong are not theoretical. On 20 December 2022 the Financial Conduct Authority and the Prudential Regulation Authority jointly fined TSB Bank plc a total of £48,650,000 — £29,750,000 by the FCA and £18,900,000 by the PRA — for operational risk management and governance failures relating to the bank's 2018 IT migration programme, in which customers lost access to accounts and the bank lost the ability to present them with reliable information about their own money. The fine was for the management of the risk, not for the technology. That is the correct place to put it.

■ Immutability, reversal, and why payment systems inherit the rule

The governing principle is that a posted entry is never amended and never deleted. Corrections are made by further entries.

The legal underpinning in the United Kingdom starts with the Companies Act 2006. Section 386 requires every company to keep adequate accounting records, defined as records sufficient to show and explain the company's transactions, to disclose with reasonable accuracy at any time the financial position of the company at that time, and to enable the directors to ensure the accounts comply with the Act. Section 386(3)(a) specifies that those records must contain "entries from day to day of all sums of money received and expended by the company and the matters in respect of which the receipt and expenditure takes place". Section 388 requires records to be preserved for three years from the date they are made for a private company, and six years for a public company.

For regulated firms the requirement is stricter. The FCA's Senior Management Arrangements, Systems and Controls sourcebook, at SYSC 9.1.1R, requires a firm to "arrange for orderly records to be kept of its business and internal organisation, including all services and transactions undertaken by it, which must be sufficient to enable the FCA to monitor the firm's compliance with the requirements under the regulatory system, and in particular to ascertain that the firm has complied with all obligations with respect to clients". SYSC 9.1.2R requires a common platform firm to retain records relating to its MiFID business for at least five years. A record that can be silently altered is not a record for these purposes.

The mechanics of correction have names. A **reversal**, known on the continent and in a good deal of banking software by the German term *storno*, posts the exact contra of the original entry: same accounts, same amount, opposite sides, with an explicit reference linking it to the entry being reversed. The original and the reversal both remain visible. An **adjustment** posts only the net correction. Reversal is generally preferred for erroneous postings because it preserves the gross history and makes the sequence of events legible; adjustment is used where the original entry was correct in kind but wrong in amount, and where the period is still open.

Around both sits the audit trail: who created the entry, who approved it, when, from what terminal, under what authority, and with what stated reason. Segregation of duties — maker and checker as

different people, with the system enforcing it — is the control that makes the audit trail meaningful, since an audit trail showing that one person did everything is a record of a risk rather than a mitigation of it.

Now the important part for this book. The same principle governs payment networks, and for the same reason: a message that has been sent to another institution cannot be unsent, because the other institution has already acted on it.

In the card world, ISO 8583 provides message type indicators 0400 for an acquirer reversal request and 0420 for an acquirer reversal advice. A reversal does not delete the original authorisation; it is a separate message referring to it, and it must be matched to it by the identifiers described in volume three. A refund is not a reversal at all — it is a fresh transaction in the opposite direction, submitted through the normal clearing process, which is precisely why a refund takes days when the original purchase appeared to take seconds. Nothing is being undone. A second thing is being done.

In the ISO 20022 world the vocabulary is explicit about this. Camt.056, the FIToFIPaymentCancellationRequest, is a *request* to cancel a payment — it can be refused, because the receiving institution may already have paid the money away. Camt.029, ResolutionOfInvestigation, carries the answer. And pacs.004, PaymentReturn, is the message that actually sends the funds back: a new payment, in the opposite direction, referencing the original. Three messages, and not one of them is an edit.

In Bacs the same logic appears as an operational rule rather than a message. Once a submission has passed its cut-off, the individual payments within it cannot be removed; a payment can be recalled only before a specified deadline through the originator's payment service provider, and after entry day the remedy lies outside the clearing — through the Direct Debit Guarantee for a collection, or through an indemnity claim for a credit sent in error.

This is why “can you cancel that payment?” is so often answered with an apologetic no. It is not obstruction and it is not usually a technical limitation. It is the same rule as the one that stops Ada rubbing out Monday, applied across institutional boundaries, where the stakes are higher and the parties do not report to the same management.

■ What this means if you are building something

Five consequences follow directly, and each of them is a mistake that is expensive to discover late.

Model postings as an append-only sequence and derive balances from them. If your data model has an `UPDATE balances SET amount = ?` in it anywhere, you have built a system that cannot explain itself, and an unexplainable balance is a regulatory problem before it is an engineering one.

Give every entry at least three dates and never conflate them: the timestamp at which the event occurred, the business date under which it is booked, and the value date from which it counts. Store the timezone explicitly. A great many reconciliation breaks are timezone bugs wearing a costume.

Make reversal a first-class operation with a mandatory link to the entry being reversed, and make it the only way an entry can be neutralised. If a support tool needs to fix something, it should issue a reversal through the same path as everything else, under the same approval controls, and appear in the same audit trail.

Reconcile sub-ledger to general ledger on a fixed cycle, automatically, with breaks aged and escalated. Do not wait for month end. In payments, a break that is one day old is a bug report and a break that is thirty days old is an investigation.

And define your business date explicitly, in one place, as a value the system knows about, rather than letting it be implied by whatever the server clock says. Every scheduled job, every report, every statement and every extract should agree on which day it is, and they will not do so by accident.

The next chapter takes the aggregate view that the general ledger produces and asks what it adds up to. A bank's balance sheet is not a summary of a bank; it is a description of what a bank fundamentally is, and once the ledger machinery is understood, the balance sheet stops being an accounting artefact and becomes the clearest available answer to the question of why your asset is somebody else's debt.

4.98 Common wrong ideas

1. **Wrong:** somewhere inside your bank there is a very large table with one row per customer, and that table is the ledger. **Right:** your balance lives in a sub-ledger, and the general ledger holds only an aggregate control account for a whole product, currency and division.
2. **Wrong:** the general ledger can look you up. **Right:** it can tell you what the bank owes in total by product, currency, legal entity and branch, because that is what regulatory returns, statutory accounts and capital calculations are built from.
3. **Wrong:** recording an entry and posting it are the same thing. **Right:** the lifecycle runs capture, validate, authorise, post, report, and only the fourth step changes a balance.
4. **Wrong:** the booking date and the value date are the same date. **Right:** they are two different facts about the same entry, which is why ISO 20022 carries `BookgDt` and `ValDt` as separate elements.
5. **Wrong:** tapping your card posts an accounting entry. **Right:** it creates a memo post that reduces the available balance while the ledger balance is untouched, with the genuine posting arriving one to three business days later and sometimes at a different amount.
6. **Wrong:** a bank has an end of day, singular. **Right:** CHAPS, Bacs, the card schemes and the bank's own accounting close are separate cut-offs, so the useful follow-up to "we missed the cut-off" is always *whose*.
7. **Wrong:** batch processing is a relic that event streams have made unnecessary. **Right:** interest accrual, external scheme cycles, periodic settlement and the need for a frozen reporting point each independently require a defined daily boundary.
8. **Wrong:** an immutable ledger is a property of the software you bought. **Right:** underneath sits an ordinary database in which a sufficiently privileged person can run an update, so immutability is restricted access, approval workflow, audit logging, sequence numbering and independent reconciliation.
9. **Wrong:** a mistake is fixed by finding the original entry and correcting it. **Right:** a reversal posts the exact contra with an explicit link, the correct entry is then made afresh, and both remain visible along with a note of who authorised the repair.
10. **Wrong:** a refund reverses the original payment. **Right:** a refund is a fresh transaction in the opposite direction through the normal clearing process, which is exactly why it takes days when the purchase appeared to take seconds.

4.99 Chapter summary in 20 lines

1. A principle that is never implemented is a slogan, and double entry is implemented with almost fanatical specificity.
2. Records take two shapes: a journal in the order things happened, and a ledger of accounts arranged by where things stand.
3. The chart of accounts is the controlled list of which accounts may be used, and the suspense account is where anything that fits none of them waits.

4. A day at Ada's Bank shows the pattern: a cash deposit, an internal transfer, a withdrawal and a fee, each written twice and each balancing.
5. Time passing is itself an economic event for a lender, which is why a day's interest accrual has to be computed and posted by somebody.
6. The trial balance lists every account in two columns and proves the totals agree, which is the oldest sanity check in commerce.
7. The rule about the eraser is absolute: a mistake is answered by a reversing entry, a fresh correct entry, and a note of who found and authorised it.
8. Once erasing is permitted no page proves anything about its own day, the auditor cannot check, the customer cannot dispute, and theft can be hidden by amending history.
9. There is no folder with your name on it in the general ledger, because your balance lives in a sub-ledger summarised into a control account.
10. That sub-ledger to general ledger reconciliation runs daily in a bank, and a £4.27 difference is somebody's afternoon.
11. Recording, posting, settlement and valuation are four events on four timings, and the gaps between them explain a large fraction of what customers ask about.
12. A card authorisation is a memo post against the available balance only, with the real entry landing at clearing days later.
13. A bank has a business date rather than a closing time, and a group operating in several countries may have several alive at once.
14. There is no single cut-off, because CHAPS, Bacs and the card schemes each impose their own and none of them is the bank's accounting close.
15. Batch exists because some facts are defined per day, because external cycles are batch by design, because settlement is periodic, and because reporting needs a frozen point.
16. Close of business is a mandatory, time-boxed sequence that ends by rolling the business date, and a failed run leaves a bank entering the next day without knowing its own position.
17. TSB's £48.65 million fine in December 2022 was for the management of operational risk rather than for the technology, which is the correct place to put it.
18. Immutability is a set of controls — restricted access, maker and checker, audit logging, reconciliation — rather than a property of the database underneath.
19. Reversal preserves the gross record and adjustment posts only the net difference; both are legitimate, and amendment in place is not.
20. Payment networks inherit the same rule, so camt.056 is only a request, pacs.004 is a new payment in the opposite direction, and nothing anywhere is an edit.

Chapter sources: Companies Act 2006, sections 386 and 388 (legislation.gov.uk). FCA Handbook, SYSC 9.1.1R and SYSC 9.1.2R. Bank of England — Payment and settlement pages covering CHAPS, Bacs, Faster Payments, the Image Clearing System, LINK and card scheme settlement; policy statement "Extending RTGS and CHAPS settlement hours – early morning extension" (2026). Bacs / Pay.UK — Bacs Direct Credit getting-started guidance on the three-day processing cycle and input window. ISO 20022 message definitions for camt.052, camt.053, camt.054, camt.056, camt.029 and pacs.004, together with the SWIFT camt.053.001.02 message definition report; SWIFT MT940 and MT942. ISO 8583 message type indicators 0400 and 0420. International Standard on Auditing 240 on journal entry testing. European Banking Authority reporting frameworks (COREP and FINREP). Temenos documentation on close-of-business processing. Bank of England and Financial Conduct Authority joint news release, "TSB fined £48.65m for operational resilience failings", 20 December 2022.

What a Bank Is

5.0 What this chapter gives you

1. You will be able to explain why your bank balance is a debt the bank owes you rather than a pile of notes held somewhere with your name on it.
2. You will be able to draw a bank's two lists and say which side a customer deposit sits on and which side a mortgage sits on, without hesitating.
3. You will be able to show, in two entries, how a bank makes a loan by typing a number into an account, and why that leaves its equity unchanged.
4. You will be able to say why the goldsmith-and-vault story and the money multiplier are wrong, and point to what the Bank of England published in 2014 instead.
5. You will be able to separate the constraints on the banking system as a whole from the constraints on one bank, and explain why a treasurer who says "we lend out deposits" is describing their job accurately.
6. You will be able to explain why a bank can report a capital ratio of fourteen per cent while the owners' money is about five per cent of its balance sheet, and why both figures are true.
7. You will be able to say why capital did not save Silicon Valley Bank, and keep the defence against losses distinct from the defence against withdrawals.
8. You will be able to describe what happens to depositors, shareholders and AT1 holders when a British bank is resolved over a weekend, using SVB UK as the worked case.
9. You will be able to state what the FSCS protects, why "per banking licence" is not the same as "per brand", and where an uncovered corporate depositor ranks in an insolvency.
10. You will be able to explain why a payment institution is not a bank, has neither the capital nor the liquidity regime, and must therefore safeguard its customers' funds rather than lend them.

Ask a hundred people in a British high street what a bank is and roughly ninety-five will describe a building where money is kept. The remaining five will say something about apps. Both answers describe the furniture.

A bank is not a place where money is kept. Very nearly the opposite is true: a bank is an institution that has spent the money you gave it, is legally obliged to give it back the moment you ask, and stays alive by being extremely careful about the gap between those two facts. Everything difficult about banking lives in that gap. Every rule, every regulator, every capital ratio, every crisis of the last four hundred years.

This is not a criticism. If banks did nothing but hold money they would be warehouses, they would charge you rent, and no house in Britain would ever be bought with a mortgage. The whole point of the institution is that it takes money people want back at a moment's notice and turns it into money other

people can borrow for twenty-five years. That transformation is useful, dangerous, and impossible to perform without the possibility of failure. Banking regulation is an attempt to make failure rare, survivable and orderly, not to abolish it, because abolishing it would mean abolishing the bank.

Two earlier chapters carry the weight here. Chapter two established that your balance is not a substance but a record. Chapter three established that no record of money may be changed by writing down one fact; every change is two entries that balance. This chapter puts them together: if your balance is a record in somebody else's book, whose book is it, what else is in it, and what happens when it stops adding up?

Along the way we will demolish the most widely taught and most thoroughly wrong story about banking: the goldsmith, the vault, and the fraction kept in reserve. It appears in school textbooks, in economics degrees, in newspaper explainers and in a great many bank employees' understanding of their own employer. The Bank of England published a paper in 2014 saying, in effect, that it is not how anything works. We will get to that paper. First, the plain version.

The plain version

■ Your money is somebody else's promise

Start with the thing in your hand. Open your banking app. It says, let us say, £1,842.60.

That number is not a picture of money sitting somewhere with your name on it. There is no drawer. There is no bundle of notes in a vault in Edinburgh labelled *Ms A. Okonkwo*. The number is a **promise**: the bank saying we owe you one thousand eight hundred and forty-two pounds and sixty pence, and we will hand it over, or send it wherever you tell us, whenever you ask.

Your bank balance is a debt. It is a debt the bank owes to you.

This feels alarming the first time you hear it, then stops feeling alarming when you notice it is exactly what you already believed. You have never once thought a specific pile of notes was yours. You have never asked which notes. What you care about is that when you tap your card for £3.20, the coffee shop gets £3.20 and your app shows £3.20 less. A promise that is always kept is indistinguishable, in daily life, from a thing you own.

■ The coat-check that lends the coats

Here is the analogy, and it is a good one for about four paragraphs.

Imagine a cloakroom at a theatre. You hand over your coat, you get a ticket, and at the end of the evening you swap the ticket back for the coat. It would be outrageous if, halfway through the second act, the attendant lent your coat to somebody walking past outside. Your coat is a particular coat. You want *that* one back.

Money is not like a coat. Any ten-pound note is as good as any other. So a cloakroom that took in money instead of coats could lend it out during the evening and nobody would mind, provided it could hand *a* ten-pound note back to each ticket-holder at the end. Nobody checks the serial numbers. That is roughly the story people tell about the goldsmiths of seventeenth-century London: they took in gold for safekeeping, issued receipts, noticed that most of the gold sat there untouched, and started lending it.

Hold that picture for one more minute, because we are going to break it. It gets you to the first important idea: a bank has two lists, not one.

■ The two lists

Every bank keeps two lists, and understanding a bank means understanding both at once.

The first is **what the bank owes**. Your £1,842.60 is on it. So is every other customer's balance, plus money the bank has borrowed from other banks, plus bonds it has issued to investors.

The second is **what the bank owns**. Mortgages it has lent out and expects to be repaid. Business loans. Credit card balances people owe it. Government bonds it has bought. Actual cash. A pile of computers and some rather expensive branch leases.

In a healthy bank, the second list is bigger than the first. It has to be. The difference between them is the owners' money — the shareholders' stake, built from what they originally put in and from profits kept rather than paid out. It is the cushion. If the bank owns £100 of things and owes £95, the cushion is £5.

That cushion is not a fund sitting in an account somewhere. It is not money set aside. It is simply the gap: what would be left over for the owners if every asset were turned into cash and every debt paid. Its entire job is to absorb bad news. When a borrower does not pay, the first list stays exactly where it is — the bank still owes you every penny of your £1,842.60 — and the second list shrinks. **The loss comes out of the owners' cushion, not out of the depositors' promises. The cushion is there precisely so that depositors never notice.** A bank fails when the losses are big enough to eat the whole cushion and start eating the promises.

■ Building a small bank with real numbers

Meet Priya, who is starting a very small, very simplified bank in a town called Little Sutton. The numbers are tiny and invented so they can be followed; everything structural about them is true of a real bank.

Priya puts in £10,000 of her own money. On day one the bank owns £10,000 of cash, owes nothing, and the owners' cushion is £10,000. Then two hundred residents open accounts and pay in £90,000 between them.

What the bank owns		What the bank owes	
Cash	£100,000	Owed to depositors	£90,000
		Owners' cushion	£10,000
Total	£100,000	Total	£100,000

Notice that the totals still match, and why. Priya received £90,000 of cash and simultaneously took on £90,000 of promises. Two entries, as chapter three insisted. The cushion did not change, because taking a deposit does not make a bank richer or poorer. It makes it bigger.

Now Priya lends. Mr Hollis wants £60,000 to expand his bakery. Here the coat-check analogy has to be retired. Priya does not hand him sixty thousand pounds in notes from the drawer. She opens him an account and types **£60,000** into it. That is the loan.

Watch the lists. The bank now owns something new: Mr Hollis's promise to repay £60,000, an asset, because it is money owed *to* the bank. And it owes something new: £60,000 sitting in Mr Hollis's account, which he can spend, because it is money owed *by* the bank.

What the bank owns		What the bank owes	
Cash	£100,000	Owed to depositors	£150,000
Loan to Mr Hollis	£60,000	Owners' cushion	£10,000
Total	£160,000	Total	£160,000

Read that table twice. The bank is now £60,000 bigger on both sides. Nothing came in the door. Nobody deposited anything. Priya created £60,000 of spendable money by typing it, and created £60,000 of debt owed to her at the same instant, and the two exactly cancel, so her cushion is still £10,000 and her books still balance.

This is not a trick and it is not fraud. It is what banking is. When the Bank of England explains this to the public it says, in plain words, that most money in the economy is created not by printing presses but by banks when they make loans, and that when a loan is repaid the money is deleted and ceases to exist. Around eight in every ten pounds in the British economy is money of this kind. Notes and coins are around three per cent of the total.

So the goldsmith story is backwards. Banks do not take in money and lend it out. They lend, and the lending is what puts the money there.

■ So what stops Priya typing numbers forever?

Four things, and each is a whole industry.

First, the money walks out. Mr Hollis did not borrow £60,000 to look at it. He spends it — on ovens from a supplier who banks with Barclays, on a builder who banks with Santander. Real money must then move from Priya's bank to those banks. She cannot type that; she has to have it, or borrow it. So while the banking system as a whole creates deposits by lending, any individual bank constantly leaks. Managing that leak is called funding, and it is why banks compete so hard for your savings account even though they do not need your savings in order to lend.

Second, some borrowers do not pay. If Mr Hollis's bakery fails and Priya recovers only £35,000, she has lost £25,000. Her cushion was £10,000. It is gone, and then some, and her bank is bust — not because it ran out of cash, but because what it owns is worth less than what it owes. This is why regulators insist the cushion be large relative to how risky the lending is.

Third, everyone might want their money at once. Priya has £100,000 of cash and £150,000 of promises. If more than £100,000 of promises are called in on the same afternoon she cannot pay, even though she is perfectly solvent and Mr Hollis is making excellent sourdough. Being unable to pay today is fatal regardless of being rich tomorrow. This is why regulators separately insist that banks hold a stock of things that can be turned into cash immediately.

Fourth, somebody is watching. In the United Kingdom that is the Prudential Regulation Authority, part of the Bank of England. It sets how big the cushion must be, how much instantly-sellable stuff must be held, and what happens if either falls short. It can and does stop banks paying dividends, and in the last resort it can close them.

■ What a run looks like

The third danger is the one that actually kills banks, and nobody believes it will happen until it is happening.

On Friday 14 September 2007, queues formed outside branches of Northern Rock across Britain. The bank had asked the Bank of England for emergency liquidity support two days earlier and the news

had leaked. Those queues are commonly described as the first run on a British bank since Overend, Gurney and Company collapsed in 1866. Northern Rock was not obviously insolvent that morning. It was illiquid: it had funded long-dated mortgages with short-dated wholesale borrowing that had suddenly stopped being available. The Chancellor guaranteed the deposits the same day to stop the queues. The bank was taken into state ownership on 22 February 2008.

Sixteen years later the same thing happened at hundreds of times the speed, because the queue was now an app. Silicon Valley Bank's customers attempted to withdraw more than forty billion dollars on 9 March 2023 and its management expected over a hundred billion more the following day. It was closed on 10 March. Its British subsidiary lost about a third of its deposits in two days. Nobody needed to stand in a queue.

■ The two words

Accountants and bankers do not say “what it owns” and “what it owes”. They say **assets** and **liabilities**. The cushion is called **equity**, or **capital**.

And now the sentence beginners find upside down, which chapter three prepared you for: **your deposit is the bank's liability, and the bank's loans are its assets**. Your asset — the money you think of as yours — is the bank's debt. The mortgage you think of as your debt is the bank's asset. Every payment in the rest of this book is a movement between somebody's assets and somebody's liabilities, and getting the direction right is the difference between a payments system that reconciles and one that does not.

Where the plain version stops being true

The plain version above is honest as far as it goes. Four things in it will actively mislead you if you carry them forward unamended.

■ Correction one: there is no fraction, and there is no reserve

The version taught in most schools says banks must keep a fixed fraction of deposits — the classic figure is ten per cent — and may lend the rest, and that this fraction, applied repeatedly as money is redeposited, “multiplies” an injection of central bank money into a larger stock of deposits. The ratio is the reserve requirement; the mechanism is the money multiplier.

Two separate things are wrong with this.

The first is empirical. **The United Kingdom has had no mandatory reserve ratio since 1981**. The number is not ten per cent; there is no number. The Bank of England does require banks above a size threshold to place non-interest-bearing cash ratio deposits with it, but that is a levy funding the Bank's policy functions, not a liquidity rule. The United States abolished its requirements outright: the Federal Reserve Board reduced all reserve requirement ratios to zero per cent effective 26 March 2020. The euro area has kept a requirement, but it is one per cent of a defined deposit base, unchanged since 18 January 2012, and since the maintenance period beginning 20 September 2023 the European Central Bank has remunerated those reserves at zero per cent. A one per cent ratio is not a binding constraint on lending; it is a small tax.

The second is causal, and matters more. The multiplier story has the arrow pointing the wrong way. It says central bank money comes first and deposits are built on top. In practice the central bank supplies reserves on demand to keep the policy rate where it has decided it should be. If banks collectively need more reserves, they get more reserves, at the policy rate. The quantity of reserves is an outcome, not a constraint.

■ **Correction two: “banks lend out your savings” is wrong, but so is “banks can lend without limit”**

Having demolished the vault, it is easy to overcorrect into an equally false position: that because lending creates deposits, a bank is unconstrained and money is conjured at will. The distinction that resolves it is between the system and the firm.

For the **banking system as a whole**, lending creates deposits and no prior pot of savings is required. That is the Bank of England’s point, and it is correct.

For an **individual bank**, the picture is much tighter, because the deposits it creates walk straight out to other banks. Priya’s £60,000 becomes Barclays’ problem within a week, and when it leaves she must settle in central bank reserves, which she cannot create. If she persistently lends more than she attracts, she must buy funding in the wholesale market, and that funding has a price and, crucially, an availability that can vanish overnight. Northern Rock died of exactly this. So a bank treasurer who tells you “we lend out deposits” is describing their daily job accurately even though the sentence is wrong as monetary economics. Practitioners and economists talk past each other about this constantly.

A further constraint the plain version understates: **capital is scarce and expensive**. Every new loan consumes capital, because the regulator requires the cushion to scale with the risk taken. A bank wanting to lend another billion pounds must find the capital to support it, and capital comes from retained profits or shareholders who want a return. This, not any reserve ratio, is the operative brake on lending in a modern bank.

■ **Correction three: capital is not cash, and capital ratios do not measure survival**

The plain version described the cushion as “what would be left over”. That is right, but three things about it routinely mislead.

Capital is **not a pot of money**. You cannot spend it, hold it or point at it. It is an accounting difference between two other numbers. When a regulator says a bank must hold more capital, it is not saying “set aside more cash”; it is saying “fund yourself with more shareholder money and less borrowed money”, which changes nothing on the asset side.

Capital ratios are measured against **risk-weighted assets**, not actual assets. A mortgage at fifty per cent loan-to-value and an unsecured loan to a start-up both sit on the balance sheet at face value, but the second absorbs far more capital. This is why a bank can report a capital ratio of fourteen per cent while its equity is about five per cent of its actual balance sheet. Both numbers are true and answer different questions. Confusing them is the commonest error in commentary about banks.

And most importantly: **capital does not stop a run**. Capital protects against losses; liquidity protects against withdrawals. They are separate defences against separate deaths, and the second kills faster. Silicon Valley Bank’s capital ratios were above the required minimums when it failed. It failed because its depositors — approximately ninety-four per cent of whom were uninsured, and who knew each other — asked for their money at once, and selling assets fast enough to meet them would have crystallised losses that then *did* exceed the capital. A bank can be killed by either.

■ **Correction four: the balance sheet is not the whole bank**

The two-list picture is complete only in the sense that a photograph is complete. It omits three things that matter.

It omits **time**. Both lists have maturities. A bank’s assets are typically long — a twenty-five-year mortgage — and its liabilities short, or instant. That mismatch is maturity transformation, and it is

not a flaw in the business model, it is the business model. It is also why banks are fragile in a way a manufacturer is not.

It omits **what is not on it**. Undrawn credit facilities, loan commitments, guarantees and letters of credit produce no balance-sheet entry until somebody draws on them — which they do precisely when things are going badly. Regulators handle this with credit conversion factors, a technical way of saying “assume some of it becomes real”.

And it omits **valuation**. “What the bank owns” is an estimate, governed by accounting standards, involving forecasts of who will default. Under IFRS 9 a bank must book expected credit losses before anything has gone wrong. Two honest banks with identical loan books can report different asset values. The balance sheet is a considered opinion presented in the grammar of fact.

The technical version

■ The identity, and a real balance sheet

Everything above is a restatement of the accounting identity introduced in chapter three:

Assets = Liabilities + Equity

For a bank, the composition of those three terms is unusual in a way no other industry shares: its principal liability is money, and its principal asset is credit. Here is a real one — Lloyds Banking Group plc, statutory consolidated position at 31 December 2025, as reported in the group’s 2025 results announcement of 28 January 2026.

Item	31 Dec 2025	31 Dec 2024
Total assets	£944.1bn	£906.7bn
Total liabilities	£896.2bn	£860.8bn
Ordinary shareholders’ equity	£41.8bn	£39.5bn
Other equity instruments	£5.9bn	£6.2bn
Non-controlling interests	£0.2bn	£0.2bn
Total equity	£47.9bn	£45.9bn

Total equity of £47.9 billion against total assets of £944.1 billion is approximately 5.1 per cent — the honest, unweighted answer to “how much of this bank is the owners’ money”.

On the **liability** side, customer deposits were £496.5 billion — retail current accounts £102.8 billion, retail savings £212.5 billion, Wealth £9.9 billion, Commercial Banking £171.1 billion. Wholesale funding (money markets, debt securities in issue, covered bonds, securitisations, subordinated debt) was £99.4 billion, of which £37.0 billion had a residual maturity under one year. Also here sit derivative liabilities, liabilities arising from insurance and investment contracts — Lloyds owns Scottish Widows — and repurchase agreements.

On the **asset** side, underlying loans and advances to customers were £481.1 billion: UK mortgages £323.1 billion, credit cards £17.3 billion, UK retail unsecured loans £10.5 billion, UK Motor Finance £16.4 billion, Business and Commercial Banking £28.3 billion, Corporate and Institutional Banking £62.0 billion. Financial assets at amortised cost totalled £554 billion; at fair value through profit or loss £240 billion; at fair value through other comprehensive income £36 billion; derivative financial assets £20 billion.

The loan-to-deposit ratio was 97 per cent — the clearest single statement of what this institution is. It has deliberately turned very nearly every pound the public entrusted to it into somebody else's long-dated debt. The expected credit loss allowance of £3,353 million, 0.7 per cent of loans and advances, is the accounting recognition of correction four: the asset side is net of a forecast.

■ Money creation: the 2014 Quarterly Bulletin

The definitive modern statement is *Money creation in the modern economy*, by Michael McLeay, Amar Radia and Ryland Thomas of the Bank of England's Monetary Analysis Directorate, in the Bank of England Quarterly Bulletin 2014 Q1, published 14 March 2014. Anyone working in payments should read it once. Its central claims, in its own words:

"In the modern economy, most money takes the form of bank deposits. But how those bank deposits are created is often misunderstood: the principal way is through commercial banks making loans."

"Whenever a bank makes a loan, it simultaneously creates a matching deposit in the borrower's bank account, thereby creating new money."

"Rather than banks lending out deposits that are placed with them, the act of lending creates deposits — the reverse of the sequence typically described in textbooks."

And, explicitly refuting both traditional models: banks "do not act simply as intermediaries, lending out deposits that savers place with them, and nor do they 'multiply up' central bank money to create new loans and deposits."

Mechanically, at the moment of origination, the lending bank posts:

	Debit	Credit
Loan asset (customer receivable)	Amount	
Customer deposit (liability)		Amount

Two entries. Balance sheet grows on both sides. Equity unchanged. Nothing has been taken from anyone. On repayment the entries reverse and the deposit ceases to exist — money destruction, which is why aggregate deposits fall when credit contracts even though nobody has withdrawn anything.

The Bank's public explainer gives the composition of UK money as approximately 3 per cent notes and coin, 18 per cent central bank reserves and 79 per cent commercial bank deposits. Reserves and deposits are different instruments circulating in different closed circuits, which is the subject of the next chapter.

The article is equally clear about what does constrain creation. Banks must lend profitably in a competitive market. They must manage the credit and liquidity risk new lending creates. Prudential regulation constrains their activities in the interest of system resilience. And "monetary policy acts as the ultimate limit on money creation" — in normal times through the price of reserves, in unconventional times through asset purchases.

■ The Basel framework in outline

The rules implementing those constraints are, in outline, international. The Basel Committee on Banking Supervision, hosted at the Bank for International Settlements, comprises 45 members from 28 jurisdictions, consisting of central banks and authorities with formal responsibility for banking supervision. Its standards have **no legal force whatsoever**. They are agreements between supervisors, which member jurisdictions then implement — or do not, or implement late, or differently — through

their own law. In the United Kingdom the binding instrument is the PRA Rulebook, made under the Financial Services and Markets Act 2000 as amended. Anyone who says a British bank is “required by Basel” to do something is speaking loosely; it is required by the PRA, which has chosen to implement Basel.

Capital: the stack

Regulatory capital is a hierarchy ordered by how readily it absorbs losses.

Common Equity Tier 1 (CET1) is ordinary shares, share premium and retained earnings, less regulatory deductions — goodwill and other intangibles, deferred tax assets relying on future profitability, prudent valuation adjustments, the excess of expected losses over accounting provisions, defined benefit pension surpluses and significant investments in financial sector entities. This is the loss-absorbing core.

Additional Tier 1 (AT1) is perpetual, deeply subordinated instruments with no maturity date and fully discretionary coupons, which convert to equity or are written down when CET1 falls through a contractual trigger — the instruments the press calls “cocos”. **Tier 2** is dated subordinated debt, absorbing losses only in resolution or liquidation.

Lloyds’ regulatory capital position at 31 December 2025 makes the structure concrete:

	31 Dec 2025
CET1 capital	£32,930m
Total Tier 1 capital	£38,053m
Total capital resources	£44,579m
Total MREL resources	£75,732m
Risk-weighted assets	£235,513m
CET1 ratio	14.0%
Tier 1 ratio	16.2%
Total capital ratio	18.9%
MREL ratio	32.2%
UK leverage ratio	5.4%

Risk weighting, and why £944 billion becomes £235 billion

Note the gap between total assets of £944.1 billion and risk-weighted assets of £235.5 billion. Two effects produce it.

First, **risk weighting proper**. Each exposure is multiplied by a weight reflecting its assessed riskiness, either from the standardised approach’s look-up tables or from the bank’s own internal ratings-based models where the PRA has granted permission. A low loan-to-value residential mortgage attracts a small weight; an unsecured corporate exposure a large one. Cash and central bank reserves attract zero.

Second, **scope of consolidation**. The regulatory consolidation is not the accounting consolidation. Lloyds’ insurance business is not consolidated for capital purposes; instead the group deducts £4,708 million of significant investments from CET1. The same effect explains why the leverage exposure measure is smaller than accounting total assets.

Because two banks with identical books could report very different RWAs, the Basel III finalisation package of December 2017 imposed an **output floor**: RWAs must be at least 72.5 per cent of what the

standardised approaches would produce, phasing in from 50 per cent over five years. It also barred the advanced IRB approach for large corporates, banks and equities, imposed input floors (probability of default at five basis points for corporates, loss-given-default at 25 per cent for unsecured corporate exposures), and removed internal models from operational risk and CVA entirely. In the United Kingdom this package is called Basel 3.1; the PRA announced on 17 January 2025 a one-year delay to 1 January 2027, with full implementation still set for 1 January 2030.

The required ratios

Basel III's Pillar 1 minimums are 4.5 per cent CET1, 6 per cent Tier 1 and 8 per cent total capital, all as percentages of risk-weighted assets. On top sit buffers, which are not minimums in the same sense: breaching a buffer does not make a bank unlawful, it triggers automatic restrictions on dividends, buybacks and discretionary bonuses through the maximum distributable amount mechanism.

Layer	Calibration
CET1 minimum	4.5% of RWAs
Tier 1 minimum	6.0% of RWAs
Total capital minimum	8.0% of RWAs
Capital conservation buffer	2.5% of RWAs, CET1 only
Countercyclical capital buffer	0–2.5% (UK rate currently 2%)
Systemic buffers	Firm-specific, for G-SIIs and O-SIIs
Pillar 2A	Firm-specific, set by the PRA for risks not captured in Pillar 1
PRA buffer (Pillar 2B)	Firm-specific, confidential, stress-test driven

The countercyclical buffer is the interesting instrument: the Financial Policy Committee raises it when it judges risk to be building and cuts it in stress, so that capital accumulated in good years can be released rather than forcing banks to shrink lending exactly when the economy needs it.

The leverage ratio

Because risk weighting can be gamed and models can be wrong, Basel III added a deliberately crude backstop: Tier 1 capital divided by a total exposure measure that includes on-balance-sheet assets, derivative exposures, securities financing transactions and off-balance-sheet items converted at prescribed factors, with no risk weighting at all.

The UK framework, set by Financial Policy Committee Direction and implemented in the Leverage Ratio (Capital Requirements and Buffers) Part of the PRA Rulebook, imposes a **minimum of 3.25 per cent** of Tier 1 capital, plus a countercyclical leverage ratio buffer and, for systemic firms, an additional leverage ratio buffer. Following PRA Policy Statement 22/25 of 12 November 2025, the retail deposits threshold for application of the requirement rose from £50 billion to £75 billion with effect from 1 January 2026, calculated on a three-year moving average rather than a point-in-time value; the separate threshold of £10 billion of non-UK assets was left unchanged. Firms below the thresholds face a supervisory expectation, in Supervisory Statement 45/15, that their leverage ratio should not ordinarily fall below 3.25 per cent. Lloyds reported a UK leverage ratio of 5.4 per cent at 31 December 2025.

■ Liquidity

Capital answers “can this bank absorb losses”. Liquidity answers “can this bank pay today”. Basel III introduced two standards.

The **Liquidity Coverage Ratio** requires the stock of high-quality liquid assets to be at least 100 per cent of total net cash outflows over a 30-calendar-day stress scenario. HQLA are assets convertible into cash quickly, by sale or by being pledged as collateral, without significant loss of value. Level 1 assets — cash, central bank reserves, certain sovereign debt — count without limit or haircut. Level 2A assets take a 15 per cent haircut. Level 2 assets in aggregate may not exceed 40 per cent of the stock, and Level 2B may not exceed 15 per cent. The standard took effect on 1 January 2015 at 60 per cent, rising ten percentage points a year to 100 per cent on 1 January 2019.

The **Net Stable Funding Ratio** works over a one-year horizon: available stable funding divided by required stable funding, minimum 100 per cent. ASF weights liabilities and capital by how likely they are to stay put — long-dated funding and stable retail deposits score highly, overnight wholesale funding zero. RSF weights assets by illiquidity and by how long they must be funded for. It became a minimum standard for internationally active banks on 1 January 2018.

Lloyds reported an LCR of 145 per cent at 31 December 2025 on LCR-eligible assets of £131.4 billion, and an NSFR of 124 per cent. From those figures the modelled 30-day net outflow can be inferred at roughly £91 billion — about 18 per cent of its customer deposit book. Hold that against Silicon Valley Bank UK, which lost about 32 per cent of its deposits in two days. That is the honest limit of the LCR: it is calibrated to a severe scenario, not the worst imaginable one, and a modern digital run can outrun it.

■ Why a bank failing is not like a company failing

A car parts manufacturer that becomes insolvent enters administration. An officeholder is appointed, trading may continue, assets are realised, creditors queue in a statutory order, and the process takes months or years. Nobody outside the supply chain much notices on the day.

A bank cannot be treated this way, for one reason: **its liabilities are the money supply and the payments system**. Freezing a bank’s balance sheet does not merely disappoint creditors, it deletes the current accounts from which salaries are paid, direct debits collected and card transactions settled. It destroys money and severs access to the payment rails simultaneously, for people who never chose to be creditors of anything.

Britain’s answer is the **Special Resolution Regime** in the Banking Act 2009, administered by the Bank of England as resolution authority, which sits alongside and largely displaces ordinary insolvency law for banks, building societies and certain investment firms. It provides five stabilisation options:

1. **Bail-in** — write-down of the claims of the bank’s unsecured creditors, including holders of capital instruments, with conversion to equity, recapitalising the firm in place.
2. **Transfer to a private sector purchaser**, executed without the consent of shareholders or counterparties.
3. **Transfer to a bridge bank** controlled by the Bank of England pending an onward sale.
4. **Transfer to an asset management vehicle** to run down a problem portfolio.
5. **Temporary public ownership**, exercised by HM Treasury as a last resort.

Where resolution is not justified — typically a small firm whose failure poses no systemic threat — the **bank insolvency procedure** applies instead: modified insolvency, with a liquidator holding two objectives in strict priority. First, work with the Financial Services Compensation Scheme so that pro-

tected depositors are paid out within seven days or their accounts transferred elsewhere. Only then, wind the firm up in the general creditors' interest.

Depositor preference

Insolvency law itself has been rewritten to put depositors near the front of the queue. Under Schedule 6 of the Insolvency Act 1986 as amended by the Banks and Building Societies (Depositor Preference and Priorities) Order 2014, paragraph 15B makes the portion of an eligible deposit not exceeding FSCS compensation an **ordinary preferential debt** — in practice the FSCS pays the depositor and stands in their place, which is why it is sometimes described as super-preferred. Paragraph 15BA makes the portion of an eligible deposit *exceeding* the compensation limit, owed to eligible persons — broadly individuals and micro, small and medium-sized enterprises — a **secondary preferential debt**. Paragraph 15BB extends secondary preference to deposits made through non-UK branches which would have been eligible had they been made in the UK.

Secondary preferential debts rank after ordinary preferential debts, but both rank ahead of floating charge holders and ordinary unsecured creditors. A large corporate depositor, by contrast, is an ordinary unsecured creditor and ranks with the bondholders.

FSCS cover

The Financial Services Compensation Scheme protects eligible deposits up to **£120,000 per eligible person per banking licence**, a limit that took effect on 1 December 2025, having previously been £85,000. Joint accounts are covered to £240,000. Qualifying temporary high balances — house sale proceeds, redundancy payments, insurance settlements — are protected up to £1.4 million for six months from the date first deposited.

The phrase “per banking licence” is doing real work and is routinely misread. Several familiar high-street brands may share a single authorisation, in which case the limit applies once across all of them, not once per brand.

MREL and loss-absorbing capacity

For resolution to work without public money, there must be enough liabilities available to be written down. That is the **minimum requirement for own funds and eligible liabilities**: own funds plus eligible liabilities, positioned appropriately in the creditor hierarchy, sufficient to absorb losses and recapitalise the firm. Lloyds reported total MREL resources of £75,732 million at 31 December 2025, an MREL ratio of 32.2 per cent of risk-weighted assets — well over twice its total capital ratio, because MREL includes £31,232 million of other eligible liabilities issued by the holding company.

Resolution carries a statutory safeguard: **no creditor worse off**. Following a bail-in or partial transfer, no pre-resolution shareholder or creditor may be left worse off than in an ordinary insolvency. An independent valuation determines this after the event, and compensation is payable if the test is failed.

The regime in action: SVB UK

The clearest worked example in British experience is the resolution of Silicon Valley Bank UK Limited, documented by the Bank of England in its report under section 79A of the Banking Act 2009.

SVB UK operated as a branch from 2012 and converted to a UK subsidiary in July 2022. Its parent was closed by the California Department of Financial Protection and Innovation on 10 March 2023. At 11pm that night the Bank of England announced its intention to place SVB UK into the bank insolvency

procedure. Over the weekend it changed course, determining that a private sector purchaser transfer better served the special resolution objectives. At 7am on Monday 13 March 2023 the Bank exercised its transfer powers and moved the whole of SVB UK's share capital to HSBC UK Bank plc.

SVB UK	Figure
Total assets at 10 March 2023	£8,756m
Total liabilities at 10 March 2023	£7,343m
Net assets	£1,413m
Customer deposits, 8 March 2023	£9,779m
Customer deposits, 10 March 2023	£6,688m
Demand deposits, 8 March 2023	£5,215m
Demand deposits, 10 March 2023	£2,140m
AT1 written down under section 6	£322m
Tier 2 written down under section 6	£33m
Consideration paid by HSBC	£1

The AT1 and Tier 2 instruments were mandatorily written down as required by section 6 of the Act, wiping out £355 million of investors' claims. The equity was transferred for one pound. Every depositor — covered and uncovered alike — retained uninterrupted access to their accounts. No UK public funds were used. Both independent valuations concluded that FSCS-covered depositors would have recovered 100 per cent of their claim in a counterfactual insolvency, so the no-creditor-worse-off test was comfortably satisfied for them.

Read that table as a summary of the chapter. A bank with £8.8 billion of assets and £1.4 billion of net assets — an owners' cushion of over 16 per cent of its balance sheet, far above any regulatory minimum — was resolved over a weekend because roughly a third of its deposits left in two days. Solvency did not save it. What saved its depositors was a statutory regime built on the assumption that this would happen to somebody, eventually.

Ring-fencing

One further piece of British structural regulation belongs here. The Financial Services (Banking Reform) Act 2013, implementing the recommendations of the Independent Commission on Banking, requires the largest UK banking groups to separate core retail banking — deposit-taking from individuals and small businesses, and the payment services attached to it — from investment banking activity, into a legally distinct, separately capitalised, separately governed **ring-fenced body**. The regime came into force on 1 January 2019, by which point approximately £1.2 trillion of core deposits had been placed inside ring-fenced entities. The threshold, set by the Financial Services and Markets Act 2000 (Ring-fenced Bodies and Core Activities) Order 2014, is now **£35 billion of core deposits**, above which a group conducting material investment banking activity must ring-fence, subject to a trading assets condition allowing exit where trading assets are less than 10 per cent of Tier 1 capital. Reforms described by the government as “smarter ring-fencing” came into force on 4 February 2025.

The practical consequence for anyone building payment systems is that “Barclays” and “HSBC” are not single legal entities. A group may contain a ring-fenced bank, a non-ring-fenced bank, an overseas subsidiary and a service company, each with its own authorisation, balance sheet and, in some cases, scheme participation. Which entity a payment is with has a real answer and real consequences.

■ What to carry forward

Your balance is the bank's liability and the bank's loans are its assets. Every payment message in Volumes III and IV is an instruction to move a liability from one institution's book to another's, and the direction of every posting depends on remembering whose book you are in.

Lending creates deposits; deposits are not lent out. This is why the money supply expands and contracts with credit, and why "where did the money come from" is usually the wrong question to ask about a payment.

Capital protects against loss and liquidity against withdrawal, and they are entirely different defences. A payments business — an authorised payment institution, an e-money firm, a fintech — is not a bank and has neither regime, which is precisely why its customers' funds must be safeguarded in a segregated account rather than lent.

And commercial bank money is not the only kind. When Priya's £60,000 walks out of Little Sutton to Barclays, something else has to move — something no commercial bank can create. That something is central bank money, and it is the subject of the next chapter.

5.98 Common wrong ideas

Wrong: A bank keeps your money in a vault and lends out a fraction of it. **Right:** A bank has already spent what you gave it; your balance is a promise, and lending creates new deposits rather than passing on old ones.

Wrong: British banks must hold ten per cent of deposits in reserve. **Right:** The United Kingdom has had no mandatory reserve ratio since 1981, and the Federal Reserve reduced all its ratios to zero per cent on 26 March 2020.

Wrong: Because lending creates deposits, a bank can lend without limit. **Right:** The deposits an individual bank creates walk out to other banks and must be settled in reserves it cannot create, and capital, funding and the regulator bind long before anything else does.

Wrong: Capital is a pot of money the bank sets aside for a rainy day. **Right:** Capital is the accounting difference between two other numbers, and requiring more of it means funding yourself with more shareholder money and less borrowed money.

Wrong: A well-capitalised bank cannot fail. **Right:** Capital absorbs losses and liquidity meets withdrawals; Silicon Valley Bank's capital ratios were above the required minimums on the day it was closed.

Wrong: A capital ratio of fourteen per cent means fourteen per cent of the balance sheet is the owners' money. **Right:** Ratios are struck against risk-weighted assets, which is why Lloyds could report 14.0 per cent while £47.9 billion of equity against £944.1 billion of assets is about 5.1 per cent.

Wrong: Basel rules bind British banks directly. **Right:** Basel standards have no legal force whatsoever; what binds a UK firm is the PRA Rulebook, made under the Financial Services and Markets Act 2000.

Wrong: FSCS cover of £120,000 applies to each of your bank's brands. **Right:** It applies per eligible person per banking licence, and several familiar high-street brands may share a single authorisation.

Wrong: A failing bank goes into administration like any other company. **Right:** Its liabilities are the money supply and the payments system, so the Special Resolution Regime in the Banking Act 2009 largely displaces ordinary insolvency law.

Wrong: The balance sheet tells you what a bank owns. **Right:** It is a considered opinion presented in the grammar of fact, net of IFRS 9 expected credit losses, silent on undrawn commitments and silent on the maturity mismatch that makes the business fragile.

5.99 Chapter summary in 20 lines

1. A bank is not a place where money is kept but an institution that has spent what you gave it and is obliged to return it the moment you ask.
2. Your balance is a record of a debt the bank owes you, not a claim on any particular notes.
3. Every bank keeps two lists: what it owns, called assets, and what it owes, called liabilities.
4. The difference between them is equity, the owners' cushion, and its whole job is to absorb losses so that depositors never notice.
5. Taking a deposit makes a bank bigger rather than richer, because the cash coming in and the promise going out are two entries of the same size.
6. When a bank lends it types a number into the borrower's account, creating an asset and a liability at the same instant and leaving its equity untouched.
7. Lending therefore creates deposits and repayment destroys them, which is the reverse of the sequence taught in most textbooks.
8. The Bank of England said exactly that in *Money creation in the modern economy* in the 2014 Q1 Quarterly Bulletin.
9. Around seventy-nine per cent of UK money is commercial bank deposits, eighteen per cent central bank reserves and about three per cent notes and coin.
10. There is no reserve ratio in the United Kingdom and no money multiplier, because reserves are supplied on demand to keep the policy rate where the central bank wants it.
11. An individual bank is nonetheless tightly constrained, since the deposits it creates leave for other banks and must be settled in reserves it cannot create.
12. Capital is scarce and expensive and must scale with the risk taken, and that, not any reserve ratio, is the operative brake on lending.
13. Capital is not cash but an accounting difference, and capital ratios are measured against risk-weighted assets rather than actual ones.
14. That is why Lloyds reported a CET1 ratio of 14.0 per cent against £235.5 billion of risk-weighted assets while its equity was about 5.1 per cent of £944.1 billion of actual assets.
15. Capital protects against loss and liquidity protects against withdrawal, and they are separate defences against separate deaths.
16. Northern Rock in 2007 and Silicon Valley Bank in 2023 both died of the second, the latter's British subsidiary losing about a third of its deposits in two days.
17. A bank cannot be wound up like a car parts manufacturer, because its liabilities are the money supply and the payment rails at the same time.
18. Britain's answer is the Special Resolution Regime of the Banking Act 2009, offering bail-in, private sector transfer, a bridge bank, an asset management vehicle and temporary public ownership.
19. SVB UK showed the regime working: £355 million of AT1 and Tier 2 written down, the share capital transferred to HSBC for one pound, and every depositor's access unbroken with no public funds used.
20. Carry forward that your deposit is the bank's liability and its loans are its assets, and that when money leaves one bank for another something no commercial bank can create has to move.

Chapter sources: McLeay, Radia and Thomas, "Money creation in the modern economy", Bank of England Quarterly Bulletin 2014 Q1; Bank of England explainer "How is money created?"; Bank of England, "The Bank

of England's approach to resolution" (2023), the section 79A report on the transfer of SVB UK to HSBC UK, the ring-fencing pages and CP10/26, and the countercyclical capital buffer pages; PRA CP2/25 and PS22/25 on the leverage ratio retail deposits threshold, PS1/26 on Basel 3.1 final rules, and the news release of 17 January 2025 on the Basel 3.1 delay; Basel Committee on Banking Supervision membership pages and the December 2017 "Basel III: Finalising post-crisis reforms" high-level summary; Financial Stability Institute executive summaries of the LCR and NSFR; Lloyds Banking Group plc 2025 results announcement, 28 January 2026; Federal Reserve Board reserve requirements page and its 28 April 2023 review of the supervision and regulation of Silicon Valley Bank; Deutsche Bundesbank and ECB on minimum reserves, including the Governing Council decision of 27 July 2023; Insolvency Act 1986 Schedule 6 as amended and the Banks and Building Societies (Depositor Preference and Priorities) Order 2014, legislation.gov.uk; Financial Services Compensation Scheme cover pages; Banking Act 2009; Financial Services (Banking Reform) Act 2013.

Two Kinds of Money

6.0 What this chapter gives you

1. You will be able to explain why a payment between two customers of the same bank needs no cooperation from anybody, while an identical-looking payment to another bank eventually requires the Bank of England to change a number in a ledger of its own.
2. You will be able to name the two kinds of sterling — central bank money and commercial bank money — and say exactly who may hold each of them.
3. You will be able to work a day's multilateral net position out of a list of customer payments, and explain why £2,245 of payments moved £450 in the upper ledger.
4. You will be able to explain why drawing £20 from a cash machine converts one kind of money into the other, and what the singleness of money means.
5. You will be able to say why reserves are not depositors' money in storage, and why the total fell by £63 billion in a year without anybody being paid a penny.
6. You will be able to decide whether a payment is genuinely on-us by asking which legal entity holds each account rather than which brand is on the app.
7. You will be able to distinguish the four account types at the Bank of England and say which are remunerated at Bank Rate and which may not hold funds overnight.
8. You will be able to explain why a Faster Payment landing at two on a Sunday morning leaves one bank exposed to another, and what prefunding does about it.
9. You will be able to describe tiering, and say why an indirect participant's settlement asset is commercial bank money rather than central bank money.
10. You will be able to answer "why is this payment behaving perversely" by asking which tier the obligation is sitting in and what has to happen before it moves.

The previous chapter established what a bank is: an institution whose fundamental product is a promise, and whose balance sheet records your deposit as its debt. This chapter takes that one step further and asks an awkward question. If your bank's promise is money to you, whose promise is money to your bank?

The question is not philosophical. It has a specific, mechanical, unromantic answer, and that answer is the hinge on which the rest of this book turns. Almost every rule about settlement that arrives later — why a CHAPS payment cannot be recalled, why Bacs takes three working days, why a Faster Payment reaches the recipient in seconds but the banks square up hours later, why a card refund is slow, why an international transfer travels through institutions nobody asked for — is a consequence of the fact that there are two distinct kinds of money in circulation at the same time, in the same currency, denominated in the same pounds, and that they are not interchangeable in the way the word "pounds" implies.

Consider two payments that look identical from the outside. Priya banks with Lloyds and pays £120 to Sam, who also banks with Lloyds. Priya's neighbour banks with Lloyds and pays £120 to a friend at Barclays. Both payments are for £120, both are sent from a phone, both arrive in seconds, both produce identical-looking lines on identical-looking statements. But the first payment never leaves Lloyds and requires no cooperation from any other institution in Britain. The second cannot be completed by Lloyds alone at any price, and eventually requires the Bank of England to change a number in a ledger of its own.

That asymmetry is invisible to customers, is largely invisible to bank staff, and governs everything. This chapter makes it visible.

The plain version

Imagine a town with two banks. One is run by Ada, the other by Mercer. Between them they hold the accounts of everybody in town.

Ada's customers include Priya, Sam and Tom. Mercer's customers include Ravi and Nadia. Each bank keeps a book with one page per customer, showing how much the bank owes that person. That is what a bank balance is: a written record of a debt owed to you. Nobody has a personal pile of coins in a personal box. There is a book, and there is a line in it with your name on.

■ The easy payment

On Monday, Priya pays Sam £120. Both of them are Ada's customers.

Ada does not need to consult anybody. She opens her book, reduces Priya's line by £120 and increases Sam's line by £120. That is the whole event. Nothing has left Ada's building. No cash has moved, no van has driven anywhere, no message has been sent to any other institution. Ada owed £120 to Priya before, and she owes it to Sam now. Her total debt to her customers is exactly the same as it was five minutes ago. What changed is the label on it.

This is worth stopping on because it is genuinely surprising the first time you see it clearly. From Ada's point of view, a payment between two of her own customers is not a transfer of anything. It is a relabelling exercise. She has rewritten who she owes, and that is all.

■ The hard payment

On Tuesday, Tom pays £800 to Ravi. Tom is Ada's customer. Ravi is Mercer's.

Ada can do her half easily. She reduces Tom's line by £800, and her total debt to her customers falls by £800. She is £800 better off.

Mercer can do his half easily too. He increases Ravi's line by £800, and his total debt to his customers rises by £800. He is £800 worse off.

And that is the problem. Mercer has just taken on a new obligation to Ravi in exchange for nothing at all. Ada has just been released from an obligation, also in exchange for nothing at all. If the story stops here, Ada has been enriched by £800 at Mercer's expense, and Mercer would be an idiot to let it happen again.

So Ada has to give Mercer something worth £800. The obvious candidate is a promise: Ada writes Mercer an IOU. But Mercer does not want Ada's IOU. Ada might be perfectly sound today, but if she is unsound in six months Mercer is left holding a piece of paper and a real debt to Ravi. If the town's banks paid each other in IOUs, every bank would end up quietly lending to every other bank all day long, without ever deciding to, and a single failure would take the lot down.

What Mercer wants is money that does not depend on Ada being solvent.

■ The vault in the market square

So the town builds one. In the market square there is an institution — call it the Old Vault — which is not a bank for people. It is a bank for banks. Ada has an account there. Mercer has an account there. So does every other bank in town, and nobody else.

Now Tuesday's payment can complete. Ada tells the Old Vault to move £800 from her account to Mercer's. The Vault reduces Ada's line by £800 and increases Mercer's by £800. Mercer has been paid in something he trusts, and he is content to owe Ravi the £800, because he has received £800 of something better than an IOU from Ada.

That is the two-tier system, in one picture. There are two ledgers, not one. Down at street level, banks keep books recording what they owe their customers. Up above, the Vault keeps a book recording what it owes the banks. Ordinary payments between customers of different banks cause changes in both books at once, and the payment is not really finished until the upper book has been changed.

■ The two kinds, named

There are therefore two kinds of money in this town, and they are not the same substance even though they are counted in the same pounds.

The first kind is what you and I have. A balance at Ada's Bank is a promise from Ada. It is only as good as Ada. If Ada fails, you are a creditor of a failed business.

The second kind is what banks have. A balance at the Old Vault is a promise from the Vault. Banks accept it from one another without hesitation for three reasons: the Vault is not trying to make money out of them, the Vault cannot run out of the town's currency because the Vault is the thing that issues it, and the Vault is not going to go bust.

In the real United Kingdom the Old Vault is the Bank of England, and the accounts banks hold there are called reserves accounts. The money in them is called central bank money. The money in your account is called commercial bank money. Once you have those two phrases, an enormous amount of otherwise baffling behaviour becomes obvious.

■ One thing ordinary people can hold

There is exactly one form of central bank money that a member of the public can hold, and it is in your wallet. A £20 note is a direct promise from the Bank of England. It is not a promise from Barclays or Nationwide or anybody else. As of the Bank's most recent published figures there are over 4.98 billion Bank of England notes in circulation, together worth about £91.5 billion.

So when you draw £20 from a cash machine, something genuinely interesting happens: you convert £20 of commercial bank money into £20 of central bank money. Your claim on your bank shrinks by £20 and you acquire a claim on the central bank instead. It does not feel like a change of substance because the two spend identically, and the fact that they spend identically — that nobody in any shop in Britain asks whether your £20 is a Lloyds £20 or a paper £20 — is the single greatest achievement of the whole arrangement. It is called the singleness of money, and it is not an accident. It is engineered, and the thing that engineers it is the shared ledger in the middle.

■ A day, with numbers

Take one full day in the town, so the arithmetic is concrete.

At the start of Wednesday, Ada holds £45,000 at the Old Vault and Mercer holds £60,000. Five payments happen.

Payer	Payer's bank	Payee	Payee's bank	Amount
Priya	Ada	Sam	Ada	£120
Tom	Ada	Ravi	Mercer	£800
Ravi	Mercer	Priya	Ada	£250
Nadia	Mercer	Tom	Ada	£1,000
Sam	Ada	Priya	Ada	£75

Five customers have had their balances changed. In total, £2,245 of payments have been made. Every payer's line went down and every payee's line went up by the full amount, exactly as they would expect.

Now look at the Old Vault's book. Two of those five payments — Priya to Sam, and Sam to Priya — were between customers of the same bank. The Old Vault never heard about them. They do not exist as far as the upper ledger is concerned. That leaves three payments that cross between banks: £800 from Ada's side to Mercer's, and £250 plus £1,000 coming back the other way.

Ada owes Mercer £800. Mercer owes Ada £1,250. There is no reason to move both amounts, so the Vault moves the difference: £450 from Mercer to Ada. Ada finishes the day with £45,450 and Mercer with £59,550. The Vault's total is unchanged at £105,000, because nothing was created or destroyed up there — one bank's pile grew and another's shrank by the same amount.

Look at what that means. Five payments, £2,245 of customer activity, and the number that actually changed hands in the upper ledger was £450, once. The rest either cancelled out or never left the bank it started in.

That single observation explains more about how payment systems are built than anything else in this book. It is why banks care so much about which of their customers pay each other. It is why the machinery that works out the £450 exists at all. And it is why, when things go wrong, the question that matters is never “did the customer's balance change” but “did the upper ledger change, and when”.

Where the plain version stops being true

The vault-in-the-square picture is the right mental model, and practitioners carry a version of it in their heads. But four things about it are wrong in ways that will bite.

■ The reserves are not your money in storage

The picture invites you to imagine that Ada's £45,000 at the Vault is the town's real money, kept safe on behalf of Ada's depositors, and that Ada's deposits are somehow backed by it. That is not what reserves are, and the error matters because it produces confident nonsense about banking.

Reserves are a liability of the central bank, created by the central bank, in whatever quantity the central bank chooses. They are not a warehouse receipt for anything. There is no ratio that must hold between a bank's reserves and its deposits — the United Kingdom has had no mandatory reserve ratio for decades. The two figures are not even the same order of magnitude. At end-February 2026, total balances held in reserves accounts at the Bank of England stood at £643,547 million, while the

UK's broad money measure M4, which is overwhelmingly commercial bank deposits, was running at around £3.3 trillion. Reserves cover something in the region of a fifth of deposits, and that fraction is not a design parameter. It is a residue of monetary policy.

You can see this directly. Reserves balances fell from £706,745 million at end-February 2025 to £643,547 million a year later. No bank paid that £63 billion to anybody. It disappeared because the Bank of England was shrinking its own balance sheet through the sale and maturity of assets held in the Asset Purchase Facility. The upper ledger's total is not a fact about the banks. It is a fact about the central bank.

The other half of the correction: reserves cannot leave the upper tier. A bank cannot lend reserves to you. Reserves circulate only among institutions that hold accounts at the central bank, plus a leak into banknotes when the public wants cash. When a bank makes a loan it creates a deposit — commercial bank money — out of nothing, as the Bank of England's own *Money creation in the modern economy* set out in its 2014 Quarterly Bulletin. It does not hand over reserves. Reserves are what it needs afterwards, if and when the borrower spends the money at a customer of a different bank.

■ “Same bank” is a legal question, not a branding question

The plain version says a payment between two customers of the same bank never leaves that bank. That is true, and it is the most useful sentence in the chapter. But “the same bank” means the same legal entity, and British banking brands routinely span several.

Since ring-fencing, the large UK groups are split. Barclays operates both Barclays Bank UK PLC and Barclays Bank PLC as separately authorised entities. HSBC operates HSBC UK Bank plc alongside HSBC Bank plc. Lloyds Banking Group contains Lloyds Bank plc, Bank of Scotland plc and Lloyds Bank Corporate Markets plc. A payment from a retail customer of one entity to a corporate customer of another entity in the same group is not an on-us payment. It has an interbank obligation in it, and it settles like any other.

The second half of this correction is about risk rather than plumbing. Even a genuine on-us payment does something real: it moves the exposure. Before Priya pays Sam, Priya is the bank's creditor for £120. Afterwards, Sam is. If the bank fails, it matters enormously which of them is holding the claim, because deposit protection is per depositor, per authorised firm. In the United Kingdom that limit rose from £85,000 — the level set in 2017 — to £120,000 on 1 December 2025, with the temporary high balance limit for qualifying life events rising from £1 million to £1.4 million on the same date. An on-us payment that pushes a customer from £90,000 to £150,000 has quietly moved £30,000 out of protection. Nothing settled, nothing cleared, and the risk profile of the bank's deposit book changed.

■ Nothing moves at the Old Vault when your payment arrives

The picture implies that a cross-bank payment causes an immediate transfer in the upper ledger. For one UK payment system that is true. For the systems that carry almost all the volume, it is false, and the gap is where a whole category of risk lives.

CHAPS settles gross and in real time: each payment moves central bank money individually, when it is made. Everything else settles net and later. Faster Payments net settlement takes place three times each business day. Bacs settles once a day. The Image Clearing System for cheques settles once a day. Visa Europe and Mastercard Europe settle on twenty-four-hour cycles, with weekend and bank holiday cycles settling the following business day. PEXA settles six times a business day. LINK settles daily.

So when a Faster Payment lands in a stranger's account at two in the morning on a Sunday, the recipient's bank has credited a real customer with real, spendable money while holding nothing but an obligation from the sending bank, and it will hold that obligation until the next settlement cycle. That interval is a credit exposure of one bank to another, created by a customer, without either bank's credit committee being consulted. The Bank of England's answer is prefunding: it offers a service that lets a payment system operator cap each participant's maximum possible net debit position at the amount that participant has locked in a segregated prefunding account in RTGS. If the participant fails, the cash is already there and settlement completes anyway.

The general rule to carry forward is this. A customer payment and its settlement are two different events with two different timestamps. Assuming they are the same event is the single most common error in payments engineering, and it is the reason Chapter 7 exists.

■ There is not one account per bank, and most institutions have none

The picture has a tidy one-to-one mapping: one bank, one account at the Vault. Reality has four account types with different rights, and most regulated firms have none of them.

The Bank of England distinguishes reserves accounts, reserves-and-settlement accounts, settlement accounts, and omnibus accounts. They differ on who may hold them, whether they are remunerated, and whether balances may sit overnight. A settlement account held by a non-bank payment service provider is unremunerated and may not hold funds beyond what settlement requires. A reserves account is remunerated at Bank Rate and may hold balances overnight. That distinction is not administrative trivia; it decides whether an institution earns anything on its central bank balances, which is a first-order fact about its business model.

And the population is small. The Prudential Regulation Authority regulates around 1,300 firms. At end-February 2026, 221 participants had access to Sterling Monetary Framework facilities. More than 70 organisations use RTGS to settle directly in one or more payment systems. Just over 35 organisations access CHAPS directly. Several thousand financial institutions reach CHAPS indirectly, by sending their payments through one of those direct participants.

For all of those indirect institutions, the two-tier picture repeats one level down, and the top tier is not the central bank. A small firm that clears through Barclays settles in a Barclays deposit. Its settlement asset is commercial bank money. It has the same exposure to Barclays that Priya has, and it has it in size. This is called tiering, it is a standing concern of financial stability regulators, and it is the mechanism behind correspondent banking and the nostro and vostro accounts that Volume IV takes apart.

The technical version

■ Definitions

Central bank money is a liability of the central bank. In sterling it takes two forms: Bank of England banknotes, which are bearer liabilities available to the general public, and balances held in accounts at the Bank of England, which are book-entry liabilities available only to institutions that meet the Bank's access criteria. Central bank money is the settlement asset for sterling: it is what an institution receives when a sterling obligation to it is finally discharged.

Commercial bank money is a deposit liability of a commercial bank, building society or credit union — an unsecured, unsubordinated claim on that institution, ranking as a depositor claim in insolvency and covered, up to the applicable limit, by the Financial Services Compensation Scheme. It is the

money the public holds, the money businesses hold, and by value it is the overwhelming majority of the money in existence.

The two-tier system is the arrangement in which the central bank provides accounts to a restricted set of institutions, those institutions provide accounts to the public, and the public's money is therefore a claim on a private firm while the private firms' money is a claim on the state. The Bank for International Settlements described this in its September 2025 report on wholesale central bank money as a "time-tested two-tier structure", and identified the central bank settlement asset as what preserves the *singleness* of money: the property that payments denominated in the sovereign unit of account settle at par regardless of which form of money is used. The foundational treatment remains the Committee on Payment and Settlement Systems report *The role of central bank money in payment systems*, published on 12 August 2003.

	Central bank money	Commercial bank money
Issuer	Bank of England	Authorised deposit-taker
Forms	Banknotes; reserves and settlement balances	Current and deposit account balances
Who may hold it	Public (notes only); eligible institutions (balances)	Anyone the bank will onboard
Credit risk	Effectively none in sterling	Full unsecured exposure to the issuer
Protection on failure	Not applicable	FSCS, £120,000 per depositor per firm since 1 December 2025
Approximate sterling stock	£91.5bn notes; £643.5bn reserves at end-February 2026	M4 around £3.3 trillion
Remuneration	Reserves at Bank Rate; notes zero	Set by the bank

■ Accounts at the Bank of England

Access is governed by the Bank's published access policy for RTGS settlement accounts and services. Four account types exist.

A **reserves account** is available to PRA-authorised UK-incorporated banks, building societies and investment firms designated by the PRA for prudential supervision, to UK subsidiaries and branches of their non-UK incorporated equivalents, and to central counterparties and international central securities depositories. Balances are remunerated at Bank Rate — 3.75% at end-February 2026, having fallen from 4.5% over the preceding twelve months — and may be held overnight. Holding a reserves account requires participation in the Sterling Monetary Framework.

A **reserves and settlement account** is the same account used for both purposes: holding funds and settling obligations arising from direct participation in a payment system that settles in RTGS.

A **settlement account** is for firms not eligible for reserves: FCA-authorised non-bank payment service providers, systemically important financial market infrastructures, and entities within a wider banking group that already holds a reserves account. Settlement accounts are unremunerated and may not hold funds beyond what is needed to support settlement. The Bank opened this route to non-bank PSPs in 2017, making it the first G7 central bank to do so.

An **omnibus account**, offered since 2021, allows the operator of a payment system recognised by HM Treasury under the Banking Act 2009, and designated under the Financial Markets and Insolvency (Settlement Finality) Regulations 1999, to co-mingle reserves belonging to multiple participants in

a single account, so that balances in the operator's own ledger are backed by central bank money. Omnibus accounts are remunerated at Bank Rate. The first user was the Sterling Finality Payment System, which processed its first live payments on 14 December 2023.

■ RTGS: the ledger that holds the upper tier

The Bank of England's Real-Time Gross Settlement infrastructure was introduced in 1996. It is, in the Bank's own words, "an accounting system" — the ledger that underpins settlement of sterling payments, rather than a payment system in its own right. Responsibility for operating CHAPS, the UK's high-value payment system, transferred to the Bank in November 2017, so that the Bank now acts as both settlement agent and system operator. On 28 April 2025 the Bank completed a multi-year renewal programme, moving to a new core ledger and settlement engine known as RT2; the participant-facing interface Enquiry Link was replaced by BERTI at the same time.

RTGS provides settlement for CHAPS, for the payment system embedded within CREST, for seven retail payment systems — Bacs, the Image Clearing System, Faster Payments, LINK, Mastercard Europe, PEXA and Visa Europe — and for the Bank's Note Circulation Scheme. It supports four settlement models: real-time gross settlement, delivery versus payment, prefunded deferred net settlement and standard deferred net settlement.

Average daily settlement in RTGS, Q2 2026:

Settlement service	Daily average value (£mn)
CHAPS	411,562
CREST DvP	477,841
Bacs (net)	6,205
Visa Europe (net)	2,356
Faster Payments (net)	2,298
Mastercard Europe (net)	1,480
LINK (net)	238
Cheque imaging (net)	37
PEXA (net)	1
Total	902,018

Two things in that table deserve attention. The first is the gulf between CHAPS at £411.6 billion a day and Faster Payments at £2.3 billion a day, when Faster Payments carries vastly more transactions. That is netting, and Chapter 8 is about it. The second is that the entire retail payments experience of sixty-odd million people — every Faster Payment, every direct debit, every card transaction settling through Visa Europe or Mastercard Europe, every cash machine withdrawal — reduces to about £12.6 billion of central bank money changing hands each day, roughly three per cent of the RTGS total.

■ The on-us case, stated exactly

An intrabank payment — variously called an on-us payment, an internalised payment, or a book transfer — is a payment in which the payer's account and the payee's account are held by the same legal entity.

The accounting is a single pair of postings within that institution's own ledger. The payer's account is debited and the payee's account is credited. Both are liability accounts. The bank's balance sheet total is unchanged; only the composition of its deposit liabilities changes. There is no interbank obligation, and therefore nothing to clear, nothing to net and nothing to settle. No reserves move. No message

reaches the Bank of England. The payment is complete at the moment the bank's ledger says it is complete, and its finality is a matter of that bank's own terms and systems rather than of any scheme rulebook or settlement finality designation.

Whether the bank routes such a payment through the relevant scheme's central infrastructure anyway — for addressing, for Confirmation of Payee, for tracking, or simply because it is easier to have one code path — is an implementation choice rather than a settlement necessity, and the Payment Systems Regulator says as much in its reporting guidance on authorised push payment scams. That guidance defines an on-us transaction as a push payment between two account holders within the same provider group, and records that, depending on a provider's commercial arrangements, such payments may pass either externally via Faster Payments or internally as a book transfer. Firms are required to report both, and to itemise the internally booked ones separately so that the regulator can check its totals against Pay.UK's own Faster Payments figures. The settlement consequence, however, is not in doubt: an on-us payment generates no interbank obligation, and no quantity of central bank money moves as a result of it.

The best evidence that this is a material phenomenon rather than a curiosity is that the Bank of England's own published statistics have to carve it out. The Bank states that CHAPS represents 0.4% of UK total payment volumes but 91% of total sterling payment values, and appends the qualification "excluding internalised flows within payment service providers". The official numbers exclude on-us traffic because it never touches the settlement layer and therefore cannot be observed there. A payment that does not settle is, from the settlement system's point of view, a payment that did not happen.

For a practitioner this has consequences beyond accounting. On-us payments do not attract interbank settlement obligations or, in card schemes, interchange between different institutions. They are not visible in scheme-level reconciliation files. They may be subject to different internal cut-offs, different limits and different fraud controls from their interbank equivalents. And in a large bank with a high market share, they can be a very large fraction of volume — which is exactly why a payments platform must be designed so that the on-us and interbank paths cannot diverge in behaviour that customers can see.

■ The interbank case, stated exactly

When payer and payee bank with different institutions, the payment creates an obligation between those institutions that must be discharged in the settlement asset. There are two ways to do it.

Real-time gross settlement. Each payment is settled individually, in central bank money, at the moment it is processed. In sterling this model is used only by CHAPS. Settlement is final and irrevocable between direct participants as it occurs, which is why CHAPS is the instrument of choice for house purchases and for the sterling leg of transactions that must not unwind: CLS, the foreign exchange settlement system, completes its final sterling settlement through direct participation in CHAPS, and LCH and Euroclear Bank are also direct participants.

The CHAPS numbers for calendar year 2025: volumes grew 1.1% to a record 53.3 million payments, an average of 210,482 each working day; total value settled rose 7.3% to £93.9 trillion, an average of £371.3 billion daily. The average payment value was £1.8 million and the median was £4,586 — a spread that tells you a great deal, since 94% of CHAPS payments were for £1 million or less while financial institution payments, carried in the pacs.009 message, were 24% of volume but 73% of value. The Bank observes that CHAPS settles the annual UK GDP every nine working days. The all-time peak volume day was 2 April 2024 with 344,099 payments; the all-time peak value day was 3 October 2022 with £642.7 billion.

The CHAPS settlement day currently runs 06:00 to 18:00, Monday to Friday. Direct participants must be able to receive by 08:00 and to send by 10:00; the cut-off for customer payments is 17:40. Participants choose whether to submit a payment as urgent, settling immediately subject to available liquidity, or non-urgent, accepting some delay in exchange for better liquidity efficiency through the Liquidity Saving Mechanism; non-urgent payments settle around eight minutes after submission on average, and are auto-promoted to urgent after 17:30. From September 2027 the Bank will open CHAPS settlement from 01:30 rather than 06:00, and it has consulted on weekend settlement not before 2029 and a 22-hours-by-6-days pattern not before 2031, with 22x7 and 23.5x7 identified as possible end states.

Deferred net settlement. The system operator collects the individual payment instructions, calculates each participant's multilateral net position at the end of a cycle, and instructs the Bank of England to settle a single figure per participant. This is dramatically cheaper in liquidity than gross settlement, and it is what makes high-volume, low-value payments economic. The cost is that customers are credited and debited before interbank settlement occurs, so participants carry an intraday exposure to one another. Prefunding removes that exposure by capping each participant's maximum net debit at cash it has already segregated in RTGS.

■ The legal layer

Two statutory mechanisms turn these arrangements from private contracts into something that survives an insolvency.

Recognition under the Banking Act 2009 brings a payment system within the Bank of England's supervisory remit, on a recognition order made by HM Treasury. Designation under the Financial Markets and Insolvency (Settlement Finality) Regulations 1999 protects transfer orders and netting arrangements within a designated system from being unwound by ordinary insolvency law — so that a settlement which has occurred stays occurred even if a participant enters administration later the same day. Both are prerequisites for an omnibus account, and both are the reason the phrase “settled in central bank money” carries legal weight rather than merely describing where the entry was made.

■ The same structure elsewhere

The two-tier arrangement is not a British peculiarity. Every major currency is organised this way, with local variations in access.

In the United States, the Federal Reserve Banks provide master accounts to legally eligible depository institutions, which are the gateway to Fedwire Funds, Fedwire Securities, the National Settlement Service, FedACH, FedCash and FedNow. Deposits of depository institutions at the Reserve Banks stood at \$2,947,630 million on the H.4.1 release of 13 August 2026. In May 2026 the Federal Reserve proposed a new, more limited “payment account” as a route for institutions that do not qualify for a full master account.

In the euro area, the Eurosystem operates TARGET Services. In 2025, T2 settled a daily average of 431,067 euro payments worth €1,932.8 billion, a total of €492,859.0 billion for the year, with an average transaction value of €4.5 million; the ECB notes that T2 settles the equivalent of euro-area GDP in around eight days of operations. The instant payment service TIPS settled a daily average of 2,735,053 euro payments worth €2.2 billion — the same inversion of volume and value that Faster Payments shows against CHAPS.

■ Central bank digital currency, briefly

A retail central bank digital currency would be central bank money held directly by households and businesses in electronic form. That is the entire proposition, and it explains why the topic generates the heat it does: it would be the first genuinely new form of tier-one money available to the general public since the banknote, and it would change the answer to the question this chapter is built on — who may hold a claim on the central bank.

The Bank of England and HM Treasury are in a design phase for a digital pound, examining technology and policy requirements. No decision to build or issue has been taken. The Bank has said the earliest it would issue a digital pound is the second half of this decade, that a decision on next steps would come in 2026, and that Parliament would have a say before any launch. It has also said that a digital pound would not replace cash, and that the Bank will continue to issue banknotes for as long as people want to use them. A Digital Pound Lab for practical experimentation was announced in October 2025.

The arguments are worth stating without taking sides. Proponents point to declining cash use, to the desirability of a public payment option that does not depend on any commercial firm, and to the anchoring role of a monetary asset available to everyone at par. Opponents and cautious central bankers point to disintermediation — if households can hold central bank money directly, deposits may leave banks, particularly in a stress — and to privacy, which is why most design work assumes holding limits and an intermediated model in which private firms, not the central bank, hold the customer relationship. These are open questions and this book does not pretend to settle them.

On the wholesale side, the framing is different, because wholesale central bank money already exists and is already digital: that is precisely what a reserves balance is. The live work there is about the form it takes rather than its existence — whether reserves can be represented on programmable ledgers so that a securities transfer and its cash leg settle atomically. Omnibus accounts are one answer already in production; the BIS report of 18 September 2025, produced with seven central banks, is the reference for the rest.

■ Why this chapter governs the rest of the book

Almost every settlement rule that follows is an answer to two questions: in whose money, and when.

CHAPS is irrevocable because it settles in central bank money at the moment of processing, and there is no mechanism by which a completed transfer of the settlement asset can be reversed by one party's regret. Bacs takes three working days because it is a deferred net batch system whose cycle was designed around the exchange of files, and the customer-visible timing is inherited from the settlement design rather than from any technical necessity. Faster Payments has a value cap because the recipient's bank extends credit until the next of the three daily settlement cycles, and that exposure has to be bounded. Card refunds are slow because the refund is a new clearing item that must travel back through the scheme and settle in its own cycle, not a reversal of the original. International transfers acquire a chain of intermediaries because there is no single central bank whose money both ends can hold, so the chain is a sequence of commercial bank claims stitched together — the two-tier problem repeated at each hop. Safeguarding rules exist for e-money and payment institutions precisely because their customers' balances are not deposits and their firms are not banks, so the money must be held somewhere that survives the firm.

None of those rules is arbitrary. Each is a consequence of a decision about which kind of money discharges an obligation and at what moment. Whenever a payment behaves in a way that seems perverse, the productive question is not “why is this system badly designed” but “which tier is this obligation sitting in right now, and what has to happen before it moves”.

6.98 Common wrong ideas

Wrong: Reserves are depositors' money kept safe at the Bank of England. **Right:** Reserves are a liability of the central bank, created in whatever quantity it chooses, and no ratio has to hold between them and anybody's deposits.

Wrong: Banks lend out reserves. **Right:** Reserves cannot leave the upper tier; a bank lends by creating a deposit and needs reserves only afterwards, if the borrower pays a customer of a different bank.

Wrong: Two accounts showing the same brand are with the same bank. **Right:** Barclays, HSBC and Lloyds each run several separately authorised entities, and a payment between two of them carries a real interbank obligation.

Wrong: An on-us payment has no consequences because nothing settles. **Right:** It moves the exposure, and it can quietly push a customer past the £120,000 protection limit without anything clearing or settling at all.

Wrong: A cross-bank payment moves central bank money at the moment it arrives. **Right:** Only CHAPS settles that way; Faster Payments settles three times each business day, and Bacs, cheques and the card schemes once.

Wrong: Every bank has an account at the Bank of England. **Right:** The PRA regulates around 1,300 firms, just over 35 organisations access CHAPS directly, and everybody else reaches it through somebody's balance sheet.

Wrong: All Bank of England accounts are the same. **Right:** Reserves accounts are remunerated at Bank Rate and may hold balances overnight; settlement accounts for non-bank payment service providers are unremunerated and may not.

Wrong: Faster Payments matters more than CHAPS because it carries far more traffic. **Right:** In Q2 2026 CHAPS settled £411.6 billion a day against Faster Payments' £2.3 billion, because netting compresses the retail systems before they reach the Bank.

Wrong: A digital pound has been decided on and will replace cash. **Right:** The Bank and HM Treasury are in a design phase, no decision to build or issue has been taken, and the Bank has said it will keep issuing banknotes for as long as people want them.

Wrong: Banknotes are issued by your bank. **Right:** A £20 note is a direct bearer liability of the Bank of England, and it is the only form of central bank money a member of the public may hold.

6.99 Chapter summary in 20 lines

1. Your bank's promise is money to you, but your bank needs money that does not depend on any other bank staying solvent.
2. That requirement produces a two-tier system: banks keep books recording what they owe customers, and the central bank keeps a book recording what it owes banks.
3. A payment between two customers of the same bank is a relabelling exercise inside one ledger and requires nobody else's cooperation.
4. A payment between customers of different banks releases one bank from an obligation and burdens the other, so something of value must pass between them.
5. That something is central bank money, held in reserves and settlement accounts at the Bank of England.

6. Commercial bank money is an unsecured claim on a private firm; central bank money carries effectively no credit risk in sterling.
7. Banknotes are the only form of central bank money the public may hold, and there are over 4.98 billion of them worth about £91.5 billion.
8. That both kinds spend identically is called the singleness of money, and it is engineered by the shared ledger in the middle rather than being an accident.
9. In the worked day, £2,245 of customer payments across two banks moved £450 once in the upper ledger, and two of the five payments never left the bank they started in.
10. Reserves are not deposits in storage: no ratio binds them, and the total fell from £706,745 million to £643,547 million in the year to February 2026 without anybody being paid.
11. Reserves cannot leave the upper tier, so a bank lends by creating a deposit and worries about reserves only when the money walks out of the door.
12. “The same bank” is a legal question rather than a branding question, and ring-fencing has split the large British groups into several authorised entities.
13. Even a genuine on-us payment moves risk, because deposit protection is per depositor per authorised firm, raised to £120,000 on 1 December 2025.
14. Only CHAPS settles gross and in real time; everything else settles net and later, and that interval is a credit exposure created by a customer.
15. The Bank of England’s answer is prefunding, which caps a participant’s maximum net debit at cash already segregated in RTGS.
16. There are four account types at the Bank — reserves, reserves and settlement, settlement, and omnibus — differing on eligibility, remuneration and overnight balances.
17. Access is narrow, and the several thousand institutions that reach CHAPS indirectly settle in a direct participant’s commercial bank money instead.
18. That is tiering: the two-tier picture repeats one level down with a commercial bank at the top of it.
19. RTGS settled about £902 billion a day in Q2 2026, of which the entire retail payments experience of the country was roughly £12.6 billion.
20. Every settlement rule in the rest of the book answers two questions — in whose money, and when — so the productive question about odd behaviour is which tier the obligation is sitting in.

Sources used: Bank of England — A brief introduction to the Real-Time Gross Settlement system and CHAPS; Payment and settlement; Payment and settlement statistics; Access policy for RTGS settlement accounts and services; Bank of England Settlement Accounts; CHAPS; Banknote statistics; Report on the Bank’s official market operations March 2025–February 2026; RTGS and CHAPS Annual Report 2024–25; Reviewing access to RTGS accounts for settlement (2024); Extending RTGS and CHAPS settlement hours – next steps (2026); Money creation in the modern economy, Quarterly Bulletin 2014 Q1; The digital pound; PRA news release and PS24/25 on depositor protection (18 November 2025); Which firms does the PRA regulate. Bank for International Settlements — Wholesale central bank money in the context of tokenisation (18 September 2025); CPSS, The role of central bank money in payment systems (12 August 2003). European Central Bank — TARGET Services Annual Report 2025; T2 facts and figures. Federal Reserve — H.4.1 statistical release, 13 August 2026; Federal Reserve Financial Services, Master Account; Federal Register, proposed revisions to the Policy on Payment System Risk and account access guidelines (26 May 2026). Pay.UK — Faster Payment System. Fnlity — Sterling Fnlity Payment System announcement, 14 December 2023.

Clearing Versus Settlement

7.0 What this chapter gives you

1. You will be able to say precisely what clearing is, what settlement is, and why “the payment went through” carries almost no information.
2. You will be able to name the four separate events that sentence collapses, and say whose record holds each one.
3. You will be able to explain why a Faster Payment credited at two o’clock on a Sunday morning is not settled between the banks until Monday.
4. You will be able to state the trade-off between real-time gross settlement and deferred net settlement as an exchange of liquidity for settlement risk.
5. You will be able to explain what happened to Bankhaus Herstatt’s counterparties on 26 June 1974, and why the loss was the whole principal rather than a fraction of it.
6. You will be able to say why the New York clearing process then ran until one in the morning for several nights running, and what that failure mode is called.
7. You will be able to explain why finality is a legally defined moment under the Settlement Finality Regulations 1999 rather than a feeling.
8. You will be able to keep finality between direct participants distinct from consumer reimbursement for authorised push payment fraud.
9. You will be able to explain why a clearing file and a settlement advice will never tie out item by item, and what reconciliation therefore has to prove.
10. You will be able to decide, before releasing goods against an incoming payment, which of the four events has actually occurred and who can tell you.

Ask anybody who works in payments what the hardest thing is about explaining their job, and a good number of them will give you the same answer. It is that the public has one word — “payment” — for what the industry treats as at least four separate events, occurring at four different times, on four different ledgers, with four different sets of consequences if something goes wrong.

The most important of those distinctions is the one this chapter is about. Clearing is the process of working out who owes what to whom. Settlement is the money actually moving. They are not the same thing, they very rarely happen at the same moment, and almost every strange behaviour in the payments world — the three-day wait, the payment that “arrived” on Sunday but was not paid for until Monday, the refund that takes a week, the transaction that shows as pending and then vanishes — is a consequence of the gap between them.

The previous chapter established that there are two kinds of money in circulation: commercial bank money, which is what sits in your account and is a promise from your bank, and central bank money,

which is what banks hold in reserves accounts at the Bank of England. That distinction is the raw material of this one. Clearing happens in the world of obligations. Settlement happens in the world of central bank money. Clearing produces a number. Settlement makes the number true.

The phrase “the payment went through” contains no information about which of these has happened. That is not pedantry. In 1974 the difference between them destroyed a German bank’s counterparties, froze the New York clearing process for several nights running, and led directly to the creation of the international committee structure that still governs payment system oversight today. The word for the risk is named after that bank.

The plain version

Imagine two shops facing each other across a small high street. On one side is a bakery run by Nadia. On the other is a hardware shop run by Owen. Their staff buy from each other constantly. The bakery’s staff need string, bin bags, light bulbs and a new mop. The hardware shop’s staff want coffee, sandwiches and a birthday cake.

At first they paid each other in cash every time. Somebody would walk across the road with a five-pound note, wait for change, and walk back. It worked, but it was tedious, and both tills ended up short of coins at awkward moments.

So they changed the arrangement. Now, when somebody from the bakery buys a mop, Owen writes it in a notebook by the till: “Bakery, mop, £8.50.” When somebody from the hardware shop buys a cake, Nadia writes it in hers: “Hardware, cake, £22.” Nobody hands over money at the moment of purchase. Goods change hands, and a note gets made.

At six o’clock, one of them walks across the road with their notebook and they sit down together for two minutes.

■ Agreeing the list

The first thing they do is read out the entries and check that both notebooks agree. Owen has “Bakery, mop, £8.50” and Nadia has the same. Owen has “Bakery, bin bags, £4.20” — Nadia has that too. Then Owen reads out “Bakery, extension lead, £11” and Nadia frowns, because she has no record of anybody buying an extension lead. They talk about it. It turns out one of Nadia’s Saturday staff bought it and forgot to write it down. Fine. Both notebooks now agree.

This step — reading out the entries, comparing them, resolving the ones that do not match, and arriving at a list both sides accept — is clearing. Nothing has been paid. Not a penny has moved. What has been produced is agreement: a shared, undisputed statement of who owes what.

That is genuinely all clearing is. It is bookkeeping and argument. It is the part of a payment where two institutions establish, to their mutual satisfaction, the size of the debt between them.

■ Moving the money

The second thing they do takes ten seconds. Nadia’s total for the day is £46.30 owed to Owen. Owen’s total is £38.90 owed to Nadia. Rather than one of them handing over £46.30 and the other handing back £38.90, they take the difference. Nadia opens the till, counts out £7.40, and puts it in Owen’s hand.

That is settlement. The obligation has been discharged. Whatever happens tomorrow — if the bakery burns down, if Owen is run over by a bus, if either of them stops trading — nobody can come back and

claim that day's trading was not paid for. It was. There is a £7.40 gap in Nadia's till and £7.40 more in Owen's, and that is not an entry in a notebook. It is the money.

Notice how much smaller the settlement is than the trading. Over the day, £85.20 of goods changed hands. The amount of actual cash that moved was £7.40. That compression is the whole reason the notebook system exists, and it is the subject of the next chapter. What matters here is the sequence: agree first, pay after.

■ The same thing, with banks

Now replace Nadia with Barclays and Owen with Lloyds, and replace mops and cakes with customer payments.

On Monday, four of Barclays' customers send money to Lloyds customers, and three of Lloyds' customers send money to Barclays customers.

Direction	Amounts	Total
Barclays customers to Lloyds customers	£180, £2,400, £65, £1,355	£4,000
Lloyds customers to Barclays customers	£900, £2,750, £150	£3,800

Seven payments. £7,800 of money as the customers experience it — seven people saw a debit, seven people saw a credit, and every one of those amounts was correct and complete from the customer's point of view.

Clearing is the process by which the two banks exchange the records of those seven payments, check that each side has the same list, sort out anything that does not match, and arrive at a single agreed statement: Barclays owes Lloyds £4,000, Lloyds owes Barclays £3,800, and therefore Barclays owes Lloyds £200.

Settlement is the moment somebody moves £200.

And here is the part that surprises people the first time they hear it. In several of the systems that move most of Britain's money, that £200 does not move at the moment the customers see their balances change. It moves later. Sometimes hours later. Sometimes three days later. In the meantime, the receiving bank has given real, spendable money to its customer on the strength of a promise from the sending bank that has not yet been honoured.

■ Two clocks

There are broadly two ways to organise this, and Britain runs both.

The first way is the notebook. Let everything accumulate, add it all up at a set time, and move one figure. This is cheap, because you move money once instead of thousands of times, and it needs very little cash sitting around waiting. Its cost is the gap. Between the moment a customer's payment is accepted and the moment the banks square up, one bank is exposed to the other. In Britain, Bacs works this way, and so does the Faster Payments Service, and so do the cheque and card systems.

The second way is to settle every payment on its own, one at a time, as it happens. No notebook, no waiting, no gap. Its cost is that each bank must have enough money sitting in its account at the Bank of England to cover each payment at the moment it is made, which is expensive, because money sitting in an account is money not being used for anything else. In Britain, CHAPS works this way — the system

used for house purchases, corporate treasury and the very large payments financial institutions make to each other.

Those two designs have names, and the names are worth learning because they will appear throughout the rest of this book. The notebook way is deferred net settlement. The one-at-a-time way is real-time gross settlement, usually shortened to RTGS.

■ Why the gap is dangerous

In June 1974 there was a bank in Cologne called Bankhaus Herstatt. It was not large. It had, however, been trading foreign currency very badly, and on 26 June the German banking supervisor took its licence away.

The timing is the whole story. The licence was withdrawn during the banking day in Germany but after the German interbank payment system had closed for the afternoon. Herstatt's counterparties — other banks around the world that had agreed to swap Deutschmarks for dollars with it — had already, that morning, paid their Deutschmarks into Herstatt's German account. That was their side of the bargain, done, irrevocably.

The dollars were supposed to come back to them in New York later the same day, because New York is five or six hours behind Cologne and its business day was still going. At half past ten New York time, Herstatt's New York correspondent bank stopped making dollar payments on its behalf.

So a set of perfectly ordinary banks had handed over the full amount of one currency and received nothing at all in return. Not a reduced amount. Not a delayed amount. Nothing, with a claim against a liquidated German bank as the only consolation.

What happened next was almost worse. Banks in New York, suddenly aware that paying first and hoping was a losing strategy, started holding back their own outgoing payments until they could see the incoming ones arrive. Everybody doing that at once meant nobody could go first. The clearing process in New York seized up and, on two or three consecutive nights, ran until one in the morning before it finished.

That is why the industry has a special word for this. When one party to a currency exchange pays out what it owes and does not receive what it is due, the loss is not a fraction of the trade. It is the entire principal. The industry calls it Herstatt risk, after a bank that had been out of business for a day.

■ What this means for “the payment went through”

Put the pieces together and you can see why that sentence carries almost no information.

It might mean the sending bank has checked the payment and taken the money off its customer. It might mean the two banks have agreed that one owes the other. It might mean the money has actually moved between them at the Bank of England. It might mean the receiving customer can see it and spend it. Those are four different facts, and depending on which payment system you are using they occur in different orders, minutes or days apart.

If somebody tells you a payment went through and you are about to hand over a car, a house or a shipping container, the useful question is not “has it gone through”. It is “has it settled, and where can I see that”.

Where the plain version stops being true

The notebook picture is the right shape, and practitioners carry something like it in their heads. Four things about it are wrong in ways that cause real errors.

■ Clearing is not a meeting at six o'clock

The high-street analogy makes clearing sound like a discrete event at the end of the day: two parties sit down, compare lists, agree. Nothing in a modern payment system works like that.

Clearing is a continuous pipeline, running all day, made of several distinct stages that the analogy compresses into one. A payment instruction is transmitted. It is validated for format and for the existence of the destination account. It is screened against sanctions lists and fraud rules. It is matched against the counterparty's record of the same instruction. It is placed in a queue and given a position in the sequence. Only at the end of that pipeline does it contribute to a figure that anyone would recognise as "what we owe you".

The Committee on Payments and Market Infrastructures, which is the body at the Bank for International Settlements that maintains the industry's shared vocabulary, defines clearing as "the process of transmitting, reconciling and, in some cases, confirming transactions prior to settlement, potentially including the netting of transactions and the establishment of final positions for settlement". Read that carefully. Netting is optional — "in some cases", "potentially including". Clearing is the whole pipeline. Netting is one thing that may happen inside it.

The same glossary then adds a sentence that ought to be printed on the wall of every payments team in the country: "Sometimes this term is also used (imprecisely) to cover settlement." The standards body itself is warning you that the industry misuses its own word. When somebody says a payment has "cleared", you cannot safely infer anything about whether money has moved.

■ Netting and clearing are not the same idea, and RTGS still clears

The analogy makes it look as though clearing exists in order to produce a net figure, and therefore that a system which settles every payment individually has no clearing at all. Both halves of that are wrong.

CHAPS payments settle one by one, gross, for full value, with no netting anywhere in the process. CHAPS nevertheless has a substantial clearing function: messages are validated, sequenced, checked against liquidity, queued when a participant is short, and released when it is not. What CHAPS does not do is defer. Clearing and settlement in CHAPS are separated by fractions of a second rather than by hours, which makes them easy to conflate, but they remain two distinct operations and the system's rules treat them as such.

Conversely, netting can happen without clearing in any interesting sense — two counterparties can agree bilaterally to offset positions under a contract with no infrastructure involved at all.

The useful axis is not "netted or not". It is when settlement occurs relative to the obligation, and in what asset. Netting is the subject of the next chapter and deserves its own treatment.

■ Your balance changing is not settlement

This is the correction with the largest practical consequences, and it is the one that catches out software teams building on top of payment rails.

When your bank shows the money as arrived, your bank has updated its own ledger — the record of what it owes you. That is commercial bank money, and it is entirely under your bank's control. It required no cooperation from any other institution and no involvement from the Bank of England. Under Faster Payments, this typically happens within seconds, at any hour, on any day of the year.

Settlement is a different ledger entirely: the settlement accounts held at the Bank of England, in central bank money. Faster Payments net settlement takes place three times each business day. So a payment

received at two o'clock on a Sunday morning has been credited to a customer, spent, perhaps withdrawn as cash, and none of that has been settled between the banks until Monday.

The receiving bank has extended credit to the sending bank in the interval, on the initiative of a customer, without anybody's credit committee being consulted. This is not an oversight; it is the deliberate design of every deferred net system, and it is why those systems need risk controls that RTGS does not.

The practical rule for anyone building or reconciling systems: the timestamp on a customer credit and the timestamp on a settlement are different fields with different meanings, and the interval between them is a real exposure held by a real institution. Treating them as the same event is the single most common error in payments engineering.

■ **"Final" is a legal status, not a feeling**

The notebook analogy ends with cash in a hand, which feels final because it is physical. In electronic systems, finality is a legal construct, and it exists because insolvency law would otherwise be able to reach back in time and unwind transfers that everybody had treated as complete.

The CPMI definition is exact: final settlement is "the irrevocable and unconditional transfer of an asset or financial instrument, or the discharge of an obligation by the FMI or its participants in accordance with the terms of the underlying contract", and it adds that "final settlement is a legally defined moment".

In the United Kingdom the relevant instrument is the Financial Markets and Insolvency (Settlement Finality) Regulations 1999, which implemented the EU Settlement Finality Directive. Systems apply to be designated under them, and the Bank of England is the designating authority for everything except systems that are, or whose operator is, a recognised investment exchange, where it is the Financial Conduct Authority. Designation gives a system protection against the ordinary operation of insolvency law, so that transfer orders already entered into the system, and the netting of them, survive the insolvency of a participant instead of being reversed.

Two further corrections follow from this. First, finality attaches to the interbank transfer, not to the customer's experience. The Bank of England states that in CHAPS "the transfer of funds is irrevocable between the direct participants" — between the direct participants. That is not a statement that a defrauded consumer cannot be reimbursed. Since 7 October 2024 the Bank has operated CHAPS reimbursement rules for authorised push payment fraud, with a maximum reimbursement level for UK retail CHAPS payments set at £85,000, aligned with the Payment Systems Regulator's approach for Faster Payments. Settlement finality and consumer redress are different mechanisms operating on different parties.

Second, and less comfortably, finality is not identical across systems. A payment can be final in one legal sense and still be sitting inside a clearing cycle whose settlement has not occurred. If you need to know whether a specific payment is beyond recall, the answer is in the rulebook of the specific system, not in a general principle.

The technical version

■ **The vocabulary, defined**

Because the terms are used loosely in commerce and precisely in the rulebooks, it is worth setting out the CPMI definitions verbatim. These are the definitions that scheme rulebooks and central bank documentation are written against.

Term	CPMI definition
Clearing	The process of transmitting, reconciling and, in some cases, confirming transactions prior to settlement, potentially including the netting of transactions and the establishment of final positions for settlement
Settlement	The discharge of an obligation in accordance with the terms of the underlying contract
Final settlement	The irrevocable and unconditional transfer of an asset or financial instrument, or the discharge of an obligation by the FMI or its participants in accordance with the terms of the underlying contract. Final settlement is a legally defined moment
Settlement asset	An asset used for the discharge of obligations as specified by the rules, regulations or customary practice for an FMI
Settlement lag	In a transfer system, the time lag between the acceptance of the transfer order by the system and its final settlement
Real-time gross settlement	The real-time settlement of payments, transfer instructions or other obligations individually on a transaction-by-transaction basis
Deferred net settlement	A net settlement mechanism which settles on a net basis at the end of a predefined settlement cycle
Net settlement system	A funds or securities settlement system in which final settlement of transfer instructions occurs on a net basis at one or more discrete, pre-specified times during the processing day
Settlement risk	The general term used to designate the risk that settlement in a funds or securities transfer system will not take place as expected. This risk may comprise both credit and liquidity risk
Principal risk	The risk that a counterparty will lose the full value involved in a transaction — for example, the risk that a seller of a financial asset will irrevocably deliver the asset, but not receive payment
Payment versus payment	A settlement mechanism that ensures that the final transfer of a payment in one currency occurs if and only if the final transfer of a payment in another currency or currencies takes place

Three of these repay close attention. “Settlement lag” names the interval this chapter is about and makes it a measurable quantity rather than a vague unease. “Principal risk” is the specific, severe form of settlement risk in which the exposure is the whole trade rather than the cost of replacing it. And “payment versus payment” is the engineering answer to principal risk in foreign exchange, which we come to below.

■ The two models, and the trade-off they encode

The choice between real-time gross settlement and deferred net settlement is a straight exchange of one scarce resource for another. RTGS spends liquidity to buy the elimination of settlement risk. DNS spends settlement risk to buy liquidity efficiency.

The Bank of England states the trade-off in exactly those terms in its description of CHAPS: “Settlement risk is eliminated between CHAPS direct participants, at the cost of an increased need for liquidity, making this model best suited to a high-value payment system with the largest potential systemic risk.”

That last clause is the design principle. You put the expensive, risk-free model where the exposures would be catastrophic, and the cheap, exposed model where the exposures are small and manageable — and then you bolt risk controls onto the cheap model until its residual exposure is acceptable.

Britain’s arrangement follows this exactly. All of it settles across accounts at the Bank of England; what differs is when, and on what basis.

System	Operator	Settlement basis	Frequency
CHAPS	Bank of England	Real-time gross	Continuous through the settlement day
CREST	Euroclear UK and International	Delivery versus payment	High-frequency cycles through the day
Bacs	Pay.UK	Deferred multilateral net	Once each business day
Faster Payments Service	Pay.UK	Deferred multilateral net, prefunded	Three times each business day
Image Clearing System	Pay.UK	Deferred multilateral net	Once each business day
LINK	LINK Scheme	Deferred multilateral net	Daily cycles; weekend and holiday cycles settle the next business day
Visa Europe	Visa Europe Limited	Deferred multilateral net	Daily cycles; weekend and holiday cycles settle the next business day
Mastercard Europe	Mastercard Europe SA	Deferred multilateral net	Daily cycles; weekend and holiday cycles settle the next business day

■ CHAPS as the worked example of RTGS

CHAPS is a sterling same-day system for high-value wholesale payments and for time-critical lower-value payments such as property purchases. Responsibility for it transferred from the CHAPS Clearing Company to the Bank of England in November 2017, which means the Bank is simultaneously the system’s operator, its settlement agent and its supervisor’s colleague — an arrangement it manages through internal separation and an exchange of letters with the Payment Systems Regulator, most recently updated in January 2025.

The mechanics are as follows. Payment instructions are routed over SWIFT to the RTGS system. Settlement occurs across settlement accounts at the Bank of England by debiting the sending participant and crediting the receiving participant, individually, for full value. Only once that has happened is the full payment message forwarded onward to the receiving bank. This is the “Y-copy” arrangement, and its consequence is worth stating plainly: in CHAPS, interbank settlement is complete before the beneficiary’s bank has even received the complete instruction. The order of events is the reverse of the retail intuition.

The scale is instructive. In 2024 CHAPS settled over £87 trillion of payments, an average of over £344 billion each working day across around 208,000 payments, giving an average payment value of about £1.7 million. The Bank's own summary of what that means: CHAPS represents around 0.5 per cent of UK total payment volumes but 92 per cent of total sterling payment values.

More recent quarterly figures show the same picture and a rising trend.

Average daily figures	Q3 2025	Q4 2025	Q1 2026	Q2 2026
CHAPS volume	206,786	213,676	205,733	215,118
CHAPS value (£mn)	359,554	387,890	399,457	411,562
CREST DvP value (£mn)	425,351	460,405	485,260	477,841
Bacs net value (£mn)	5,370	5,653	5,800	6,205
FPS net value (£mn)	2,008	2,104	2,287	2,298
Cheque imaging net value (£mn)	41	38	37	37
LINK net value (£mn)	232	235	213	238
Visa net value (£mn)	2,165	2,273	2,173	2,356
Mastercard net value (£mn)	1,379	1,451	1,385	1,480
Total RTGS value (£mn)	796,099	860,050	896,612	902,018

There are over 35 CHAPS direct participants, with several thousand financial institutions reaching the system indirectly through them. The list is more varied than the phrase “high-street banks” suggests: alongside Barclays UK, Lloyds Bank plc, NatWest, HSBC UK Bank plc and Santander UK plc, the direct participants include CLS Bank International, LCH Limited, Euroclear Bank SA/NV, Fidelity UK Limited, ClearBank Limited, iFAST Global Bank Limited and Banking Circle S.A., plus a long list of London branches of foreign banks.

The liquidity cost of gross settlement is managed by two devices. The first is the CHAPS Throughput Rules, which oblige direct participants to have settled specified proportions of their daily CHAPS value by fixed points in the day — currently 12:00, 15:00 and 17:00. Without such a rule, every participant's individually rational strategy is to delay outgoing payments and fund them from incoming ones, which collectively produces the New York failure mode of 1974 in slow motion. The second is queueing and central scheduling inside RTGS, so that payments a participant cannot yet fund wait in an ordered queue rather than being rejected.

CHAPS currently opens at 06:00 and closes at 18:00 for bank-to-bank payments, with customer payments to be submitted by 17:40. Direct participants must be open to receive by 08:00 and to send by 10:00. In 2026 the Bank confirmed an early-morning extension under which CHAPS will open at 01:30, with all direct participants required to receive payments from that time whether or not they choose to send; the target is September 2027, and the Throughput Rules will continue to be measured from 06:00.

■ Bacs as the worked example of DNS

Bacs is the opposite design and demonstrates every property of deferred net settlement.

Payments submitted to Bacs go through a three-day clearing and processing cycle. The deadline for receipt of payment instructions from service users is 22:30 on Day 1. Processing and distribution to the receiving banks happens on Day 2. Debits and credits reach customer accounts on Day 3.

The interbank obligations arising in Bacs are settled at the Bank of England on a multilateral net basis on Day 3 of the cycle. Settlement occurs at 09:30, through the posting of multilateral net settlement

positions directly to settlement accounts in the RTGS processor. One posting, once a day, for the whole system.

The compression this achieves is the point of the design. In 2025 Bacs facilitated 6.86 billion payments with a value of £6.05 trillion, including a record 5.0 billion Direct Debits. Spread over roughly 250 working days, that is on the order of £24 billion of customer payments a day. The corresponding amount of central bank money that actually moved averaged £5,587 million a day across 2025 to date, and £6,205 million a day in the second quarter of 2026. Those two figures come from different publications measuring slightly different things, so the ratio should be treated as indicative rather than exact, but the order of magnitude is not in doubt: something like three-quarters of the value entering Bacs never needs to be settled at all, because it offsets against value going the other way.

Since its inception in 1968, over 183.6 billion transactions have been debited or credited to British bank accounts through Bacs.

The cost of that efficiency is the settlement lag. From 22:30 on Day 1 to 09:30 on Day 3 there exist agreed obligations between banks that have not been discharged. If a participant failed inside that window, the surviving participants would be holding claims rather than money.

■ The risk controls that make DNS survivable

The industry's answer to that window is not to shorten it but to collateralise it.

The Bank of England's RTGS supports both prefunded and non-prefunded deferred net settlement models. Under the prefunded model, each participant's maximum possible net debit position is capped, and the participant "must hold funds in a special prefunding account in RTGS equal to that cap". If a participant fails, the money required to complete settlement is already segregated at the central bank and settlement proceeds regardless. The exposure between participants is not reduced; it is removed, by being funded in advance.

Faster Payments operates on this basis, with each participant subject to a Net Sender Cap limiting its indebtedness to other participants per settlement cycle. The design has an obvious cost — prefunding is idle liquidity — and that cost is a barrier to entry, particularly for non-bank payment firms. In 2026 Pay.UK introduced a more flexible Net Sender Cap model intended to reduce the amount of prefunding required and widen direct access to Faster Payments.

The 2012 CPSS Red Book chapter on the United Kingdom recorded the three Faster Payments clearing cycles as settling at 07:15, 13:00 and 15:45. The Bank of England's current published description states only that net settlement takes place three times each business day, so the precise times should be taken from the scheme's current documentation rather than from the Red Book.

■ Herstatt, precisely

The canonical settlement failure deserves stating in exact terms, because it is invoked constantly and described loosely.

On 26 June 1974 the Bundesaufsichtsamt für das Kreditwesen, the German federal banking supervisory office, withdrew the banking licence of Bankhaus Herstatt of Cologne. The BIS account is specific about the timing: the withdrawal occurred during the banking day but after the close of the interbank payments system in Germany. At 10:30 New York time on the same day, Herstatt's New York correspondent suspended dollar payments on its account.

The consequence was that Herstatt's counterparty banks were exposed for the full value of the Deutschmark deliveries they had already made. They had performed their side of foreign exchange

contracts irrevocably and received nothing. The Basel Committee's own case study of the failure, published in April 2004 as Working Paper No. 13, Bank Failures in Mature Economies, records that by June 1974 Herstatt's losses on its foreign exchange operations amounted to DM 470 million, and that a special audit authorised by the supervisor in March of that year had already found open exchange positions of DM 2 billion against an internal limit of DM 25 million.

The 1996 BIS report on settlement risk in foreign exchange transactions — usually called the Allsopp Report — sets out the resulting definition: “One party to an FX trade could pay out the currency it sold but not receive the currency it bought. This principal risk in the settlement of foreign exchange transactions is variously called foreign exchange settlement risk or cross-currency settlement risk. It is also referred to as Herstatt risk.” The report itself notes that naming the risk after one institution is in some respects inappropriate, since the circumstances of subsequent failures differed.

The second-order effect is the one the Bank of England's own archives illuminate, in research published on the Bank Underground blog in June 2015. Banks in New York reacted by holding up outgoing payments until their accounts were sufficiently in credit. The Federal Reserve's Charles Coombs recorded that this “had the effect of virtually freezing up the clearing process, which on two or three consecutive days ran until 1am”. A payment system in which everybody waits for everybody else does not fail loudly; it congeals. The archives show a chain reaction of payment delays across international financial centres, threatening the stability of the international payment system as a whole.

The institutional response was durable. The G10 members, together with Luxembourg and Spain, established a standing committee under the auspices of the Bank for International Settlements. That lineage runs through the Basel Committee on Banking Supervision and the Committee on Payment and Settlement Systems, now the Committee on Payments and Market Infrastructures, whose glossary supplied the definitions above.

■ Payment versus payment, and how much risk is left

The engineering fix for principal risk in foreign exchange is payment versus payment: a settlement mechanism that ensures the final transfer of one currency occurs if and only if the final transfer of the other does. Neither leg can complete alone, so neither party can be left having paid without being paid.

The dominant implementation is CLS Settlement, operated by CLS Bank International, which provides PvP settlement for 18 of the world's most traded currencies. CLS reports settling an average daily value of over USD 7 trillion of payment instructions for more than 70 settlement members and over 37,000 indirect participants; its product materials cite a figure of over USD 8.0 trillion of payments each day. CLS Bank International is a CHAPS direct participant and Continuous Linked Settlement is designated under the UK Settlement Finality Regulations, which is a compact illustration of how these arrangements interlock: sterling PvP legs settle in central bank money at the Bank of England.

The 2025 BIS Triennial Survey put numbers on how much of the problem remains. Of roughly USD 14 trillion of gross financial obligations settled daily in April 2025, USD 5.2 trillion — 36 per cent — settled via PvP, eliminating settlement risk. About USD 7.6 trillion, 54 per cent, settled by methods that mitigate but do not eliminate the risk, such as pre-settlement netting, intragroup settlement and timing controls. And more than USD 1.4 trillion, 10 per cent, settled on a gross bilateral basis, fully exposed. More than four fifths of daily settlement — USD 12.2 trillion — involved CLS-eligible currency pairs, of which 40 per cent settled via PvP.

Fifty-one years after Herstatt, on an ordinary day in April 2025, well over a trillion dollars a day of foreign exchange still settled in exactly the manner that destroyed Herstatt's counterparties.

■ Settlement finality in UK law

The Financial Markets and Insolvency (Settlement Finality) Regulations 1999 allow payment and settlement systems to apply for protection against the normal operation of insolvency law. The designating authority is the Bank of England, except where the system is, or its operator is, a recognised investment exchange for the purposes of the Financial Services and Markets Act 2000, in which case it is the FCA.

Systems designated under UK law include CHAPS, Bacs, the Faster Payments Service, the Image Clearing System, Continuous Linked Settlement, Visa Europe, Euroclear UK and International, LCH Limited, ICE Clear Europe, LME Clear Limited, SIX x-clear and the Sterling Finality Payment System. A further set of non-UK-law systems is designated for these purposes, including EBA Clearing's STEP2-T, EURO1 and RT1, Euroclear Bank SA/NV, Eurex Clearing AG, LCH SA, Clearstream Banking S.A., the Norges Bank Settlement System and RIX operated by Sveriges Riksbank.

Designation is not a badge. It is the legal machinery that makes settlement stick when a participant fails, and being a participant within the meaning of those Regulations is one of the stated conditions for direct participation in CHAPS.

■ The four events, and how to talk about them

The chapter's opening claim can now be made precisely. "The payment went through" collapses at least four distinct events.

Event	What has happened	Whose record	Can it be undone?
Authorisation or acceptance	The sending institution has validated the instruction and typically debited its customer	Sending bank's own ledger	Yes, by the sending bank, subject to scheme rules
Clearing	The obligation has been established and, where applicable, netted into a position	Scheme infrastructure and both banks' records	The position can change until the cycle closes
Settlement	Central bank money has moved between settlement accounts at the Bank of England	The Bank of England's RTGS ledger	Not between the direct participants once final
Customer credit	The receiving bank has increased the balance it owes its customer	Receiving bank's own ledger	Yes, in some circumstances, by the receiving bank

The order of those four is not fixed. In CHAPS, settlement precedes the receiving bank getting the full message, which precedes the customer credit. In Faster Payments, the customer credit comes first, often within seconds, and settlement follows in the next of three daily cycles. In Bacs, the customer entries and the settlement both fall on Day 3, near each other but not identical, with settlement at 09:30. In the card systems, authorisation may precede clearing by days and settlement by days more — a sequence Volume III takes apart in detail.

This is why the diagnostic question in an incident is never "did the payment go through". It is a sequence of narrower questions. Did the instruction get accepted, and by whom? Did it enter a clearing cycle, and which one? Has that cycle settled? And has the beneficiary's bank posted to the customer

account? Four questions, four different systems of record, and frequently four different organisations to ask.

■ What this means for people building on these rails

Two practical consequences follow, and both cost money when ignored.

The first concerns reconciliation, the subject of a later chapter. A clearing file and a settlement advice are not two views of the same fact. The clearing file lists individual items; the settlement advice reports a position. They will not tie out item by item, and they are not supposed to. Reconciliation is the discipline of proving that a set of items produces the position that was settled, which is a different exercise from matching one list against another and considerably more subtle.

The second concerns the meaning of an incoming payment for risk purposes. If your business releases goods, credits an account, or allows a withdrawal on the strength of a payment that has been cleared but not settled, you have taken on a share of a bank's settlement exposure without a bank's balance sheet. For most retail-scale activity that is a reasonable commercial judgement, because prefunding and net sender caps stand behind the retail systems. For unusual amounts, unusual timing, or unusual counterparties, it is worth knowing exactly which of the four events has occurred before acting.

And that, ultimately, is why this distinction is worth a chapter. The public has one word for payment. The industry has a pipeline with several stages, several ledgers and several legally defined moments, and it built that pipeline the way it did because in June 1974 it discovered, expensively, what happens when the gap between agreeing and paying is left unmanaged.

7.98 Common wrong ideas

Wrong: Clearing and settlement are two words for the same thing. **Right:** Clearing establishes who owes what and settlement discharges the obligation, and the CPMI glossary itself warns that the industry uses "clearing" imprecisely to cover both.

Wrong: A payment that has cleared has been paid for. **Right:** Clearing produces a number and only settlement makes the number true, and in Bacs those events are two days apart.

Wrong: Your balance going up means the banks have settled. **Right:** Your balance is your own bank's record of what it owes you, updated without any cooperation from another institution or from the Bank of England.

Wrong: Clearing exists in order to net, so a system that settles gross does not clear. **Right:** CHAPS nets nothing yet validates, sequences, checks liquidity and queues, because clearing is the whole pipeline and netting is only one optional step inside it.

Wrong: Clearing is a discrete event at the end of the day. **Right:** It is a continuous pipeline of transmission, validation, sanctions and fraud screening, matching, queueing and position-building that runs all day.

Wrong: Herstatt risk is a currency risk. **Right:** It is principal risk: paying out the full amount of one currency and receiving nothing of the other, which is why payment versus payment was invented.

Wrong: CLS has solved foreign exchange settlement risk. **Right:** In April 2025 only 36 per cent of roughly USD 14 trillion of daily obligations settled by PvP, and more than USD 1.4 trillion settled on a gross bilateral basis, fully exposed.

Wrong: Settlement finality means a defrauded consumer cannot be reimbursed. **Right:** Finality is irrevocability between the direct participants, and since 7 October 2024 CHAPS has had its own reimbursement rules for authorised push payment fraud, capped at £85,000 for UK retail payments.

Wrong: Finality works the same way in every system. **Right:** A payment can be final in one legal sense and still sit inside a clearing cycle that has not settled, so the answer is in the specific rulebook rather than in a general principle.

Wrong: In CHAPS the beneficiary's bank receives the message and then the banks settle. **Right:** Under the Y-copy arrangement interbank settlement completes first, and only then is the full payment message forwarded onward.

7.99 Chapter summary in 20 lines

1. The public has one word, "payment", for what the industry treats as at least four events on four ledgers at four different times.
2. Clearing is the process of working out who owes what to whom; settlement is the money actually moving.
3. Clearing produces a number and settlement makes the number true.
4. Clearing happens in the world of obligations and settlement in the world of central bank money at the Bank of England.
5. The CPMI defines clearing as transmitting, reconciling and sometimes confirming transactions prior to settlement, potentially including netting, and warns that the term is often used imprecisely to cover settlement.
6. Clearing is a continuous pipeline of validation, screening, matching, queueing and sequencing rather than a meeting at six o'clock.
7. Netting is one optional step inside clearing rather than its purpose, and CHAPS clears substantially while netting nothing at all.
8. There are two settlement models: real-time gross settlement, which spends liquidity to eliminate settlement risk, and deferred net settlement, which spends settlement risk to buy liquidity efficiency.
9. Britain puts the expensive model where the exposures would be catastrophic and the cheap model where they are small, then bolts risk controls onto the cheap one until the residue is acceptable.
10. CHAPS settles each payment individually across accounts at the Bank of England for full value, and forwards the complete message only afterwards, which is the Y-copy arrangement.
11. CHAPS is around 0.5 per cent of UK payment volumes and about 92 per cent of sterling payment values, with Throughput Rules at 12:00, 15:00 and 17:00 to stop every participant waiting for everybody else.
12. Bacs is the opposite design: submission by 22:30 on Day 1, processing on Day 2, and customer entries with settlement at 09:30 on Day 3.
13. Bacs compresses something like £24 billion of daily customer value into roughly £6 billion of central bank money moving once a day.
14. The price of that efficiency is the settlement lag, a window in which agreed obligations exist between banks that have not been discharged.
15. The industry's answer is not to shorten that window but to collateralise it, through prefunding in a segregated RTGS account and through net sender caps.
16. Your balance changing is not settlement, and treating the customer credit timestamp and the settlement timestamp as one event is the commonest error in payments engineering.

17. On 26 June 1974 Bankhaus Herstatt lost its licence after the German interbank payment system had closed, and its counterparties lost the full principal of the Deutschmarks they had already delivered.
18. Banks in New York then held back their own payments until they could see incoming ones arrive, freezing the clearing process until one in the morning for several nights and prompting the committee structure that governs payment oversight to this day.
19. Finality is a legally defined moment, conferred in the United Kingdom by designation under the Financial Markets and Insolvency (Settlement Finality) Regulations 1999, and it is distinct from consumer redress.
20. The diagnostic question in an incident is therefore never “did the payment go through” but which of authorisation, clearing, settlement and customer credit has occurred, and in whose records it can be seen.

Sources used: Bank for International Settlements — CPMI, A glossary of terms used in payments and settlement systems; Settlement risk in foreign exchange transactions (Allsopp Report, 1996); Payment, clearing and settlement systems in the United Kingdom (CPSS Red Book country chapter); Uncovering FX settlement risk: new measures from the 2025 BIS Triennial Survey, BIS Quarterly Review. Bank of England — CHAPS; A brief introduction to the Real-Time Gross Settlement system and CHAPS; Payment and settlement; Payment and settlement statistics; Extending RTGS and CHAPS settlement hours: early morning extension (2026); Financial market infrastructure supervision — Who are we (designated systems lists); Bank Underground, “BoE archives reveal little known lesson from the 1974 failure of Herstatt Bank” (June 2015). Pay.UK — Bacs Payment System; Bacs Payment System statistics; Understanding changes to the Faster Payment System Net Sender Caps. CLS Group — Settlement; Reimagining same-day FX (ShapingFX series). legislation.gov.uk — The Financial Markets and Insolvency (Settlement Finality) Regulations 1999. Payment Systems Regulator — PS24/5, CHAPS APP scams reimbursement requirement.

Netting

8.0 What this chapter gives you

1. You will be able to explain how about £23 billion of daily Bacs payments is discharged by moving about £5.3 billion between settlement accounts, with nobody short-changed.
2. You will be able to compute bilateral and multilateral net positions from a matrix of flows, and say why multilateral almost always wins.
3. You will be able to explain why the sum of all net debit positions must equal the sum of all net credit positions.
4. You will be able to say why Faster Payments nets harder than Bacs, and connect the answer to the direction of the traffic rather than to the technology.
5. You will be able to explain why netting does not reduce risk but concentrates it, turning diversified credit risk into liquidity risk at a single moment.
6. You will be able to work through an unwind and show why a bank with no meaningful exposure to the failed participant can suddenly owe five times more.
7. You will be able to distinguish position netting, netting by novation and close-out netting, and say which of them survives an administrator's cherry-picking.
8. You will be able to state the Lamfalussy "cover one" standard and name the controls that implement it: caps, prefunding and loss sharing.
9. You will be able to describe what a central counterparty does through novation, and the order in which its default waterfall is consumed.
10. You will be able to give the real reasons Bacs takes three days, and explain why the debit and the credit both fall on day three.

On an average working day in 2024, the Bacs system carried about 26.8 million payments worth roughly £23.0 billion between British banks. At the end of that day, the amount of money that actually moved between those banks in the Bank of England's settlement accounts averaged £5,265 million.

Nobody was short-changed. Every salary was paid in full, every direct debit was collected in full, every recipient saw the exact amount they were owed. Twenty-three billion pounds of value was transferred and about five billion pounds of money moved.

That is netting. It is arithmetic so simple that a child can do it, and its consequences run through every payment system on earth. It is the reason interbank settlement is affordable at all. It is also the reason a payment system can fail in a way that a pile of individual payments cannot, and the reason central banks spent the 1980s and 1990s writing rules about it.

The previous chapter separated clearing from settlement. This chapter is about what happens in the gap: the operation that turns a very large number of small obligations into a very small number of large ones.

The plain version

Four friends share a house: Ama, Ben, Cara and Dev. Over the course of a week they keep buying things for each other, because that is what people who live together do. Somebody gets to the front of the queue first and pays for two coffees. Somebody covers a train fare. Somebody buys the bin bags.

By Sunday night the situation looks like this.

Who owes	To whom	Amount
Ama	Ben	£12
Ben	Ama	£9
Ama	Cara	£5
Cara	Ama	£7
Ama	Dev	£8
Dev	Ama	£2
Ben	Cara	£15
Cara	Ben	£4
Ben	Dev	£6
Dev	Ben	£11
Cara	Dev	£20
Dev	Cara	£3

Twelve separate debts. If everybody paid everybody exactly what they owed, twelve payments would be made and £102 in cash would change hands.

Nobody does that. It would take all evening and everybody would need to have the right cash on them.

■ The first improvement: settle up in pairs

The obvious short cut is to deal with one pair at a time. Ama owes Ben £12 and Ben owes Ama £9, so instead of two payments totalling £21, Ama hands Ben £3 and both debts are gone.

Do this for every pair and you get six payments instead of twelve.

Pair	Owed one way	Owed the other	Who pays	Amount
Ama and Ben	£12	£9	Ama pays Ben	£3
Ama and Cara	£5	£7	Cara pays Ama	£2
Ama and Dev	£8	£2	Ama pays Dev	£6
Ben and Cara	£15	£4	Ben pays Cara	£11
Ben and Dev	£6	£11	Dev pays Ben	£5
Cara and Dev	£20	£3	Cara pays Dev	£17

Twelve payments have become six, and £102 of cash movement has become £44. Every debt has been discharged in full. Nobody has been given a discount and nobody has taken a haircut. The debts cancelled each other out because they pointed in opposite directions.

This is called bilateral netting. Bilateral means two-sided: you look at one pair of people at a time.

■ The second improvement: one pot

There is a better way, and it takes about a minute with a pencil.

Instead of pairing people up, work out for each person the total they owe everybody and the total everybody owes them. Then take the difference.

Person	Owes in total	Is owed in total	Difference
Ama	£25	£18	Owes £7
Ben	£30	£27	Owes £3
Cara	£31	£23	Owes £8
Dev	£16	£34	Is owed £18

Ama puts £7 on the kitchen table. Ben puts £3 on the table. Cara puts £8 on the table. There is now £18 on the table, and Dev picks it up.

Four movements of cash instead of twelve. Eighteen pounds instead of one hundred and two. And, once again, every single one of those twelve original debts has been paid in full. Nobody is out of pocket by a penny more or less than they should be.

Notice that the amounts owed and the amounts owing balance exactly. Ama, Ben and Cara owe £7, £3 and £8, which comes to £18, and Dev is owed £18. That is not a coincidence and it is not luck. Every debt has two ends. If you add up everything everybody owes and subtract everything everybody is owed, the answer must be zero, because you have counted each debt once as a plus and once as a minus. So the money going into the pot always equals the money coming out of it.

This is called multilateral netting. Multilateral means many-sided: everybody at once.

■ Now do it with banks

Replace the four friends with four banks, and replace the coffees and train fares with customer payments.

When a Barclays customer pays a Lloyds customer £180, that is not really a payment from a person to a person. Underneath, it is Barclays owing Lloyds £180. The customer's balance goes down at Barclays, the recipient's balance goes up at Lloyds, and now Barclays is £180 short with Lloyds and has to make it good.

There are about 27 million of those a day on Bacs alone, and they run in every direction at once, because Barclays customers pay Lloyds customers all day and Lloyds customers pay Barclays customers all day. The great majority of the value cancels out, exactly the way Ama's £12 and Ben's £9 cancelled out.

So the banks do precisely what the flatmates did. They add everything up, work out one number for each bank, and move that. Twenty-seven million payments become a few dozen movements in the Bank of England's books.

■ Why this is such a spectacular deal

Money sitting in a bank's settlement account at the Bank of England is money doing nothing else. It is not lending and it is not investing; it is sitting there so that payments can be made out of it. A bank that had to fund every outgoing payment as it happened would need a very large pile of cash standing still all day.

Netting means it needs a very much smaller pile. In the flatmates example the cash needed fell from £102 to £18. If you had asked all four to have the full £102 available on Sunday evening, somebody

would have said no. At £18 nobody even thinks about it. Multiply that by a national payment system and netting stops being convenient and becomes the difference between a payment system a country can afford to run and one it cannot.

■ And why it is dangerous

Here is the catch, and it is worth understanding before we go any further.

Look at Dev. He is owed £18 and he is expecting to be handed £18 on Sunday night. But he is not being handed £18 by any one person. He is being handed £18 by an arrangement. If Cara turns out to have nothing in her wallet at the crucial moment, the pot has £10 in it, not £18, and Dev is short.

Worse: if Cara's debts have to be taken out of the calculation entirely, everybody's number changes. It is not just that Cara's £8 goes missing. Ama, who owed £7, might now owe more, or less. Ben, who owed £3, might turn out to be owed money. Nobody can predict which way their position moves until the sum is redone.

That is the trade the whole payments industry has made. Netting makes settlement cheap by making everybody depend on one calculation. When the calculation works, it is enormously efficient. When it does not, the failure lands on everybody at once rather than on the person who was unlucky enough to be owed money by the party that failed.

■ And why waiting is useful

One last thing before we take the analogy apart.

The longer you wait before squaring up, the more cancels out. If the flatmates settled every debt the moment it arose, nothing would ever net, because there would never be two debts alive at the same time to cancel against each other. If they wait a day, some of it cancels. If they wait a week, most of it does.

This is not a small effect. It is the reason a payment system that waits three days can move a fifth of the money that a payment system settling instantly would have to move.

That is the thread that runs to the end of this chapter. When you eventually ask why a Bacs payment takes three days, part of the answer is: because three days is how long you have to wait for enough of it to cancel out.

Where the plain version stops being true

The kitchen table is the right picture, and it is very close to what actually happens. Four things about it are wrong in ways that matter.

■ Nobody is paid less, and the netting is invisible to customers

This is the misunderstanding that most often survives a first explanation, and it is worth killing immediately.

Netting operates on the obligations between banks. It does not touch the payments between customers. If your employer pays you £2,400 and your landlord takes £1,100 by direct debit on the same morning, you are credited £2,400 and debited £1,100. You are not credited £1,300. Your bank's ledger records both movements in full, because your bank's ledger is a record of what your bank owes you, and it owes you both things separately.

What is netted is the layer underneath: the amount your bank owes the other banks. The customer layer is gross. The interbank layer is net. They are different ledgers with different arithmetic, and the

whole efficiency of the system comes from the fact that one can be compressed without touching the other.

The same is true within a bank. If a Lloyds customer pays another Lloyds customer, that payment never reaches an interbank system at all. It is settled by moving a number from one row of Lloyds' own ledger to another. The Bank of England's published payment figures explicitly exclude payments internalised within a single institution, for exactly this reason.

■ Netting is a legal act, not just arithmetic

Adding up is easy. Making the answer stick when somebody goes bankrupt is not, and this is where most of the industry's difficulty actually lies.

Suppose Cara is declared insolvent on Sunday afternoon. Her administrator's job is to maximise the return to her creditors. A sufficiently aggressive administrator will look at the twelve debts and say: the £20 that Cara owes Dev is a claim in the insolvency, to be paid at whatever pence in the pound the estate can manage; the £3 that Dev owes Cara is an asset of the estate and must be paid in full, today. That is called cherry-picking, and it converts a tidy £17 net obligation into a £20 unsecured claim and a £3 immediate demand.

If that were allowed, netting would provide no protection whatever, and every bank in the country would have to hold capital and liquidity against its gross exposures rather than its net ones. So it is not allowed, but only because specific law says so. In the United Kingdom the relevant instrument is the Financial Markets and Insolvency (Settlement Finality) Regulations 1999, SI 1999/2979, which implemented Directive 98/26/EC. Systems designated under those regulations get statutory protection: transfer orders entered into a designated system, and the netting of those orders, take precedence over ordinary insolvency law and cannot be unwound by an administrator.

Designation is not automatic and not universal. There is a real difference between netting that is legally robust and netting that is merely a spreadsheet, and practitioners care about which one they are looking at.

■ There is usually no pot

The kitchen table suggests a fund sitting in the middle holding money. In most net settlement systems there is no such thing.

The system operator calculates the net positions and sends the Bank of England a set of instructions. The Bank debits the accounts of the banks in net debit and credits the accounts of the banks in net credit, in a single simultaneous posting. Nothing is held in the middle at any point. The "pot" is a bookkeeping convenience, not a place.

There is an important exception, and it is exactly where the risk is. Some systems do require participants to put money aside in advance, in a segregated account at the Bank of England, so that there is something to fall back on if a participant cannot pay. That is called prefunding, and the Bank of England introduced it for both Bacs and Faster Payments in September 2015. Where prefunding applies there genuinely is a pot, and it exists because the arithmetic on its own was judged not safe enough.

■ Netting does not reduce risk; it concentrates it

The flatmates example makes netting look like pure gain. It is not.

Before netting, Dev has three separate exposures: to Ama, to Ben and to Cara. If Cara fails, Dev loses what Cara owed him and keeps what the other two owed him. The damage is proportionate and contained.

After netting, Dev has one exposure: to the arrangement. If any participant fails, the amount Dev receives changes, and by an amount that has nothing to do with his direct dealings with the party that failed. He has traded a set of small, understandable, diversified exposures for a single large exposure to a mechanism.

That is not a criticism of netting. It is the deal. Netting converts credit risk that was spread across many counterparties into liquidity risk concentrated in one moment, and the entire architecture of modern payment system regulation exists to manage that one moment. The technical version works through exactly what happens when it goes wrong.

The technical version

■ Vocabulary

The Committee on Payments and Market Infrastructures, the standard-setting body at the Bank for International Settlements, defines clearing as “the process of transmitting, reconciling and, in some cases, confirming transactions prior to settlement, potentially including the netting of transactions and the establishment of final positions for settlement”. Netting is therefore a sub-process of clearing, not a synonym for it, and the qualifiers matter: “in some cases”, “potentially including”. A system can clear without netting.

The working definitions used throughout the rest of this book are as follows.

Gross settlement means each payment obligation is settled individually, for its full amount. **Net settlement** means obligations are offset against each other and only the resulting balance is settled.

Bilateral netting offsets obligations between one pair of participants. Each pair produces one net figure and one settlement movement.

Multilateral netting offsets each participant’s obligations against all other participants simultaneously. Each participant produces one net figure: a **net debit position** if it owes the system, a **net credit position** if the system owes it. The sum of all net debit positions equals the sum of all net credit positions, always, by construction.

Deferred net settlement (DNS) means the netted figure is settled at a scheduled time after the underlying payments have been exchanged. **Real-time gross settlement** (RTGS) means each payment is settled individually as it arises. The Bank of England currently supports four settlement models in its RTGS infrastructure: real-time gross settlement, delivery versus payment, prefunded deferred net settlement and unfunded deferred net settlement. Only CHAPS uses the first. The compression achieved is expressed as a **netting ratio**: the netted value as a proportion of the gross value, or its complement, the percentage of value eliminated.

■ Four banks, worked in full

Take four settlement participants and one clearing cycle. The matrix below shows the total value of payment instructions sent by each bank to each other bank during the cycle, in millions of pounds. Rows are payers, columns are payees.

Sending bank	To Northbank	To Eastbank	To Southbank	To Westbank	Total sent
Northbank	—	420	180	260	860
Eastbank	380	—	310	150	840
Southbank	210	290	—	340	840
Westbank	160	120	270	—	550

Sending bank	To Northbank	To Eastbank	To Southbank	To Westbank	Total sent
Total received	750	830	760	750	3,090

Twelve directional flows. Aggregate gross value: **£3,090 million**.

If this system settled gross, twelve settlement movements totalling £3,090 million would be required, and each bank would need enough liquidity in its settlement account to cover its full outgoing total: £860 million for Northbank, £840 million each for Eastbank and Southbank, £550 million for Westbank. Aggregate liquidity requirement: £3,090 million.

Bilateral netting

Offset each pair.

Pair	Owed one way	Owed the other	Net payer	Net amount
Northbank / Eastbank	420	380	Northbank	40
Northbank / Southbank	180	210	Southbank	30
Northbank / Westbank	260	160	Northbank	100
Eastbank / Southbank	310	290	Eastbank	20
Eastbank / Westbank	150	120	Eastbank	30
Southbank / Westbank	340	270	Southbank	70

Six settlement movements totalling **£290 million**. That is 9.4 per cent of the gross value: bilateral netting has eliminated 90.6 per cent of the value that would otherwise have had to move. Note the liquidity consequence. Northbank must fund £40 million to Eastbank and £100 million to Westbank, £140 million in all, while separately receiving £30 million from Southbank, and it cannot use the incoming £30 million to fund the outgoing £140 million, because bilateral netting produces six independent settlements rather than one.

Multilateral netting

Now offset everything against everything.

Bank	Total sent	Total received	Net position
Northbank	860	750	Net debit 110
Eastbank	840	830	Net debit 10
Southbank	840	760	Net debit 80
Westbank	550	750	Net credit 200

Northbank, Eastbank and Southbank are debited £110 million, £10 million and £80 million respectively. Westbank is credited £200 million. Debits sum to £200 million; credits sum to £200 million.

Four postings to settlement accounts, moving **£200 million**. That is 6.5 per cent of the gross value: multilateral netting has eliminated 93.5 per cent.

The three side by side

Settlement basis	Movements	Value settled	Percentage of gross	Largest single obligation
Gross	12	£3,090m	100.0%	£420m
Bilateral net	6	£290m	9.4%	£100m
Multilateral net	4	£200m	6.5%	£110m

Multilateral netting beats bilateral netting here, as it almost always does, because it allows a claim on one participant to fund an obligation to a different participant. Northbank's £30 million claim on Southbank reduces its aggregate obligation from £140 million to £110 million, which bilateral netting cannot achieve. The efficiency grows with the number of participants and with how balanced the flows are, and is weakest when one participant is a persistent net sender to everybody.

■ What netting buys, in real numbers

The four-bank example is illustrative. The following are not.

Bacs and Faster Payments

Pay.UK, which operates Bacs, the Faster Payment System and the Image Clearing System, reported that in 2024 Bacs processed 6,811 million payments worth £5,838,567 million, an average of 26.8 million payments and £22,986 million per working day. In 2025 the figures rose to 6,864 million payments worth £6,049,248 million, an average of 27.1 million payments and £23,910 million a day.

The Bank of England, which settles those obligations, reported that daily average Bacs net settlement values in RTGS in 2024 were £5,265 million.

So roughly £23 billion of customer payment value per working day was discharged by moving roughly £5.3 billion between settlement accounts. Netting eliminated of the order of three quarters of the value. The two figures come from different publishers and may be computed over slightly different numbers of days, so the ratio should be read as approximately 77 per cent eliminated rather than to three significant figures. The order of magnitude is not in doubt.

The compression in the number of movements is more dramatic still. At the end of 2025 Bacs had 33 direct customers, over 330 indirect customers and around 117,000 service users submitting payment files. Twenty-seven million payment instructions a day, originating from a hundred and seventeen thousand organisations, resolve to at most thirty-three postings in the Bank of England's ledger, once per business day.

Faster Payments compresses harder. Pay.UK reported total Faster Payments value of £4,242,380 million in 2024, while the Bank of England reported daily average Faster Payments net settlement values of £1,929 million. Faster Payments settles three times each business day, so the annual settled value is of the order of £490 billion against gross value of £4,242 billion. Netting is eliminating something in the region of 88 per cent.

The reason Faster Payments nets harder than Bacs is instructive. Bacs traffic is dominated by direct credits — payroll and benefits — which flow overwhelmingly in one direction on particular days of the month. Faster Payments traffic is person-to-person and person-to-business, running in both directions continuously. Balanced flows net well. Directional flows do not.

For comparison, the Bank of England's 2024 daily averages across the whole of RTGS were:

Settlement service	Daily average value 2024
CHAPS (gross)	£344,409m
CREST delivery versus payment	£391,715m
Bacs (net)	£5,265m
Visa Europe (net)	£2,193m
Faster Payments (net)	£1,929m
Mastercard (net)	£1,406m
LINK (net)	£241m
Image Clearing System (net)	£44m
Total	£747,202m

CHAPS settles gross and therefore appears at its full value. Every other retail system in the table has already been compressed before it reaches the Bank of England. That single table is the clearest statement in print of what netting does to a national payment system: seven retail systems carrying the overwhelming majority of the country's payment volume account for about 1.5 per cent of the value settled.

Securities: the DTCC

In the United States, the National Securities Clearing Corporation, part of the Depository Trust and Clearing Corporation, clears the great majority of US equity trades. DTCC states that NSCC netting reduces the value of payments that need to be exchanged by an average of 98 to 99 per cent every day.

That is a higher ratio than any payment system achieves, and the reason is structural: securities trading involves the same instruments changing hands many times in a day between the same set of intermediaries, so the offsetting is extraordinarily dense. NSCC's continuous net settlement service reduces each member's obligations in each security to a single net long or short position. For scale, DTCC reported that on 9 April 2025 NSCC reached a peak value of \$5.55 trillion, against a previous peak of \$5.22 trillion on 20 December 2024.

Foreign exchange: CLS

CLS Bank settles foreign exchange on a payment-versus-payment basis in eighteen currencies, holding accounts with all eighteen central banks. It settles the gross value of each instruction, but funds the settlement on a multilaterally netted basis: each settlement member transfers only the net amount of its combined obligations in each currency.

CLS reports that this, combined with its in/out swap liquidity tool, produces an average funding requirement of less than 1 per cent of the total value of all trades for participating settlement members. In the first half of 2022 CLS settled an average of over \$6.5 trillion every day, with over 70 settlement members and over 30,000 third-party participants.

Less than one per cent. That is the number to remember when somebody claims that netting is a marginal optimisation.

The 2025 BIS Triennial Central Bank Survey put a broader figure on the same phenomenon. More than \$14 trillion of gross foreign exchange obligations were settled on an average day in April 2025. More than \$2 trillion of that, about 15 per cent, was subject to pre-settlement netting, and that \$2.2 trillion was reduced to \$337 billion of actual settlement — a reduction of roughly 85 per cent on the netted portion alone.

■ The legal forms of netting

The distinctions below are not academic. They determine whether a netting arrangement survives an insolvency, and therefore whether a bank may recognise net rather than gross exposures for regulatory capital purposes.

Position netting, sometimes called payment netting or settlement netting, computes a net figure for settlement purposes but does not legally discharge the underlying obligations. The gross obligations continue to exist until they are settled. If a participant fails before settlement, the net calculation has no independent legal force and the gross obligations are what an administrator sees. This is the form that gives rise to unwind risk.

Netting by novation replaces each new obligation between a pair of parties with a single new obligation that discharges and supersedes the previous one. At any moment there is exactly one obligation between the pair, of the net amount. Because the gross obligations have been legally extinguished as they arose, there is nothing for an administrator to cherry-pick.

Close-out netting operates on default rather than on schedule. On a defined termination event, all outstanding contracts between the parties are terminated, valued at market, and reduced to a single net sum payable one way. This is the mechanism in section 6 of the ISDA Master Agreement and it is why the Master Agreement insists that all transactions under it form a single agreement.

Multilateral netting cannot be achieved by novation between pairs alone, because there is no bilateral contract between Northbank and the system. It requires either a central counterparty that becomes the legal counterparty to everybody, or statutory protection of the sort the Settlement Finality Regulations 1999 confer on designated systems, or both.

■ Where the risk goes: the unwind, worked through

Return to the four banks and suppose Northbank fails during the cycle, before settlement, and cannot meet its net debit of £110 million.

In a system with position netting and no prefunding, the classic remedy is to remove the failed participant's payments and recalculate. This is called an unwind, and it is worth seeing what it does.

Delete every flow to and from Northbank. The remaining matrix is:

Sending bank	To Eastbank	To Southbank	To Westbank	Total sent
Eastbank	—	310	150	460
Southbank	290	—	340	630
Westbank	120	270	—	390
Total received	410	580	490	1,480

Recalculate the net positions and compare.

Bank	Original net position	Position after unwind	Change
Eastbank	Debit 10	Debit 50	Must find £40m more
Southbank	Debit 80	Debit 50	Owes £30m less
Westbank	Credit 200	Credit 100	Receives £100m less

Look carefully at what has happened.

Eastbank had nothing to do with Northbank's failure in any meaningful sense — its bilateral position with Northbank was a net £40 million payable. Yet its settlement obligation has increased fivefold, from £10 million to £50 million, and it must find the difference at the moment of settlement, with no notice. Southbank's obligation has fallen. Westbank, which had planned its day around receiving £200 million, receives £100 million.

Two properties of that table are the reason central banks stopped tolerating unwind as a risk control. First, the direction and size of the change bear no reliable relationship to a participant's dealings with the failed party; Eastbank cannot manage its exposure to an unwind by managing its exposure to Northbank, because the effect propagates through everybody else's positions too. Second, the change is discovered at exactly the worst moment: a bank learns that it needs £40 million more, in central bank money, at the instant a major counterparty has failed and every other bank in the country is scrambling for liquidity. That is the mechanism by which a single failure becomes a system-wide liquidity event.

■ The rules written in response

In November 1990 the central banks of the Group of Ten published the Report of the Committee on Interbank Netting Schemes, universally known as the Lamfalussy Report after its chairman, Alexandre Lamfalussy. It set six minimum standards for netting schemes, which remain the foundation of every rule written since.

In summary, the six standards require that netting systems have a well-founded legal basis in all relevant jurisdictions; that participants clearly understand the effect of the scheme on each financial risk; that multilateral netting systems have clearly defined procedures for managing credit and liquidity risk, specifying who is responsible for what and placing limits on the maximum credit exposure any participant can create; that multilateral netting systems be capable, at a minimum, of ensuring timely completion of daily settlement in the event of an inability to settle by the participant with the largest single net debit position; that admission criteria be objective, publicly disclosed and permit fair and open access; and that the operational reliability of technical systems and the availability of backup facilities be assured.

The fourth is the one that changed the industry. It is the requirement now universally known as “cover one”: the system must be able to complete settlement on time even if its single largest debtor cannot pay. Everything below follows from it.

Net debit caps and net sender caps. A cap limits how large a net debit position a participant may build up between settlements. When a participant reaches its cap it can no longer send payments until incoming payments reduce its position or settlement occurs. In the Faster Payment System these are called Net Sender Caps, and Pay.UK announced in July 2026 a move away from a prescriptive formula, based on a Minimum Net Sender Cap and a Peak Contingency Value, towards participants setting their own caps against quarterly plans reviewed by Pay.UK. The purpose is unchanged: to bound the size of the hole any single participant can leave.

Prefunding. The Bank of England introduced prefunding for Bacs and Faster Payments in September 2015 to address the settlement risk arising from the build-up of obligations in deferred net systems. Under prefunding, each settlement participant holds cash in a segregated prefunding account in RTGS equal to its cap. If it defaults, the cash set aside is used to complete settlement. As the Bank puts it, this eliminates credit risk between settlement participants: the money is already there, in central bank money, before the obligation arises.

Prefunding is the reason the modern British answer to “what if a bank cannot settle its Bacs position” is no longer “unwind the cycle”. It is “use the money it already put aside”.

Loss-sharing arrangements. Where prefunding is not used, systems typically operate survivors-pay arrangements: a defaulter's shortfall is allocated among the remaining participants according to a pre-agreed formula, usually related to their bilateral exposures to the failed party. These arrangements produce certainty of settlement completion at the cost of mutualised loss, and their existence is one of the reasons access criteria to net settlement systems are strict.

■ Central counterparties in outline

A central counterparty is the industrial version of the kitchen table. It is worth understanding in outline because it is where netting is taken to its logical conclusion, and because central counterparties are now the largest single concentration of financial risk in the world.

Novation

The defining act of a CCP is novation. When two parties agree a trade and submit it for clearing, the original contract between them is legally extinguished and replaced by two new contracts: one between the buyer and the CCP, and one between the CCP and the seller. The CCP becomes buyer to every seller and seller to every buyer.

That single legal move accomplishes two things at once. It makes multilateral netting legally trivial, because every participant now has exactly one counterparty and can net everything it has against that one counterparty. And it removes the need for any participant to assess the creditworthiness of any other participant, because it deals only with the CCP.

Margin and the default waterfall

Because everybody's credit risk is now concentrated in one place, that place must be made extremely difficult to break. CCPs do this with a layered stack of financial resources, consumed in a fixed order known as the default waterfall.

Initial margin is collateral posted by each clearing member against the potential future loss on its portfolio, sized to cover the likely movement in that portfolio over the period it would take to close it out. LCH's SwapClear service calculates initial margin using its PAIRS methodology — Portfolio Approach to Interest Rate Scenarios — drawing on ten years of historical market data.

Variation margin is collected daily, and intraday where required, to settle the actual mark-to-market change in each portfolio. Variation margin is not collateral; it is payment of losses as they accrue, which prevents losses from accumulating.

The default fund is a mutualised pool contributed by all clearing members.

Skin in the game is a tranche of the CCP's own capital, deliberately placed in the waterfall ahead of the surviving members' contributions so that the CCP shares the loss before its members do.

At LCH the consumption order is the defaulting member's posted margin, then the defaulter's default fund contribution, then LCH's own capital, then the non-defaulting members' default fund contributions. If all of that is exhausted, variation margin gains haircutting is applied as a final measure: members who have made gains have those gains cut to fill the hole.

The international standard governing the size of these resources is the CPMI-IOSCO Principles for Financial Market Infrastructures, published in 2012. Principle 4 requires an FMI to maintain financial resources sufficient to cover the default of the participant to which it has the largest exposure in extreme but plausible market conditions — cover one — and requires CCPs involved in activities with

a more complex risk profile, or that are systemically important in more than one jurisdiction, to cover the two largest participants and their affiliates. That is the origin of “cover two”.

The default management process, and the one time it was tested at scale

When a clearing member defaults, a CCP does not simply liquidate. LCH’s SwapClear process runs in three stages: porting non-defaulting clients to surviving clearing members; neutralising the residual risk by hedging, with the assistance of a default management group drawn from member firms; and then auctioning the defaulter’s portfolio, split by product and currency, to the surviving members.

The definitive test was Lehman Brothers. At default, Lehman’s cleared interest rate derivatives portfolio at LCH comprised 66,390 trades with a notional value of \$9 trillion, against which it had posted initial margin of the order of \$2 billion. LCH completed the default management within about three weeks, using 100 per cent of the variation margin posted by the defaulting member and only 35 per cent of its initial margin. No mutualised default fund resources were consumed and no surviving member lost money.

That outcome is the strongest argument ever made for central clearing, and it is the reason the G20 leaders committed at Pittsburgh in September 2009 to central clearing of standardised over-the-counter derivatives, a commitment implemented in Europe through the European Market Infrastructure Regulation and retained in UK law after withdrawal from the European Union.

The concentration objection

The obvious criticism follows directly from the plain version of this chapter. Netting concentrates risk, and a CCP is netting taken to its limit. A CCP that fails does not fail one counterparty; it fails an entire market simultaneously, and the markets it clears may have no functioning alternative.

This is not a fringe concern. It is why CCPs are supervised directly by central banks — in the United Kingdom, by the Bank of England, which supervises central counterparties alongside central securities depositories and recognised payment systems — why the PFMI impose cover-two requirements, why CCP recovery and resolution regimes exist, and why the question of who bears loss beyond the waterfall is among the most contested in financial regulation.

One institutional fact captures the point. LCH Ltd is a CHAPS Direct Participant: a clearing house sitting directly on the Bank of England’s real-time gross settlement system, alongside the high street banks, because when a clearing house needs to move money it cannot afford to depend on anybody else’s netting.

■ Why Bacs takes three days

We can now answer the question properly.

The Bacs cycle has three named days, and the scheme documentation is precise about them.

Day 1, input day. This is the latest day on which a service user or bureau may submit a payment file for a given processing cycle. Files must be transmitted between 07:00 and 22:30. Submissions may be made up to 30 days in advance of the payment date, and a payment can be recalled after submission provided the submitter’s payment service provider is notified before a specified cut-off. Through the day, submitted data is validated and sorted into bank order by the central infrastructure.

Day 2, processing day. The sorted files are delivered to the recipient payment service providers, which process each payment into their own systems.

Day 3, entry day. Payments are simultaneously credited to recipients' accounts and debited from the payer's account. Interbank settlement of the multilateral net positions takes place at the Bank of England once that business day.

Now: what does each day buy?

The netting window. Netting only compresses obligations that coexist. A wider window means more offsetting. The Bacs cycle accumulates a full day of national payroll, benefits, utility collections and supplier payments before it computes anybody's position, and the result is the compression already quoted: about £23 billion of gross value discharged with about £5.3 billion of settlement.

Simultaneity of debit and credit. This is the point most commentary misses entirely. Bacs is often described as "three days of the money being in limbo". It is not. Nothing leaves the payer's account on day 1 or day 2. The debit and the credit both occur on day 3, at the beginning of the operating day, and the interbank settlement occurs on the same day. The three-day figure is the length of a pipeline, not the length of an exposure. That is a materially different risk profile from a system where the credit happens instantly and settlement happens later.

The exception window. Direct debits and direct credits can go wrong in specific, catalogued ways: a closed account, an insufficient balance, a cancelled instruction, a changed account number. Bacs provides a set of automated advice services for exactly these cases — ADDACS for direct debit amendments and cancellations, AUDDIS for mandate lodgement, ARUCS and AWACS for returned and amended credits. Those services require a defined window in which the receiving bank can inspect an item, decide whether it can be applied, and generate a report back to the originator. A cycle with no processing day has nowhere to put that work.

Predictability, for both sides. An originator running payroll for 12,000 people needs to know precisely which day the money lands, weeks in advance; a fixed cycle with a published calendar delivers that, and "as fast as possible" does not, because it is not a date. A bank, equally, that knows its Bacs settlement occurs once a day on a known set of items can forecast its central bank money position, arrange collateral and fund accordingly, which is far cheaper than holding a buffer against a continuous unpredictable drain.

So is three days a design choice or a defect?

The honest answer is: both, and the proportions have changed over time. Bacs began in 1968 as the Inter-Bank Computer Bureau, in an era when payment files travelled between banks as reels of magnetic tape moved by van. Three days had a physical meaning then that it does not have now, and it would be dishonest to pretend the current cycle length was derived from first principles in the present decade.

But the evidence for it being load-bearing rather than merely vestigial is strong, and it is this: when the United Kingdom decided it needed same-day retail payments, it did not speed Bacs up. It built the Faster Payment System alongside it in 2008, and left Bacs running exactly as it was. Seventeen years later Bacs is still carrying 27 million payments a day and, in 2025, 5.03 billion direct debits — an all-time annual high for the scheme.

That is what a design choice looks like from the outside. Two systems, two different points on the same trade-off curve, both running at national scale, because the trade-off is real and different users sit at different points on it.

	Bacs	Faster Payments	CHAPS
Settlement basis	Deferred multilateral net	Deferred multilateral net	Real-time gross
Settlement frequency	Once each business day	Three times each business day	Continuously
Customer credit timing	Day 3 of cycle	Within seconds, 24/7	Within the day
Value cap	None	£1,000,000	None
Liquidity cost to banks	Lowest	Low	Highest
Settlement risk between banks	Managed by prefunding	Managed by prefunding and net sender caps	Eliminated
Typical use	Payroll, benefits, direct debits	Person-to-person, bill payment	House purchase, wholesale, treasury

Read that table across and the logic is visible. CHAPS eliminates settlement risk and pays for it in liquidity. Bacs eliminates most of the liquidity requirement and pays for it in time and in the machinery — caps, prefunding, cycle discipline — needed to keep the risk contained. Faster Payments sits in between, buying speed for the customer while keeping the netting benefit for the banks, and paying for it with a value cap and three settlement cycles a day instead of one.

There is no free position on that curve. There never was. Netting is the mechanism by which the industry chooses where on it to sit, and every timing behaviour that irritates the public — the three days, the cut-off, the settlement that happens on Monday for a payment made on Saturday — is a consequence of that choice being made deliberately, by people who understood what they were trading away.

The next chapter leaves the interbank layer entirely and goes down to something much smaller and much more likely to break your code: what a currency amount actually is, and why the type you use to hold it matters more than almost anything else you will write.

8.98 Common wrong ideas

Wrong: Netting means somebody gets paid less. **Right:** Netting operates only on the obligations between banks; the customer layer stays gross and every payment is credited and debited in full.

Wrong: Netting is just arithmetic. **Right:** It is a legal act, and without designation under the Settlement Finality Regulations 1999 an administrator could cherry-pick the gross obligations and ignore the net figure.

Wrong: There is a pot in the middle holding the netted money. **Right:** In most net settlement systems the operator sends positions to the Bank of England, which debits and credits in a single simultaneous posting with nothing held anywhere in between.

Wrong: Netting reduces risk. **Right:** It concentrates it, replacing several small diversified credit exposures with one large exposure to a mechanism and to a single moment.

Wrong: A bank with little exposure to the failed participant is safe in an unwind. **Right:** In the worked example Eastbank's obligation rises fivefold, from £10 million to £50 million, because the effect propagates through everybody else's positions.

Wrong: Bilateral and multilateral netting come to much the same thing. **Right:** Multilateral netting lets a claim on one participant fund an obligation to another, taking the four-bank example from £290 million down to £200 million.

Wrong: Unwinding the cycle is the modern remedy when a participant cannot settle. **Right:** Since September 2015 the British answer for Bacs and Faster Payments has been prefunding: use the money the participant already put aside.

Wrong: Bacs means the money sits in limbo for three days. **Right:** Nothing leaves the payer's account on Day 1 or Day 2; the debit, the credit and the interbank settlement all fall on Day 3.

Wrong: Three days is purely a relic of magnetic tape moved around by van. **Right:** It is both relic and design choice, and the proof is that when Britain wanted same-day retail payments it built Faster Payments alongside Bacs in 2008 rather than speeding Bacs up.

Wrong: A central counterparty takes risk out of the market. **Right:** It removes bilateral credit assessment and makes multilateral netting legally trivial, but it becomes the largest single concentration of risk, which is why cover-two requirements and CCP resolution regimes exist.

8.99 Chapter summary in 20 lines

1. On an average working day in 2024 Bacs carried about 26.8 million payments worth roughly £23.0 billion, and about £5,265 million actually moved between settlement accounts.
2. Nobody was short-changed, because netting operates on the obligations between banks and never on the payments between customers.
3. Bilateral netting offsets obligations one pair at a time; multilateral netting offsets each participant against all the others simultaneously.
4. In the four-flatmate example, twelve debts and £102 of cash became four movements and £18 on the kitchen table, with every debt discharged in full.
5. The sum of all net debit positions always equals the sum of all net credit positions, because every debt is counted once as a plus and once as a minus.
6. In the four-bank matrix, £3,090 million of gross flows became £290 million bilaterally and £200 million multilaterally, or 9.4 and 6.5 per cent of gross.
7. Multilateral netting wins because a claim on one participant can fund an obligation to a different one, which bilateral netting cannot achieve.
8. Netting works best when flows are balanced, which is why Faster Payments eliminates something like 88 per cent of value while directional Bacs payroll traffic eliminates about 77 per cent.
9. The longer the window before squaring up, the more cancels out, so waiting is itself part of the compression.
10. Money in a settlement account is money doing nothing else, so netting is the difference between a payment system a country can afford to run and one it cannot.
11. The cost is that netting concentrates risk: each participant swaps several small exposures for one large exposure to a calculation.
12. Netting is a legal act as well as an arithmetical one, and only statutory protection stops an administrator cherry-picking the gross obligations.
13. Position netting leaves the gross obligations legally alive and gives rise to unwind risk, while netting by novation and close-out netting extinguish them.
14. Multilateral netting cannot be built from bilateral novation alone and needs either a central counterparty or statutory designation, or both.

15. An unwind redistributes obligations unpredictably, and it does so at the worst possible moment, when every bank in the country is already scrambling for liquidity.
16. The Lamfalussy Report of November 1990 set six minimum standards, of which cover one — settling on time despite the largest single debtor failing — changed the industry.
17. That standard is implemented through net debit and net sender caps, through prefunding introduced for Bacs and Faster Payments in September 2015, and through survivors-pay loss sharing where prefunding is absent.
18. A central counterparty is netting taken to its limit: novation makes it buyer to every seller and seller to every buyer, defended by initial margin, variation margin, its own skin in the game and a mutualised default fund.
19. LCH resolved Lehman's 66,390 cleared interest rate trades of \$9 trillion notional in about three weeks, using 35 per cent of the posted initial margin and no mutualised default fund resources at all.
20. Bacs takes three days because the netting window, the exception-handling window and a published calendar are worth more to its users than speed, and Britain proved the trade-off was real by building Faster Payments alongside it rather than instead of it.

Sources used: Bank of England, "Payment and settlement" and "A brief introduction to the Real-Time Gross Settlement system and CHAPS" (including Table A, daily average RTGS settlement values and volumes 2024); Bank of England, "Bank of England Settlement Accounts" and "RTGS and CHAPS Service Description", December 2018; Pay.UK, "Annual summary of payment statistics 2025"; Pay.UK, "Bacs Payment System" and "Understanding changes to the Faster Payment System Net Sender Caps"; Bacs, "Bacs Direct Credit — Getting started" (processing cycle); Committee on Payments and Market Infrastructures, Glossary, and Report of the Committee on Interbank Netting Schemes (the Lamfalussy Report), November 1990; CPMI-IOSCO, Principles for Financial Market Infrastructures, 2012; CPMI Red Book, "Payment, clearing and settlement systems in the United Kingdom"; The Financial Markets and Insolvency (Settlement Finality) Regulations 1999, SI 1999/2979, and Directive 98/26/EC; DTCC, UST1 FAQs and press release of 23 April 2025; CLS Group, "What are the necessary ingredients for any PvP arrangement to work?"; BIS Quarterly Review, "Uncovering FX settlement risk: new measures from the 2025 BIS Triennial Survey"; LSEG, LCH SwapClear risk management; CCP Global, "The Lehman Case".

Currency, Decimals and the Money Type

9.0 What this chapter gives you

1. You will be able to explain why ten 10p coins added together as pounds come to 0.9999999999999999, and why a till written to notice exactly £1.00 will never notice.
2. You will be able to say why multiplying a yen amount by 100 charges the cardholder a hundred times too much, and why a Kuwaiti amount loses its last digit in a two-decimal system.
3. You will be able to derive a currency's minor unit from ISO 4217 rather than assuming 100, and name the exceptions: seventeen zero-decimal codes, seven three-decimal, two four-decimal and thirteen with no minor unit at all.
4. You will be able to show that the 42-pence gap across 3,000 merchant fees came from an unwritten decision about where the rounding line falls, and not from a data type.
5. You will be able to split a total across parts by the largest remainder method so that the parts always sum back to the total exactly.
6. You will be able to say why bankers' rounding is not "the correct rounding for money", and cite the euro conversion rule and the VAT Notice 700 treatments that override it.
7. You will be able to explain why a JSON number with a decimal point is the one wire convention that guarantees eventual trouble, and name the two that do not.
8. You will be able to carry one internal representation — an integer count of minor units with an explicit currency — across ISO 8583, ISO 20022, SWIFT MT and Bacs Standard 18.
9. You will be able to record a foreign exchange conversion so that an independent party can re-derive it, and say why a reference rate is not a transaction rate.
10. You will be able to name the eight recurring failure modes of money handling and the single test that catches each one.

There is a category of bug that does not crash anything. The service stays up, the tests pass, and every individual transaction looks correct when you open it. Then, at the end of the month, the settlement report is three pence away from the bank statement and nobody can say why. Three pence is not a material amount of money. Explaining three pence to an auditor, and proving that it is not in fact three pence multiplied by every account you have not checked yet, can consume a fortnight.

This chapter is about the machinery that produces those three pence, and how to build a system in which they cannot occur. It is the most directly practical chapter in this volume for anyone who writes code, and it is also the chapter that finishes the volume's argument. Chapter two established that a balance is a record rather than a substance. Chapter three established that every payment is two facts rather than one. This chapter establishes what kind of thing that record can be made of, and the answer is surprisingly constrained. Money is not a number. Money is a number, plus a currency, plus

a rule about how finely that currency may be cut, plus a written decision about what happens when a calculation lands between two of the permitted cuts.

Leave out any one of those four and you have not built a ledger. You have built something that usually agrees with a ledger.

The plain version

Think about a school tuck shop. Everything is priced in whole pennies: a chocolate bar at 65p, a packet of crisps at 48p, a drink at £1.20. At the end of the day the person running the shop tips the cash box onto the table and counts.

Nobody counts in pounds. They count in coins — four £1 coins, seven 50p pieces, twenty-three 20p pieces — and then convert everything into a single number of pennies, because pennies are the smallest thing the shop can actually receive. Nine hundred and forty-two pennies. Only at the very end, writing the figure on the sheet, does anybody turn 942 into “£9.42”.

That is the whole trick, and it is worth saying plainly because almost every serious money system in the world works this way. **Count in the smallest coin. Convert to pounds only when you print something a human will read.**

■ Why not just count in pounds?

Because computers cannot hold “one tenth” exactly, and this is not a matter of them being sloppy. It is a matter of the ruler they use.

Imagine a ruler with no centimetre marks. Instead, someone has marked the halfway point, then the halfway points of each half, then the halfway points of those, and so on forever. You can mark exactly one half, one quarter, three eighths, eleven sixteenths. You can get extremely close to one tenth — closer than any eye could see. But you can never land exactly on it, because one tenth is not reachable by halving.

The ordinary fractional numbers a computer uses are built on exactly that ruler. And so when a computer stores the number 0.1, it does not store one tenth. It stores the nearest mark on the halving ruler, which is:

```
0.1000000000000000055511151231257827021181583404541015625
```

That is a real number, not an exaggeration for effect. It is the exact value your laptop puts in memory when you type 0.1. Now watch what happens. Ask a computer to add 0.1 and 0.2 and it produces:

```
0.3000000000000000444089209850062616169452667236328125
```

which it will politely display to you as 0.30000000000000004. Ask it whether 0.1 plus 0.2 equals 0.3, and it says no.

Try something even more homely. Put ten 10p coins into a till, one at a time, using pounds:

0.1, then 0.2, then 0.30000000000000004, then 0.4, then 0.5, then 0.6, then 0.7, then 0.7999999999999999, then 0.8999999999999999, and finally 0.9999999999999999.

Ten ten-pence coins have not made a pound. They have made a hair less than a pound. If the till software is written to notice when the drawer contains exactly £1.00, it will never notice.

Count in pennies instead and the problem does not arise. Ten lots of 10 is 100. Whole numbers on a computer are exact. There is nothing to go wrong.

■ Not every country cuts money the same way

Here is the second idea, and it catches people out constantly because it is invisible from Britain. We assume every currency has a hundred small units in a big one. A pound has 100 pence, a dollar has 100 cents, a euro has 100 cents. So a programmer writes $\text{pence} = \text{pounds} * 100$ and moves on.

Japan does not work like that. The yen has no subdivision in use. A price of 2,500 yen is 2,500 yen; there is no such thing as 2,500.00. If you take a ¥2,500 payment and helpfully multiply by 100 before sending it, you have just asked for ¥250,000 — a hundred-fold error, and it will go through, because the number is perfectly valid.

Kuwait goes the other way. The Kuwaiti dinar is divided into 1,000 fils, so amounts carry three decimal places: 12.345 dinars is an ordinary price. Bahrain, Jordan, Oman, Iraq, Libya and Tunisia do the same. If your system assumes two decimal places, it will quietly shave the last digit off every Kuwaiti amount it touches.

So the rule is not “multiply by 100”. The rule is “look up how many places this particular currency has, and multiply by that”. There is an official list. It is called ISO 4217, and it gives every currency a three-letter code — GBP, USD, JPY, KWD — and tells you how many decimal places that currency uses.

■ A number on its own is not money

Which brings us to the third idea. “£19.99” is money. “19.99” is not; it is a quantity of something unspecified. This sounds like pedantry and it is the single most useful discipline in the chapter. If a value in your system can be 19.99 without also carrying the letters GBP, then sooner or later somebody will add it to a euro amount, or send it to Japan, or store it in a column another team reads as dollars. The currency is not decoration on the amount. It is part of the amount, the way “miles” is part of “sixty miles”.

■ Splitting things up

The last idea is the one that produces the three pence.

Three friends split a £5.00 bill. Five divided by three is £1.6666... Round each share to the nearest penny and each pays £1.67. Three times £1.67 is £5.01: they have paid a penny more than the bill.

Try it the other way. Three people split £100.00. A hundred divided by three is £33.3333... Round each to the nearest penny and each pays £33.33. Three times £33.33 is £99.99. A penny has vanished.

There is no clever arithmetic that avoids this. Five pounds genuinely does not divide into three equal whole numbers of pennies. What a well-built system does is not “avoid the problem” — it is **decide, in advance and in writing, who gets the odd penny, and then make sure the parts always add back up to the whole**. In the £5.00 case: two people pay £1.67 and one pays £1.66. In the £100.00 case: two pay £33.33 and one pays £33.34.

That is it. Count in the smallest coin. Look up how many small coins are in a big one for that particular currency. Never let an amount travel without its currency. And when you divide a total, force the pieces to add back up to it.

Where the plain version stops being true

The plain version is the right instinct and it will keep a small system honest for years. Four things about it are wrong in ways that show up at scale, and the fourth is the one that costs money.

■ “Count in whole pennies” is only true at the settlement boundary

The tuck-shop picture implies that every number in a financial system is a whole number of minor units. It is not, and insisting that it is will produce worse errors than it prevents.

Plenty of legitimate financial quantities are finer than a penny. Petrol in the United Kingdom is priced at, for example, 149.9 pence per litre. Interest accrues daily at rates like 4.37% per annum. Telecommunications and cloud services bill per second or per gigabyte at rates with four, six or eight decimal places. Foreign exchange rates carry six significant figures by law in the euro area. HM Revenue and Customs explicitly contemplates VAT calculated per article to three decimal places of a pound.

The correct statement is narrower and more useful. **The amount that settles must be a whole number of minor units. Intermediate quantities need not be, and the design question is where the rounding line falls.** A rate is not money. A price per unit is not money. A running accrual is not money until it is capitalised. Money is what leaves one account and arrives in another, and that has to be a whole number of the smallest thing the payment system can carry.

A system that rounds too early loses accuracy for no reason. A system that rounds too late presents a settlement instruction that no rail will accept. Both failures come from not writing down where the line is.

■ Floating point is not “imprecise” — it is exact in the wrong base, and that is worse

The plain version says computers cannot hold one tenth. True, but the framing invites the wrong conclusion: that the errors are random, tiny, and can be mopped up by rounding at the end.

First, binary floating point is not approximate. Every value it holds is exactly a whole number multiplied by a power of two, and the arithmetic is exactly specified and perfectly repeatable. The problem is that the set of values it can hold does not include the numbers a ledger is made of, so a calculation that ought to be exact acquires a small, deterministic, unpredictable-by-eye residue. That is worse than randomness, because it survives testing. Every one of your test cases will pass and the one combination you did not try will be off by a hair.

Second, and more practically: the damage almost never arrives as a visibly wrong total. It arrives in three other ways.

It arrives as a failed equality test. Debits of £1,234.56, £78.90, £4,321.00, £0.07 and £999.99 against a single credit of £6,634.52 plainly balance. Add the debits in binary floating point and you get 6634.5199999999995, so `total_debits == total_credits` returns false. Your double-entry validator, the one from chapter three that is supposed to be the last line of defence, now rejects a correct journal — or, if somebody has “fixed” it with a tolerance, now accepts an incorrect one.

It arrives as rounding that goes the wrong way. Ask a computer to round 1.15 to one decimal place and it returns 1.1, not 1.2 — not because the rounding rule is wrong but because the value actually stored for 1.15 is 1.149999999999999118..., genuinely below the midpoint. The function did exactly what it was asked. The input was never what you thought.

And it arrives at serialisation boundaries, because JSON has no decimal type and a JSON number is in practice parsed into a binary double. An amount that was exact in your database becomes inexact the moment it crosses an API.

And here is the part that reframes the whole subject. In a run of 3,000 sales at £19.99, the floating-point error in the total is about 1.6 billionths of a penny — real, but not what will cost you money. Now apply a merchant fee of 1.4% plus 20p to each of those sales. If you round the fee per transaction, each is

27.986p plus 20p, rounded to 48p, and the total is £1,440.00. If you compute the fee on the whole run, it is £1,439.58. **The gap is 42 pence, and floating point had nothing to do with it.** It came entirely from an unwritten decision about where the rounding line falls. That is the shape of nearly every real money bug: a rounding policy nobody specified, not a data type nobody understood.

■ **Minor units are a convention, and the convention is not universally agreed**

The plain version implies that “how many decimal places does this currency have” is a fact you can look up once. It is a convention recorded by a standards body, and there are at least three ways it drifts from reality.

First, ISO 4217 itself rounds off awkward cases. The Malagasy ariary is notionally divided into five *iraimbilanja* and the Mauritanian ouguiya into five *khoums* — the only two circulating currencies whose subdivision is not a power of ten. ISO 4217 assigns both a minor unit of 2 anyway, because a field that holds an exponent has nowhere to record “one fifth”. The standard is doing something sensible; it is simply not describing the currency.

Second, processors disagree with the standard, and with each other, for backwards-compatibility reasons. Stripe’s published list of zero-decimal currencies includes MGA, which ISO 4217 gives a minor unit of 2. In the other direction, ISO 4217 gives both the Icelandic króna and the Ugandan shilling a minor unit of 0, but Stripe’s documentation states that both “transitioned to a zero-decimal currency” while “backward compatibility requires you to represent it as a two-decimal value, where the decimal amount is always 00” — so to charge 5 ISK you submit 500. Stripe further treats the Hungarian forint and the New Taiwan dollar as zero-decimal for payouts while accepting two-decimal charges in them, requiring payout amounts to be evenly divisible by 100. None of this is wrong. All of it means “the minor unit” is a property of a currency *within a particular scheme’s rules*, not of the currency alone.

Third, the smallest unit you can *account* in and the smallest unit you can *pay* in are different things. Several countries have withdrawn their smallest coin and round cash transactions to the nearest five units at the till while still recording electronic payments to the single unit. A ledger that models one granularity cannot represent the gap between the invoiced amount and the cash tendered.

The practical consequence: derive the exponent from a table you control and can version, validate it against the scheme you are actually talking to, and never from a hard-coded 100 or from whatever your standard library’s locale data says this month.

■ **Bankers’ rounding is not “the correct rounding for money”**

“Use bankers’ rounding for money” circulates as settled wisdom. It is not settled, and applying it where the law says otherwise is a compliance problem rather than a style preference.

Round-half-to-even exists for a good reason: rounding halves consistently upwards introduces a systematic upward bias. Take the ten values 0.005, 0.015, 0.025 and so on up to 0.095. Their true sum is exactly 0.50. Round each half upwards and they sum to 0.55; round each to the nearest even last digit and they sum to 0.50. Over a large book of similar figures the half-up rule drifts and the half-even rule does not. That is a genuine and important property.

But it is a statistical argument, and statistics do not override statute. Council Regulation (EC) No 1103/97, which still governs conversion between the euro and the legacy national currencies it replaced, says in Article 5: “If the application of the conversion rate gives a result which is exactly half-way, the sum shall be rounded up.” Not to even. Up. In the United Kingdom, VAT Notice 700 permits an invoice trader to round the *total* VAT on an invoice **down** to a whole penny — an explicit concession

in the taxpayer's favour — while stating that the same concession “is not appropriate to retailers”, who “must not round the VAT figure down”.

So there is no single rounding mode for money. There are modes mandated by specific regimes for specific calculations, and the job is to know which regime you are in. A system that applies half-even everywhere because a blog post recommended it will produce VAT figures HMRC's own guidance does not describe.

The second half of this correction is the one that generates reconciliation breaks. **Rounding each component and then adding is not the same as adding and then rounding.** Three items at £8.99 with VAT at 20%: each line's VAT is £1.798, which rounds to £1.80, so the invoice total from line rounding is £5.40. Compute VAT on the £26.97 order total and you get £5.394, which rounds to £5.39. Both figures are defensible. They are one penny apart. If your invoice engine uses one method and your settlement engine the other, you will be reconciling that penny for the life of the product.

The technical version

■ ISO 4217: the code, the number and the exponent

ISO 4217 is the international standard for representing currencies and funds; the current edition is ISO 4217:2015. The maintenance agency is SIX Financial Information AG, acting on behalf of the Swiss Association for Standardization, which publishes the authoritative tables and issues amendments as currencies are introduced, redenominated or withdrawn.

The standard defines three tables. **List One** contains the currently active currency and funds codes. **List Two** contains fund codes. **List Three** is the historical list of withdrawn codes — a list that exists precisely because currency codes are not permanent, and a system that stores only a code, with no effective date, cannot always reconstruct what an old record meant.

Each entry in List One carries the country or entity, the currency name, a three-character alphabetic code, a three-digit numeric code, and the minor unit — the number of decimal places conventionally used. The alphabetic code is normally the two-letter ISO 3166-1 country code followed by the initial of the currency name: GB plus P gives GBP, US plus D gives USD, JP plus Y gives JPY. Supranational and non-country codes begin with X, which ISO 3166-1 reserves for exactly this purpose. That is why gold is XAU, silver XAG, platinum XPT and palladium XPD; why the IMF's Special Drawing Right is XDR; why the CFA francs are XOF and XAF; and why XTS is reserved for testing and XXX denotes “no currency involved”. The three-digit numeric code exists because a great deal of financial messaging is numeric-only: GBP is 826, USD is 840, EUR is 978, JPY is 392, INR is 356, KWD is 414.

As published on 1 January 2026, List One contains 178 distinct alphabetic codes spread over 280 country-and-currency rows. The overwhelming majority — 139 codes — have a minor unit of 2. The exceptions are the ones your code has to handle.

Minor unit	Count	Codes (with numeric code)
0	17	BIF 108, CLP 152, DJF 262, GNF 324, ISK 352, JPY 392, KMF 174, KRW 410, PYG 600, RWF 646, UGX 800, UYI 940, VND 704, VUV 548, XAF 950, XOF 952, XPF 953
2	139	GBP 826, USD 840, EUR 978, INR 356 and 135 others

Minor unit	Count	Codes (with numeric code)
3	7	BHD 048, IQD 368, JOD 400, KWD 414, LYD 434, OMR 512, TND 788
4	2	CLF 990 (Unidad de Fomento), UYW 927 (Unidad Previsional)
N.A.	13	XAG 961, XAU 959, XBA 955, XBB 956, XBC 957, XBD 958, XDR 960, XPD 964, XPT 962, XSU 994, XTS 963, XUA 965, XXX 999

Three observations. The four-decimal entries are not ordinary currencies: CLF, Chile's Unidad de Fomento, and UYW, Uruguay's Unidad Previsional, are inflation-indexed units of account, and they appear in any complete currency table you load, so a validator assuming nothing exceeds three decimals will reject them. The thirteen "N.A." entries have no minor unit at all — XAU is not a currency but a troy ounce of gold, and if your money type looks up its exponent and defaults to 2 on a miss, you have silently invented centigrams of gold. And XXX exists: it is the correct code for a transaction where no currency is involved. It is not a null and it is not an error.

■ Representing an amount: the money type

A money value is a triple: an integral quantity, a currency, and the exponent that relates them. Systems usually store the first two and derive the third, which is defensible provided the currency table is versioned and the derivation happens once, at the edge.

The canonical representation is **an integer count of minor units**. GBP 19.99 is stored as 1999 with currency GBP; JPY 2500 as 2500 with currency JPY; KWD 12.345 as 12345 with currency KWD. There is no decimal point anywhere in storage. Formatting for display happens at the last possible moment, and parsing user input at the first.

The range question is usually settled quickly. A signed 64-bit integer holds up to 9,223,372,036,854,775,807, which in pence is £92,233,720,368,547,758.07 — comfortably more than the money supply of any country. It is not unlimited: currencies with large nominal values at institutional scale, hyperinflationary histories, and any field that also stores intermediate accruals at higher precision will get there sooner than you expect.

Where an integer is not enough, use a decimal type — a numeric type whose radix is ten, so that one tenth is exactly representable.

Type	Representation	Practical limit
<code>java.math.BigDecimal</code>	Arbitrary-precision unscaled BigInteger with a 32-bit scale	Effectively unbounded; scale is per-value and preserved
<code>decimal.Decimal</code> (Python)	Arbitrary-precision, context-controlled	Default context precision of 28 significant digits, adjustable
<code>System.Decimal</code> (.NET)	128 bits: a 96-bit integer plus sign and a scaling factor of 10 to the power 0 through 28	Maximum 79,228,162,514,264,337,593,543,950,335
IEEE 754 <code>decimal32</code> / <code>decimal64</code> / <code>decimal128</code>	Decimal floating point	7, 16 and 34 significant decimal digits respectively
SQL <code>NUMERIC</code> / <code>DECIMAL</code>	Implementation-defined exact numeric	PostgreSQL: up to 131,072 digits before and 16,383 digits after the decimal point

The PostgreSQL documentation is unusually direct about which of its numeric types to use: numeric “is especially recommended for storing monetary amounts and other quantities where exactness is required”, while real and double precision “are inexact, variable-precision numeric types” where “some values cannot be converted exactly to the internal format and are stored as approximations, so that storing and retrieving a value might show slight discrepancies.”

Two traps around that. PostgreSQL also ships a type literally called money, whose fractional precision comes from the database’s `lc_monetary` setting, so its behaviour depends on server configuration and it carries no currency identifier. And an object-relational mapper that maps a NUMERIC column to a language-level double undoes the benefit of the column type without producing any error at all. Check the mapping, not the schema.

The serialisation boundary deserves the same suspicion. JSON has no decimal type; JSON numbers are commonly parsed into IEEE 754 binary64, and in JavaScript integers above `Number.MAX_SAFE_INTEGER`, which is 9,007,199,254,740,991, lose exactness. The two safe conventions are to transmit an integer count of minor units alongside a currency code, as most card-processing APIs do, or to transmit the amount as a JSON string and parse it into a decimal type, as many banking APIs do. A JSON number with a decimal point is the one option that guarantees eventual trouble.

■ Why binary floating point cannot carry money

IEEE 754 binary64 — double in C, Java and JavaScript, float in Python — is 64 bits: one sign bit, an 11-bit exponent, and 52 stored significand bits giving 53 bits of effective precision. Every finite value it represents is of the form m times 2 to the power e , with m and e integers, so the representable numbers are exactly the dyadic rationals within range: fractions whose denominator is a power of two.

One tenth is not a dyadic rational. Neither is one hundredth. The two quantities a currency system is built out of are precisely the quantities binary floating point cannot hold. The nearest binary64 value to 0.1 is

```
0.1000000000000000055511151231257827021181583404541015625
```

and the nearest to 0.2 is

```
0.200000000000000011102230246251565404236316680908203125
```

Their exact sum, correctly rounded to the nearest binary64 value, is

```
0.3000000000000000444089209850062616169452667236328125
```

whereas the nearest binary64 value to the literal 0.3 is

```
0.299999999999999988897769753748434595763683319091796875
```

These are different numbers. The arithmetic is not wrong; it is exactly right about a question you did not intend to ask.

The consequences in a ledger are specific and testable:

Equality fails. Adding the debits £1,234.56, £78.90, £4,321.00, £0.07 and £999.99 in binary64 yields 6634.519999999995 against a credit of 6634.52 — a residue of about minus 9.1 times 10 to the minus 13. Double-entry validation by exact comparison rejects a correct journal.

Accumulation drifts. Adding 0.01 one million times gives 10000.000000171856. Adding 0.07 one thousand times gives 69.9999999999966.

Rounding lands on the wrong side. `round(1.15, 1)` gives 1.1 and `round(2.675, 2)` gives 2.67, because the stored values are 1.149999999999999118... and 2.67499999999999982236..., both genuinely below the midpoint. The nearest-even rule is not the culprit; the input never was a midpoint.

Truncation compounds. The Vancouver Stock Exchange index is the standard historical illustration, and because it is a truncation error rather than a binary-representation error it is the more general lesson. The index was established at 1,000.000 in January 1982 and recalculated on every trade, with each updated value truncated rather than rounded to three decimal places. Losing a fraction of a thousandth on each of thousands of daily recalculations, it drifted downwards relentlessly, and by the close on Friday 25 November 1983 it stood at 524.811. Over the weekend of 25 to 28 November the calculation was corrected and the index restated at 1,098.892 — more than double the published figure. Nothing had happened in the market. The index had simply been throwing away a sliver, several thousand times a day, for twenty-two months.

The deeper argument against floating point is not accuracy at all. A ledger has to be **reproducible by an independent party**. Your merchant, your bank, your auditor and your regulator must be able to take the same inputs and the same stated rules and arrive at your figure, digit for digit, using different software. Exact decimal arithmetic on integer minor units makes that trivial. Binary floating point makes it a research project.

■ Rounding: modes, mandates, and the invariant

The named rounding modes a practitioner needs are few.

Mode	Behaviour at a tie	Also called
Round half up	Away from zero	Round half away from zero, commercial rounding, arithmetic rounding
Round half down	Towards zero	
Round half to even	To the nearest even last digit	Bankers' rounding; IEEE <code>roundTiesToEven</code>
Round towards zero	Not applicable; always truncates	Truncation, chopping; IEEE <code>roundTowardZero</code>
Round towards positive infinity	Always upward	Ceiling; IEEE <code>roundTowardPositive</code>
Round towards negative infinity	Always downward	Floor; IEEE <code>roundTowardNegative</code>

IEEE 754 specifies `roundTiesToEven` as the default attribute for binary formats and requires `roundTiesToAway` — round half away from zero — to be available for decimal formats, which is the standard acknowledging that commercial practice does not universally use ties-to-even.

The bias argument for ties-to-even is worth being able to demonstrate. The ten values 0.005, 0.015, 0.025, 0.035, 0.045, 0.055, 0.065, 0.075, 0.085 and 0.095 sum to exactly 0.500. Rounded half up to two places they become 0.01 through 0.10 and sum to 0.55, an overstatement of ten per cent of the total. Rounded half to even they become 0.00, 0.02, 0.02, 0.04, 0.04, 0.06, 0.06, 0.08, 0.08, 0.10 and sum to exactly 0.50.

Now the mandates, which override the argument.

Euro conversion. Council Regulation (EC) No 1103/97 fixes the arithmetic in law. Conversion rates “shall be adopted as one euro expressed in terms of each of the national currencies” and “shall be

adopted with six significant figures”. They “shall not be rounded or truncated when making conversions”. “Inverse rates derived from the conversion rates shall not be used” — you may not divide by the published rate to go the other way. Amounts moving between two national currency units “shall first be converted into a monetary amount expressed in the euro unit, which amount may be rounded to not less than three decimals and shall then be converted into the other national currency unit” — mandatory triangulation through the euro. And on ties, Article 5: “If the application of the conversion rate gives a result which is exactly half-way, the sum shall be rounded up.”

United Kingdom VAT. VAT Notice 700 sets out the permitted treatments. At paragraph 17.5, “You may round down the total VAT payable on all goods and services shown on a VAT invoice to a whole penny”, with the caveat that this “concession to round down amounts of VAT is designed for invoice traders and applies only where the VAT charged to customers and the VAT paid to HMRC is the same.” At 17.5.1, where VAT is worked out per line, the trader must round “down to the nearest 0.1 pence” or “to the nearest 1 pence or 0.5 pence”, and “Whatever you decide, you must be consistent.” At 17.5.2, VAT per unit or per article must be worked to “4 digits after the decimal point and then round to 3 digits” or to the nearest penny or half penny, with an explicit prohibition on rounding down to nil on any unit liable at the standard or reduced rate. At 17.6, retailers who calculate VAT at line or invoice level “must not round the VAT figure down”, though they “may round (up and down) each VAT calculation”.

Read those two regimes together and the design conclusion is inescapable: **the rounding mode is a per-calculation configuration item with a legal basis, not a global constant.**

The invariant that must hold regardless of mode is this. For any total that is broken into parts, the parts must sum exactly to the total, in integer minor units, with no residual. Enforce it as an assertion, not as a hope.

■ Allocation: dividing a total without creating or destroying money

Naive per-part rounding violates that invariant in both directions. Splitting GBP 5.00 three ways gives 166.666... minor units per part; rounding each gives 167, and three parts of 167 total 501 — a penny created from nothing. Splitting GBP 100.00 three ways gives 3,333.33... per part; rounding each gives 3,333, totalling 9,999 — a penny destroyed.

The standard fix is the **largest remainder method**, borrowed from apportionment in electoral mathematics. Compute each part’s exact share as a rational number of minor units, take the floor of each, and subtract the sum of the floors from the total to get the residual — always a non-negative integer strictly less than the number of parts. Distribute that residual one minor unit at a time to the parts with the largest fractional remainders, breaking ties by a deterministic rule, such as line sequence, recorded in the specification.

Worked, with a GBP 10.00 order-level discount to be apportioned across three lines of GBP 3.99, GBP 12.50 and GBP 7.51:

Line	Line total (minor units)	Exact share of 1000	Floor	Fractional remainder	Final
A	399	166.2500	166	0.2500	166
B	1250	520.8333	520	0.8333	521
C	751	312.9167	312	0.9167	313
Total	2400	1000.0000	998	residual 2	1000

The floors total 998, leaving a residual of two minor units, which go to lines C and B in descending order of remainder. The allocation is GBP 1.66, GBP 5.21 and GBP 3.13, summing exactly to the GBP 10.00 discount.

The alternative is the **running-remainder method**: process the parts in a fixed order, and for each allocate the difference between the rounded cumulative target and the cumulative amount already allocated. It also guarantees the invariant and is order-dependent by construction, which is fine as long as the order is deterministic and documented.

Whichever you choose, the properties to test are the same. The sum of the parts equals the total exactly. Every part carries the same sign as the total, so an allocation of a positive amount never emits a negative part. The function is deterministic. And the result is stable under reversal: allocating a refund of the same total across the same parts must produce the same pieces, or your refund will not net your original to zero.

The places this bites in production are consistent across the industry: marketplace payouts split among sellers, order-level discounts and shipping apportioned across lines for VAT, subscription proration on mid-cycle upgrades, interchange and scheme fees attributed to individual transactions from a netted invoice, and partial refunds allocated back across multiple tenders.

■ Multi-currency, in outline

Two rules first, because they eliminate most multi-currency defects before they exist. **An amount never exists without its currency**, enforced in the type system where the language permits it. **Two amounts in different currencies are not addable**, and an attempt should be an error rather than a number.

A ledger account is denominated in exactly one currency. A multi-currency ledger is therefore not a ledger of mixed rows; it is a set of single-currency ledgers, plus explicit foreign exchange gain and loss accounts to absorb the difference when a position is revalued. IAS 21, *The Effects of Changes in Foreign Exchange Rates*, requires foreign currency monetary items to be translated at the closing rate at each reporting date, with the resulting exchange differences recognised in profit or loss in the period in which they arise. That is what those accounts exist to satisfy.

Any conversion must be recorded with five things, or it cannot be re-derived: source amount and currency, target amount and currency, the rate as an exact decimal, the quotation direction, and the timestamp with the identity of the rate source. A stored rate of “1.1642” with no direction is not information. Market convention writes a pair as base and quote — EUR/USD 1.0850 means one euro buys 1.0850 dollars — but conventions are not universal across systems, so store the direction rather than infer it.

Reference rates are not transaction rates, and the institutions that publish them say so plainly. The European Central Bank’s euro foreign exchange reference rates are, in the ECB’s own words, “usually updated at around 16:00 CET every working day, except on TARGET closing days”, based on “the daily concertation procedure between central banks across Europe, which normally takes place around 14:10 CET”. The ECB adds that they “are published for information purposes only” and that “Using the rates for transaction purposes is strongly discouraged.” A rate you can deal on comes from a counterparty with a bid and an offer around it, and the mid-point between the two is not a price anyone will give you.

Rounding at conversion follows the target currency, not the source: converting into JPY must produce a whole number of yen, converting into KWD a multiple of one fils. Converting an amount, rounding it, and converting it back will not return the original, and any process that assumes it does — a refund path, a reversal, a reconciliation — is broken. So is chained conversion: A to B to C does not in general

equal A to C, which is exactly why Regulation 1103/97 mandates a specific triangulation route with a specified intermediate precision rather than leaving it to the implementer.

■ How the amount travels on the wire

The four message families a UK payments engineer meets carry amounts four different ways, and the differences are instructive.

ISO 8583, the card authorisation format taken apart in Volume III, carries the amount in data element 4, *Amount, Transaction*, a fixed-length twelve-digit numeric field. There is no decimal point and no exponent in the field itself; the position of the decimal is implied entirely by data element 49, *Currency Code, Transaction*, which carries the three-digit ISO 4217 numeric code. A charge of GBP 19.99 travels as 000000001999 with DE49 set to 826. A charge of JPY 2,500 travels as 000000002500 with DE49 set to 392 — and if the terminal has multiplied by 100 on the way in, the message is still perfectly well-formed and the cardholder is charged a hundred times too much. Nothing in the format can detect that error.

ISO 20022, the XML-based successor, carries the amount as a decimal with the currency as an attribute on the element: an instructed-amount element whose Ccy attribute is GBP and whose text content is 19.99. The schema type restricts the value to at most 18 total digits and 5 fractional digits, with a minimum of zero. That schema is deliberately permissive — 5 fractional digits exceeds any List One currency — and the real constraint is a business rule stated in the Message Definition Reports: “The number of fractional digits (or minor unit of currency) must comply with ISO 4217.” A validator accepts 10.21, 10.2 and 10 for EUR and rejects 10.403. Schema validation alone is not sufficient validation.

SWIFT MT messages carry the amount in fields such as 32A, *Value Date, Currency Code, Amount*, with the format 6!n3!a15d — six numeric characters of date in YYMMDD, three alphabetic characters of currency code, and an amount of up to fifteen characters. The network validated rules are explicit: “The integer part of Amount must contain at least one digit. The decimal comma ‘,’ is mandatory and is included in the maximum length. The number of digits following the comma must not exceed the maximum number allowed for that specific currency as specified in ISO 4217.” The decimal separator is a comma, it is mandatory even for a whole amount — five hundred thousand is written 500000, — and it consumes one of the fifteen characters.

Bacs Standard 18, the fixed-length format carrying UK Direct Debit and Direct Credit submissions, takes the strictest line of the four. The amount occupies eleven characters at positions 36 to 46 of the record and is expressed in pence: all numeric, not all zeros, right-justified and zero-filled. There is no currency field, because the service is sterling only, and no decimal point anywhere in the file. A payment of £19.99 is written 00000001999.

One internal representation; four adapters. Hold the amount as an integer count of minor units with an explicit currency, convert at each boundary, and look the exponent up from a versioned table rather than assuming it.

■ The failure modes, and the test that catches each

Failure mode	Symptom	Test
Binary floating point in a ledger field	Reconciliation off by fractions of a penny; balance checks fail sporadically	Assert exact equality of debits and credits with no tolerance

Failure mode	Symptom	Test
Hard-coded exponent of 2	Yen charged 100 times too much; files truncated on Kuwaiti amounts	Round-trip a JPY, a KWD and a CLF amount through every boundary
Naive per-part rounding	Allocated parts sum to one minor unit more or less than the total	Property test: for random totals and weights, sum of parts equals total
Unspecified rounding line	Two subsystems disagree by a penny per transaction	Assert invoice total equals the sum of line totals
Amount without currency	Cross-currency addition produces a plausible wrong number	Make currency a required field; forbid addition of unlike currencies
Conversion without stored rate	Historical figures cannot be re-derived	Persist source, target, rate, direction, timestamp and rate source
JSON number for an amount	Value changes across an API round trip	Transmit integer minor units or strings; assert identical round trip
Currency code without effective date	Redenominated historical records misread	Store the effective date; retain List Three codes rather than remapping

The point of all this

This chapter belongs in a volume called *What Money Is*, rather than in an appendix on software engineering, because the money type is where the volume's argument becomes physical.

If money were a substance, its representation would be a detail. You would weigh it, and the scales would either be accurate or not. But money is a record — a claim, recorded twice, denominated in a unit some authority defines. And a record has to be *exact*, not merely accurate, because its whole function is that two parties who do not trust each other can compute the same answer from it and be bound by the result. A ledger uses integer minor units and a documented rounding rule not out of fastidiousness, but because the alternative — a number that is very nearly right, arrived at by a process the counterparty cannot reproduce — is not a record of anything.

Three pence is not a material amount of money. But three pence you cannot explain is a defect in the only property a ledger has.

9.98 Common wrong ideas

Wrong: Every number in a financial system must be a whole number of minor units. **Right:** Only the amount that settles must be; rates, per-unit prices and running accruals are legitimately finer, and the design question is where the rounding line falls.

Wrong: Floating point is merely imprecise, so the small random errors can be mopped up by rounding at the end. **Right:** It is exact in the wrong base, so the residue is deterministic and survives testing — every case you thought to write passes, and the one combination you did not try is off by a hair.

Wrong: The damage from floating point shows up as a visibly wrong total. **Right:** It shows up as a failed equality test in your double-entry validator, as rounding that lands on the wrong side, and at serialisation boundaries where JSON has no decimal type.

Wrong: `round(1.15, 1)` returning 1.1 is a fault in the rounding rule. **Right:** The value actually stored is 1.14999999999999991118..., genuinely below the midpoint; the function did what it was asked and the input was never what you thought.

Wrong: Money bugs come from choosing the wrong data type. **Right:** The 42-pence gap on a run of 3,000 sales came entirely from computing the merchant fee per transaction rather than per run, which is a rounding policy nobody specified.

Wrong: How many decimal places a currency has is a fact you can look up once. **Right:** ISO 4217 rounds off the ariary and the ouguiya, processors disagree with the standard and with each other, and the smallest unit you can pay in is not always the smallest you can account in.

Wrong: Bankers' rounding is the correct rounding for money. **Right:** Council Regulation (EC) No 1103/97 requires an exact half to be rounded up, and VAT Notice 700 lets an invoice trader round total VAT down while forbidding retailers from doing so.

Wrong: Rounding each line and then adding gives the same answer as adding and then rounding. **Right:** Three items at £8.99 with VAT at 20% give £5.40 by line and £5.39 by order total, and if your invoice engine and your settlement engine disagree you will reconcile that penny for the life of the product.

Wrong: Rounding each share to the nearest penny divides a total fairly. **Right:** £5.00 three ways creates a penny from nothing and £100.00 three ways destroys one, so the parts must be forced to sum back to the whole.

Wrong: A currency whose exponent is missing can safely default to 2. **Right:** XAU has no minor unit at all, so defaulting it silently invents centigrams of gold; and XXX, meaning no currency involved, is neither a null nor an error.

9.99 Chapter summary in 20 lines

1. Money is not a number: it is a number, plus a currency, plus a rule about how finely that currency may be cut, plus a written decision about what happens between two permitted cuts.
2. The working discipline is to count in the smallest coin and convert to pounds only when printing something a human will read.
3. Binary floating point cannot hold one tenth or one hundredth, because every value it represents is a whole number multiplied by a power of two.
4. Ten ten-pence coins added as pounds make 0.9999999999999999, so a till checking for exactly £1.00 never fires.
5. The arithmetic is not approximate but exact in the wrong base, which is worse, because a deterministic residue survives every test you thought to write.
6. The damage arrives as failed equality tests, as rounding that lands on the wrong side, and at API boundaries where an exact amount is parsed into a binary double.
7. Not every currency has a hundred minor units: the yen has none in use and the Kuwaiti dinar has a thousand fils, so the exponent must be looked up rather than assumed.
8. ISO 4217 gives every currency a three-letter code, a three-digit numeric code and a minor unit, and its List Three of withdrawn codes exists because codes are not permanent.
9. Of 178 active alphabetic codes, 139 have a minor unit of two, seventeen have none, seven have three, two have four, and thirteen have no minor unit at all.
10. The minor unit is a property of a currency within a particular scheme's rules rather than of the currency alone, so derive it from a table you control and version.

11. The canonical representation is an integer count of minor units with an explicit currency, and there is no decimal point anywhere in storage.
12. Where an integer will not stretch, use a decimal type whose radix is ten, and check that no object-relational mapper quietly turns the column back into a double.
13. The named rounding modes are few, and the bias argument for ties-to-even is genuine, but statistics do not override statute.
14. Euro conversion is fixed in law: six significant figures, no inverse rates, mandatory triangulation through the euro, and exact halves rounded up.
15. United Kingdom VAT permits an invoice trader to round total VAT down and forbids a retailer from doing so, which makes the rounding mode a per-calculation configuration item with a legal basis.
16. Rounding each component and then adding is not the same as adding and then rounding, and one penny of disagreement between two subsystems is permanent.
17. Whatever the mode, the parts of any divided total must sum exactly to the total in integer minor units, enforced as an assertion rather than hoped for.
18. The largest remainder method achieves that: take the floor of each exact share, then hand the residual out one minor unit at a time to the largest fractional remainders, with a deterministic tie-break.
19. An amount never travels without its currency, an account is denominated in exactly one currency, every conversion records source, target, rate, direction, timestamp and rate source, and a published reference rate is not a rate you can deal on.
20. A ledger must be exact rather than merely accurate, because its whole function is that two parties who do not trust each other can compute the same answer from it — and three pence you cannot explain is a defect in the only property a ledger has.

Sources used: SIX Financial Information — ISO 4217 currency code List One, published 1 January 2026; ISO — ISO 4217 and ISO 4217:2015. EUR-Lex — Council Regulation (EC) No 1103/97, Articles 4 and 5. HM Revenue and Customs — VAT Notice 700, paragraphs 17.5, 17.5.1, 17.5.2 and 17.6. European Central Bank — Euro foreign exchange reference rates. IEEE 754 Standard for Floating-Point Arithmetic. PostgreSQL documentation, section 8.1 Numeric Types. Microsoft .NET API documentation, System.Decimal. Stripe — Supported currencies. ISO 20022 — ActiveCurrencyAndAmount schema restriction and Message Definition Report rule on fractional digits, with XMLdata's validation notes. SWIFT field 32A network validated rules, via IBM Financial Transaction Manager documentation. Bacs Standard 18 record layout as reproduced in submitter documentation. The Vancouver Stock Exchange restatement of 25 to 28 November 1983, reported in The Wall Street Journal and catalogued in the numerical-analysis collections at TU Delft and the University of Texas at Austin. Floating-point values here were computed directly in IEEE 754 binary64 and Python's decimal module.

Reading a Statement

10.0 What this chapter gives you

1. You will be able to read a bank statement line by line and say, for each line, which system wrote it, when, on whose authority, and what could still change it.
2. You will be able to explain why Priya's booked balance was £47.30 too high and her available balance £52.70 too low on the same screen at the same moment, with no bug anywhere.
3. You will be able to separate booking date, value date and the underlying transaction's own date, and say why a report that chooses between them implicitly books entries into the adjacent day.
4. You will be able to explain why a pending card item is not an entry at all, why it can vanish leaving no reversal line, and why matching it to a posted entry is inference rather than lookup.
5. You will be able to check the two arithmetic proofs on any statement — opening plus movements equals closing, and unbroken statement sequence numbers — before accepting the file.
6. You will be able to choose between camt.052, camt.053, MT940 and BAI2 for a controlled reconciliation and defend the choice.
7. You will be able to read the ISO 20022 balance type codes and say whether a given screen is showing CLBD, CLAV or XPCD.
8. You will be able to explain why a payer-supplied reference is evidence of intent rather than a key, and design allocation that does not break when it is truncated or wrong.
9. You will be able to run a three-way reconciliation between internal ledger, processor report and bank statement, and name the characteristic failure of each pair.
10. You will be able to say why twelve failed safeguarding firms averaged a 65 per cent shortfall in client funds, and what the rules coming into force on 7 May 2026 require instead.

Everything in this volume so far has been an argument about what money is. Money is a record. A balance is a row. Every payment is two facts. A bank is a balance sheet whose liabilities you happen to call your savings. Clearing is not settlement. Netting turns twenty-seven million obligations into a few dozen. And an amount of money is an integer in minor units, never a floating-point number.

All of that is true, and none of it is in your hand.

This chapter is. Open your banking app, or take the PDF the bank sent you at the start of the month, and put it in front of you. Every abstraction in the preceding nine chapters is printed on it somewhere, usually in a column you have never looked at. By the end of this chapter you will be able to read the document line by line and say, for each line, which system wrote it, when, on whose authority, and what could still change it.

Then we will do the harder thing, which is to prove that your record of these events and the bank's record of these events are the same record. That exercise is called reconciliation, and it is where young

payments companies actually die — not in a dramatic outage, but quietly, over eighteen months, because nobody could ever say with certainty how much customer money the firm was holding.

The plain version

Here is the first thing to understand, and it is not obvious.

Your bank statement is not your money. It is not even a picture of your money. It is a letter from your bank saying what it believes it owes you, written at a particular moment, by a system that was not finished processing.

Think of a school report. The report says what the teacher thinks of your work. It is written on a Thursday, printed on a Friday, and posted on a Monday. If you handed in a brilliant essay on the Friday afternoon, it is not in the report, and the report is not wrong — it was written before the essay arrived. The report is a true description of a moment that has already passed.

A statement is the same. It is written at a moment. The world kept going.

■ Three dates, not one

Every line on your statement has at least three dates attached to it, and most statements only show you one or two of them.

Imagine you write an essay on Friday. You hand it in on Monday. The teacher marks it as though it had arrived on Sunday night, because Sunday night was the deadline. Three dates:

The date it happened. Friday, when you wrote it.

The date it was recorded. Monday, when the teacher took it off the pile and wrote your name in the mark book.

The date it counts from. Sunday, because that is the date the rules say applies.

Money works exactly like this. You buy a coffee on Saturday. The shop tells its bank about it on Monday. Your bank writes it into your account on Tuesday. And your bank may decide that, for the purposes of working out interest, the money left you on Saturday.

Three dates for one coffee. Most people have never noticed there is more than one, and then one day they see a purchase they made on the first of the month sitting on a line dated the fifth, and think the bank has made a mistake. It has not. It has shown a different one of the three dates than the one they were thinking of.

■ Two balances, not one

The second thing that is not obvious is that you have two balances at the same time, and they are usually different.

One is the balance the bank has written down: the sum of everything it has actually recorded. Call it the written-down balance.

The other is the balance you are allowed to spend: the written-down balance, minus anything the bank has promised somebody else on your behalf but has not yet recorded. Call it the spendable balance.

Here is the classic case. You drive onto a forecourt and use the pay-at-pump. The pump has no idea how much fuel you are about to take. It could be five pounds, it could be ninety. So before it releases a drop, it asks your bank a question: if this person tries to spend one hundred pounds, will you cover it?

Your bank says yes, and — this is the important part — it sets one hundred pounds aside. Not spent. Set aside.

You put in forty-seven pounds and thirty pence of fuel and drive away. For the next couple of days, your written-down balance has not moved at all, because nothing has been recorded yet. But your spendable balance is a hundred pounds lower, because a hundred pounds is being held back for a bill that has not arrived. Then the real bill arrives for forty-seven pounds thirty, the bank records it, the written-down balance finally drops, and the hundred-pound reservation is released.

Notice what happened in between. For two or three days, your written-down balance was too high — you had spent forty-seven thirty and it did not show. And your spendable balance was too low — a hundred pounds was reserved when only forty-seven thirty was owed. Both balances were wrong at the same time, in opposite directions, and neither of them was a bug.

■ A month in one account

Let us make this concrete. Here is a statement. Priya has a current account. This is March.

Date	Description	Out	In	Balance
1 Mar	Opening balance			1,842.60
3 Mar	Direct debit, Thames Water	34.00		1,808.60
5 Mar	Card, Caffè Nero, 1 Mar	3.40		1,805.20
9 Mar	Payment in, J Raman, ref BIRTHDAY		50.00	1,855.20
12 Mar	Card, Sainsbury's, 10 Mar	62.15		1,793.05
18 Mar	Card, Esso, 15 Mar	47.30		1,745.75
20 Mar	Refund, Zara, 14 Feb		28.99	1,774.74
25 Mar	Salary, KedByte Technologies		2,150.00	3,924.74
26 Mar	Standing order, rent	1,100.00		2,824.74
31 Mar	Interest		0.42	2,825.16
31 Mar	Closing balance			2,825.16

Ten lines. Now read them properly.

The coffee on line three was bought on 1 March and written down on 5 March. Four days. Nothing went wrong; that is simply how long it took the coffee shop's bank to send the bill to Priya's bank.

The fifty pounds on line four came from Priya's brother. The word BIRTHDAY is on the statement because her brother typed it into his banking app. No bank wrote it, no bank checked it, and no bank knows what it means. If he had typed nothing, the line would say nothing. If he had typed the wrong thing, the line would say the wrong thing and the money would still have arrived, because the money is addressed by the account number, not by the note.

The Esso line is the forecourt. Purchase 15 March, recorded 18 March. On 16 March, Priya's app would have shown a written-down balance of £1,793.05 and a spendable balance of £1,693.05, because £100 was reserved. Her real position — everything she had actually spent — was £1,745.75. Her written-down balance was £47.30 too high. Her spendable balance was £52.70 too low. On one screen. At the same time.

The refund on line seven is for something bought on 14 February. It is not the February purchase being undone. It is a completely separate payment, going the other way, made by the shop to Priya. This is why refunds are slow when purchases feel instant: nothing is being erased, a second thing is being done.

And the whole statement has one property that matters more than any single line: the opening balance plus everything in, minus everything out, equals the closing balance. £1,842.60 plus £2,229.41 in, minus £1,246.85 out, is £2,825.16. If that sum does not come out, something is missing from the page.

■ Reconciliation, without the word

Now the second half of the chapter, in plain terms.

Suppose you run a market stall. At the end of Saturday you have a notebook: every sale, what it was, how much. Your card machine's provider also has a record of every sale. And your bank has a record of what landed in your account.

Three records of the same day. Reconciliation is the work of proving all three describe the same Saturday, and explaining — line by line, not in total — every place they differ.

It sounds trivial. It is not. Your notebook says you took £1,240 on Saturday. Your bank account shows £1,213.90 arriving on Tuesday. Those numbers are different and both are correct: the card company kept £26.10 in fees, and it paid you two days later. If your test is “does the bank number equal my number”, you fail a test you should pass. And the reverse also happens: your notebook says £1,240, the bank says £1,240, and you are missing a sale, because two errors cancelled out. The total matched and the record is wrong.

That is the whole discipline in one paragraph. Matching totals proves nothing. Reconciliation means matching individual events, one by one, to individual events, and having a defensible explanation for every single one that does not match. Not most of them. All of them.

Do this every day and a mistake is a small puzzle. Do it every month and a mistake is an investigation. Do it never — and this is the thing that kills companies — and one day somebody asks how much of the money in your bank account belongs to your customers, and the honest answer is that nobody knows.

Where the plain version stops being true

The report-card picture is a good one, and the three dates and two balances are real. Four things it hides matter enough to correct.

■ There is no such thing as “the statement”

The plain version implies one document. In practice, at least three different artefacts are called your statement, they are produced by different systems at different moments, and they routinely disagree.

The screen in the mobile app is usually a merged view: recorded entries from the ledger, plus card authorisations that have not been recorded at all, presented in one list so it looks like a single sequence of events. The PDF or paper statement is normally the ledger alone. The machine-readable file that a business downloads — an ISO 20022 camt.053, a SWIFT MT940, a BAI2 file — is the ledger alone as well, but it is generated at a fixed cut-off and delivered on a schedule, so it can be hours behind the app and days behind reality.

When a customer says “the app shows a payment but the statement does not”, nobody has made an error. Two different systems were asked two different questions. This is also why software consuming bank data must be explicit about which artefact it is reading. Building a customer balance from the app's merged view and building it from the camt.053 will give different numbers on most days, and only one of them is a balance in any accounting sense.

■ **“Pending” is not a status; it is the absence of an entry**

This is the correction that matters most for anyone building systems.

In the plain version, the pending coffee “becomes” the posted coffee, as though a single record changed state from grey to black. That is not what happens in a card transaction. The pending item is not an entry in the ledger at all. It is a memorandum — an amount held against your available funds, recorded in an authorisation store, with no double entry behind it and no place in the sum that produces your balance. The posted item, days later, is a genuinely new ledger entry created from a completely separate message that arrived through the clearing system.

Three consequences follow, and each of them breaks naive software.

The amounts need not match. The forecourt held £100 and the entry was £47.30. A restaurant may authorise the bill and clear the bill plus a tip. A hotel may authorise nightly and clear once.

The pending item can vanish with no trace whatsoever. If the merchant never submits the transaction for clearing, the hold simply expires. Nothing appears on the statement. There is no reversal line, because there was never a line.

And there is usually no field on your statement that links the two. The authorisation and the clearing record carry identifiers that can be matched by the issuer’s systems, but those identifiers are frequently not surfaced to the customer, and are not present in a camt.053 in any standardised place. Matching a pending item to a posted item from the outside is inference, not lookup.

■ **Value date is not “when I can spend it”**

The plain version compressed three ideas into “the date it counts from”. They come apart.

The value date is an interest and accounting concept: the date from which the amount is treated as being at the account holder’s disposal for the purpose of calculating interest. It is not a promise about availability of funds, and in the United Kingdom the two are regulated separately. The Payment Services Regulations 2017 require, at regulation 89, that “the credit value date for the payee’s payment account must be no later than the business day on which the amount of the payment transaction is credited to the account of the payee’s payment service provider”, and, separately, that the payee’s provider “must ensure that the amount of the payment transaction is at the payee’s disposal immediately after that amount has been credited to that payment service provider’s account”. Two obligations, one about interest, one about access.

Value dates can also point backwards. An entry booked today can carry a value date of last Tuesday, which means interest is recalculated as though the money had been there since last Tuesday. Back-valuation is normal, particularly for corrections, and it means that yesterday’s interest figure was not final when it was printed.

Finally, the money-can-still-go-back problem. A recorded credit is not necessarily a final credit. Cheques paid in under the United Kingdom’s Image Clearing System, live since 2019, are typically available to withdraw by 23:59 on the next weekday, but availability and finality are different questions and a cheque can still be returned. Direct debit collections can be represented or refunded. Payments sent in error can be pursued through indemnity and recall processes. So a balance made entirely of booked entries can still overstate what you will keep.

■ Reconciliation is not matching, it is explaining

The market stall gave the right instinct and the wrong scale. Reconciliation in a real payments business is not a comparison of two lists; it is a continuous process of accounting for differences, most of which are expected.

Differences fall into recognisable species. Timing differences, where both records are right and one is early. Netting differences, where one side reports gross and the other net of fees. Granularity differences, where one side records a thousand sales and the other records one settlement. Sign and direction differences, especially around refunds and chargebacks. Currency differences, where the amount instructed and the amount settled are not in the same currency and neither party rounded the same way.

None of those is an error. Each still has to be identified, categorised and proved, because the only way to find the one real break — the missing payment, the double credit, the customer whose money went to the wrong internal account — is to have rules that eliminate every harmless difference and leave the harmful one standing alone.

The technical version

■ What the document actually is

An account is a ledger: an ordered, append-only sequence of postings, each of which is one half of a double entry, as chapters three and four established. A statement is not that ledger. A statement is a **report over an interval of that ledger**, and it has a fixed logical shape regardless of format:

An identification of the account and the account servicing institution. A statement identifier and a sequence number. A period, with a start and an end. One or more balances, each labelled with a type and a date. Zero or more entries. And, implicitly, an arithmetic proof: opening booked balance, plus the signed sum of the entries in the period, equals the closing booked balance.

That proof is the single most useful property of the document and it is the first thing an integration should check. If it does not hold, the statement is either truncated, mis-parsed, or reporting a balance type you have misidentified.

The sequence number is the second most useful property. In ISO 20022 it is the `Elct rncSeqNb` element inside `Stmt`; in SWIFT MT940 it is field `:28C:`, the statement number and sequence number. Consecutive statements must have consecutive numbers. A gap means a statement was never delivered, and a firm reconciling only the statements it received will reconcile cleanly to a false position. BAI2 provides no equivalent statement sequence number, which is one of several reasons it is a weaker basis for controlled reconciliation than `camt.053`; it relies instead on control totals in its account, group and file trailer records.

■ The formats, and which one you are actually reading

Format	Standard body	Content	Typical timing
camt.052	ISO 20022	Bank-to-customer account report, intraday	Periodically through the day
camt.053	ISO 20022	Bank-to-customer statement, end of day	Once per business day, after cut-off

Format	Standard body	Content	Typical timing
camt.054	ISO 20022	Bank-to-customer debit/credit notification	Per event or per batch
MT940	SWIFT MT	Customer statement message	End of day
MT942	SWIFT MT	Interim transaction report	Intraday
BAI2	Banking Administration Institute	Cash management balance reporting	Previous-day, and intraday variants
PDF / CSV / OFX	None binding	Human or spreadsheet consumption	Varies

The distinction between camt.052 and camt.053 is the one that causes the most trouble in practice. camt.053 is the official end-of-day statement for the previous business day and its entries are booked. camt.052 is an intraday report and may contain items that are not final. Reconciling a firm's books against camt.052 will produce a position that changes underneath you; reconciling against camt.053 produces a position that does not.

An MT940 message is a tagged text format. The fields that matter are :20: transaction reference number, :25: account identification, :28C: statement number and sequence number, :60a: opening balance, :61: statement line, :86: information to account owner, :62a: closing balance, :64: closing available balance and :65: forward available balance.

Field :61: carries the format 6!n[4!n]2a[1!a]15d1!a3!c16x[/16x] with an optional [34x] supplementary details line. Read left to right, that is: a six-digit value date as YYYYMMDD; an optional four-digit entry date as MMDD; a debit/credit mark of C, D, RC or RD; an optional funds code; the amount; a transaction type identification code of one letter and three characters, where S denotes a SWIFT transfer and the three characters are the message type, and N denotes a non-SWIFT transfer with a mnemonic such as TRF or CHG; a reference for the account owner; optionally, after a double slash, the account servicing institution's own reference; and optionally a supplementary details line.

Look at that entry date again. MMDD. Four digits. **There is no year.** A parser must infer the year from context, and every year in the last week of December and the first week of January, somewhere in the world, a reconciliation engine books an entry twelve months out because it inferred wrongly. This is not a hypothetical class of bug; it is a scheduled one.

BAI2 is a comma-delimited record format used mainly by North American banks and by global banks for cash management reporting. Its records are 01 file header, 02 group header, 03 account identifier, 16 transaction detail, 88 continuation, 49 account trailer, 98 group trailer and 99 file trailer. Balances appear as type codes inside the 03 record: 010 opening ledger balance, 015 closing ledger balance, 040 opening available balance and 045 closing available balance, with 100 and 400 carrying total credits and total debits.

■ The dates, exactly

ISO 20022 gives an entry three date elements and they are not interchangeable.

BookgDt is the booking date: the date the entry was posted to the account. This is the date that determines which statement the entry appears on and which day's closing balance it affects.

ValDt is the value date: the date from which the amount is taken to be at the account holder's disposal for interest and accounting purposes.

Dt inside the transaction details, and the various dates inside RltdDts, carry the underlying transaction's own dates — for a card entry, the date of purchase; for a direct debit, the requested collection date.

The relationship between them is institution-specific and rail-specific, and the only safe assumption is that they are three separate facts. For a Bacs direct credit, booking date and value date will typically both fall on entry day. For a card purchase, the transaction date precedes the booking date by anything from hours to a week. For a back-valued correction, the value date precedes the booking date deliberately.

Store all three, plus the timezone in which the booking date was determined. A very large proportion of reconciliation breaks that appear to be missing transactions are, on inspection, entries booked into the adjacent day because two systems disagreed about when the day ended. Business date is not a derived quantity. It is a value the institution declares, and it changes at a cut-off, not at midnight.

■ Pending, posted, and the balance type code set

ISO 20022 gives an entry a status element, Sts, from a code set whose members are B00K (booked, posted to the account), PDNG (pending, not yet posted), INFO (informational, no financial impact) and, in current versions, FUTR (future). Market practice for an end-of-cycle camt.053 requires B00K; PDNG and INFO entries belong in intraday reporting. If you are receiving PDNG entries in a document you believe to be an end-of-day statement, you are reading a camt.052 or a bank's proprietary variant.

Balances are where the standard is most useful and most ignored. A Bal element carries a type code, and the code set is precise:

Code	Meaning
OPBD	Opening booked balance. Always equals the previous report's closing booked balance
CLBD	Closing booked balance. Opening booked balance plus all entries booked in the period
OPAV	Opening balance of money at the account owner's disposal on the date specified
CLAV	Closing balance of money at the account owner's disposal on the date specified
ITBD	Booked balance calculated during the business day, at the time specified, subject to change
ITAV	Available balance calculated during the business day, at the time specified, subject to change
FWAV	Forward available balance at the disposal of the account owner on the date specified
PRCD	Balance at the previously closed reporting period
XPCD	Expected balance: booked entries plus pending items known at the time of calculation, projecting the end-of-day balance if everything is booked and nothing further is posted
INFO	Balance for information only

Two of these deserve emphasis. XPCD is the standard's name for the number a customer thinks of as "my balance" in a mobile app — booked plus pending — and it is explicitly labelled as a projection, not a position. And FWAV is the mechanism by which a bank tells you that money is arriving on a future

date under a value-dating arrangement; it is a forward-looking figure and belongs in cash forecasting, not in a ledger.

The available balance itself is not defined by any standard, because it is a commercial construct. In broad terms:

Available balance equals the booked balance, minus authorisation holds not yet cleared, minus credits booked but not yet available under the institution's funds availability policy, plus any agreed overdraft or credit facility, plus any same-day credits the institution has chosen to release early.

Every term after the first varies by institution, by product and sometimes by customer. This is why two accounts holding identical entries can show different available balances at the same bank, and why "available balance" is never a safe input to an accounting process.

■ Authorisation holds

A card authorisation is a message, not an entry. Its economic effect on the cardholder is a reduction in available funds; its accounting effect on the issuer's ledger is nothing at all until a clearing record arrives.

The scheme rules control how long that gap may be, and the answer depends on how the merchant characterised the authorisation. A **final authorisation** asserts that the amount is exactly the amount that will be captured. An **estimated** or **pre-authorisation** asserts that the amount is a good-faith estimate, which is what a fuel pump, a hotel and a car hire desk necessarily send.

Timeframes published by the acquirer Worldline for the European region, as of writing, illustrate the shape of the rules:

Scheme	Authorisation type	Validity period
Visa	Final authorisation	10 days
Visa	Pre-authorisation, lodging, cruise line, vehicle rental	30 days
Visa	Pre-authorisation, other merchant categories	10 days
Mastercard	Final authorisation	3 days, with reversal or clearing still possible up to 7 calendar days subject to fees
Mastercard	Pre-authorisation	30 days
Diners	Final authorisation / pre-authorisation	7 days / 30 days
Discover	Final authorisation / pre-authorisation	10 days / 30 days

Visa's own acceptance documentation states the merchant-side obligation in the same terms: an authorisation flagged as final "must be submitted for capture within seven calendar days of its request", while a pre-authorisation must be submitted for capture within 30 calendar days, and an authorisation that is not captured must be reversed.

Three practical points follow.

First, the hold on the cardholder's available balance and the scheme's authorisation validity period are related but not identical. An issuer may release a hold on its own schedule, typically shorter than the scheme maximum, precisely because customers complain. Once released, the funds are spendable

again — and the clearing record may still arrive afterwards, producing an entry against a balance that no longer expects it.

Second, an estimated authorisation can be adjusted. Incremental authorisations increase the reserved amount as a hotel stay lengthens; partial reversals reduce it when the final amount is known. Whether an issuer reflects those adjustments promptly in the customer-visible available balance is an implementation matter, not a rule.

Third, the amount that clears may legitimately differ from the amount authorised, within tolerances that vary by merchant category. Software that treats such a mismatch as an anomaly will generate an unmanageable volume of false positives against fuel, hospitality and transport merchants.

■ Why a balance can be wrong in two directions at once

Return to Priya's account on the afternoon of 16 March.

Measure	Amount	Why
Booked balance (CLBD as at 15 March)	1,793.05	Every posted entry, correctly summed
Authorisation hold	100.00	Estimated authorisation at the pump
Available balance	1,693.05	Booked minus the hold
Economic position	1,745.75	What she has actually committed to spend

The booked balance overstates her position by £47.30, because she has bought fuel that the ledger does not yet know about. The available balance understates it by £52.70, because £100 is reserved against a £47.30 obligation. There is no number on the screen that is correct, and every number on the screen is produced correctly.

Generalise it. A balance overstates in three situations: authorised but uncleared card spending; credits booked that remain revocable, including cheques within the return window, direct debits within the indemnity window and credit transfers subject to recall; and fees and interest accrued but not yet applied. It understates in three others: holds that exceed the eventual clearing amount; holds that persist after the corresponding clearing record has already posted, producing a temporary double count; and incoming credits received by the institution but not yet released under its availability policy.

Both categories are usually live at once on a busy account. The correct engineering response is not to pick one balance and trust it. It is to hold, separately, a booked position derived from entries, a set of outstanding holds with their own lifecycle, and a set of credits with a finality date. Any single number you present to a user is a summary of those three, and you should know which one you are presenting.

■ The reference field, and who wrote it

The most misunderstood column on any statement is the narrative, and the misunderstanding is this: people believe the bank wrote it. Almost none of it is written by the bank.

Different rails give the payer different amounts of room, and the room is small.

For **Bacs**, using the Standard 18 file layout described in the *Bacs Electronic Funds Transfer, File Structures* document, the payment instruction record carries a service user's reference in field 10, and it is **18 characters**. The originating account name is field 9, the destination account name field 11, and the

amount in pence field 8, itself limited to 11 digits. Eighteen characters is the whole channel through which a payroll bureau, a utility or an insurer can tell the receiving party what a payment is for.

For **Faster Payments**, which uses an ISO 8583 message, the Open Banking Read/Write specification's field mapping shows field 62 END TO END REFERENCE at 31 characters, field 120 REFERENCE INFORMATION at 18 characters and field 121 REMITTANCE INFORMATION at 140 characters. Note the asymmetry: 140 characters of remittance information exist in the message, but the field most receiving systems key on is the 18-character one, and if the creditor account carries a secondary identification such as a building society roll number, that secondary identification takes field 120 and the payer's reference does not get it.

For **CHAPS**, carried as an MT103, the beneficiary reference is field 70, four lines of 35 characters, with structured tags such as /ROC/ for the original customer reference.

In a camt.053, the references that reach the statement live in NtryDtls/TxDtls/Refs and comprise MsgId, PmtInfId, EndToEndId, TxId and, for direct debits, MndtId. Remittance data lives in RmtInf, as Ustrd free text of up to 140 characters per occurrence or as Strd structured remittance. The entry itself carries AcctSvcrRef, the account servicing institution's own reference, and AddtlNtryInf, free text written by the bank.

So the ownership of each string is as follows:

String	Written by	Reliability
EndToEndId	The originating party's system	High for payments you sent; unvalidated for payments you received
RmtInf/Ustrd	The originating party, usually a human typing	Free text; no validation of any kind
Bacs field 10 reference	The originating service user's file	18 characters, uppercase, restricted character set
AcctSvcrRef	The account servicing bank	The bank's own key; stable within that bank
AddtlNtryInf	The account servicing bank	Descriptive; format not guaranteed stable across releases
Card merchant descriptor	The acquirer, from the merchant's registered acceptor name and location	Frequently a trading entity's legal name, not its shop sign

The consequence for anybody building a payment system is a single hard rule. **A payer-supplied reference is evidence of intent, not a key.** It can be truncated, uppercased, stripped of punctuation, rewritten by an intermediary, mistyped by the payer or omitted entirely, and the payment will still arrive, because the payment is addressed by account identifier. Systems that allocate incoming funds by exact-matching a reference will misallocate. Systems that allocate by exact match, then fall back to fuzzy match, then fall back to a human queue with the amount, date, payer name and remitting sort code available for judgement, will misallocate far less.

The corresponding rule for outbound payments is to generate the reference yourself, make it structurally checkable — a prefix plus a check character is enough — and keep it within the tightest field length on any rail you use, which in the United Kingdom means 18 characters, not 35 and not 140.

■ Classification: bank transaction codes

Each entry may carry a BkTxCd, a hierarchical classification with a domain, a family and a sub-family, drawn from the ISO 20022 external code set, plus an optional proprietary code.

A received SEPA credit transfer, for example, is PMNT / RCDT / ESCT: payments domain, received credit transfer family, SEPA credit transfer sub-family. The family codes within the payments domain distinguish received from issued and credit transfers from direct debits — RCDT, ICDT, RDDT, IDDT — with MCOP and MDOP for miscellaneous credit and debit operations such as fees and interest.

Two warnings. Banks vary enormously in how carefully they populate this structure, and a great deal of real-world classification lives in the proprietary code rather than the standard one. And the American equivalent, the BAI2 numeric type code, has a different and larger vocabulary, so a firm consuming both formats needs a mapping layer and should expect that mapping to be incomplete.

■ Reconciliation, properly

Reconciliation is the process of proving that two or more independently maintained records of the same set of events agree, and identifying, categorising and resolving every difference between them.

In a payments business there are rarely two records. There are usually three, and the discipline is described as three-way reconciliation:

The **internal ledger**: what your own system says each customer is owed and what your own system says you instructed.

The **scheme or processor report**: what the acquirer, processor or scheme says happened — authorisations, captures, refunds, chargebacks, fees, and the settlement amounts derived from them.

The **bank statement**: what money actually arrived in or left your bank accounts.

Each pair must tie, and each fails in a characteristic way.

Internal ledger against processor report fails on events the processor knows about and you do not: chargebacks raised, fees levied, and transactions your API call timed out on but which the processor nonetheless completed. That last case produces duplicate charges, and it is why every outbound instruction must carry an idempotency key that survives a retry.

Processor report against bank statement fails on netting. The processor reports gross transactions; the bank receives a net settlement after fees, refunds, chargebacks and reserve movements. If your matching logic looks for a bank credit equal to the sum of the day's sales, it will never find one. The correct approach is to reconstruct the settlement: sum the gross items in the batch, subtract each reported deduction, and prove the result equals the credit on the statement to the penny. Then, separately, prove every gross item in the batch exists in your ledger.

Internal ledger against bank statement fails on timing and on entries you never created: bank charges, interest, returned direct debits, and credits from parties you were not expecting.

A working process has five properties. It runs daily, not monthly. It matches at item level, never at total level alone. It classifies every break into a named category with an owner. It ages breaks and escalates them, so that a break surviving three days becomes somebody's problem by policy rather than by luck. And it treats the arithmetic proofs described earlier — opening plus movements equals closing, sequence numbers unbroken, BAI2 control totals consistent — as automated preconditions, not as things a human notices.

■ Why this is where firms die

This is not a metaphor and there are figures.

For investment firms holding client money, the FCA's client assets rules are explicit. CASS 7.15.15R requires an internal client money reconciliation "each business day", performed on the previous business day's closing records. CASS 7.15.22R requires an external reconciliation against third-party records "as regularly as is necessary but without allowing more than one month to pass" between reconciliations. And CASS 7.15.29R requires that where an internal reconciliation reveals a shortfall, the firm determines the reason and pays in the difference "by the close of business on the day that the reconciliation is performed". A firm that cannot reconcile daily cannot fund a shortfall daily, and a firm that cannot fund a shortfall daily is in breach on the day it first fails.

For payments and e-money firms, the rules were weaker for years, and the FCA's own evidence on what that produced is stark. In its consultation paper CP24/20, published in September 2024, the regulator recorded that between the first quarter of 2018 and the second quarter of 2023, twelve payments firms that safeguarded relevant funds became insolvent, and that for those firms there was an average shortfall of 65 per cent in funds owed to clients. Two thirds of customer money, gone, on average, across twelve failures. The same paper identified the causes in terms that are almost a description of this chapter: "a lack of documented processes for consistently identifying which funds must be safeguarded" and "inadequate reconciliation procedures to ensure that the correct sums are protected on an ongoing basis". In 2023 alone the FCA opened supervisory cases relating to approximately 15 per cent of firms that safeguard.

The response is Policy Statement PS25/12, published on 7 August 2025, which introduces the interim Supplementary Regime through new Handbook chapters CASS 15, CASS 10A, SUP 3A and SUP 16.14A, together with amendments to the *Payment Services and Electronic Money — Our Approach* document. The rules come into force on **7 May 2026**. Their substance is this chapter's subject matter turned into obligations: internal and external safeguarding reconciliations at least once each reconciliation day, excluding weekends, public holidays and days when relevant foreign markets are closed; a resolution pack sufficient to return customer funds promptly on insolvency; an annual safeguarding audit by a qualified auditor, with an exemption for firms that have not been required to safeguard more than £100,000 of relevant funds over a period of at least 53 weeks; a monthly safeguarding return to the regulator; and documented due diligence on third parties holding relevant funds.

Read that list next to the failure statistic and the causal story is unambiguous. Firms did not fail because their payments went wrong. Their payments mostly worked. They failed because nobody could prove, on any given day, how much of the money in the bank belonged to customers — and by the time anyone asked the question properly, two thirds of it was not there.

Reconciliation is therefore not back-office hygiene. It is the control that converts a database of intentions into a defensible statement of obligation. It is the difference between a firm that knows what it owes and a firm that believes what it owes.

■ A short list of things to build

For anyone implementing against bank statement data, the practical conclusions of this chapter compress into seven rules.

Parse and store all three dates on every entry, plus the timezone, and never let a report choose between them implicitly.

Verify the arithmetic proof on every statement before accepting it, and verify statement sequence continuity, alarming on a gap rather than proceeding.

Key on the account servicing institution's own reference for idempotency of ingestion, and on your own generated end-to-end identifier for outbound payments. Never key on payer free text.

Model holds as a separate object with their own lifecycle, linked to entries where possible and explicitly unlinked where not.

Distinguish booked, available and projected balances in your data model with different names, and never store one in a field named after another.

Reconcile daily, at item level, three ways, with categorised and aged breaks.

And treat every credit as revocable until the rail-specific finality point has passed — a subject the remaining volumes take apart rail by rail.

Where Volume I leaves you

Trace the argument backwards from the document in your hand.

The statement is a report over a ledger. The ledger is a sequence of double entries, each recording two facts about one event, because a payment is never one fact. Those entries adjust a balance, and the balance is not a quantity of anything; it is the size of a debt a bank owes you, held as a row in the bank's records. The bank can owe you that debt because a bank is an institution whose liabilities circulate as money, and its own money — the money it uses to settle with other banks — is a different and superior kind, issued by a central bank. When your payment reaches another bank, the obligation between the two of them is established by clearing and discharged by settlement, which are different events at different times, and those obligations are compressed by netting before any central bank money moves. Every amount involved is an integer count of minor units in a currency identified by an ISO 4217 code, because it has to be exact.

That is Volume I. It is a description of what money is, and every part of it happened inside one institution's records, or between an institution and the central bank that stands above it.

The rest of this book is about what happens **between** institutions, and it starts with the most familiar object in the whole system.

There is a piece of plastic in your pocket that weighs about five grams. It carries, in several different physical forms, one piece of information: which account to charge. It talks to a terminal it has never met, through a network owned by neither of them, using a message format designed in the 1980s and cryptography designed in the 1990s, and it does this several hundred billion times a year without anybody thinking about it.

Volume II is about that card, and it begins by taking it apart.

10.98 Common wrong ideas

Wrong: A statement is a picture of your money. **Right:** It is a letter from your bank saying what it believes it owes you, written at a particular moment by a system that was not finished processing.

Wrong: There is one document called “the statement”. **Right:** At least three artefacts carry that name — the app’s merged view, the PDF, and the machine-readable file at a fixed cut-off — produced by different systems at different moments, and they routinely disagree.

Wrong: A pending item becomes the posted item, changing state from grey to black. **Right:** The pending item is a memorandum in an authorisation store with no double entry behind it; the posted item is a genuinely new ledger entry created from a separate clearing message.

Wrong: A cleared amount that differs from the authorised amount is an anomaly worth alerting on. **Right:** The forecourt held £100 against a £47.30 entry, and treating such differences as anomalies produces an unmanageable volume of false positives against fuel, hospitality and transport merchants.

Wrong: If a hold disappears there will be a reversal line to reconcile it against. **Right:** If the merchant never submits the transaction for clearing the hold simply expires, and nothing appears on the statement, because there was never a line.

Wrong: The value date tells you when you can spend the money. **Right:** It is an interest and accounting concept, regulated separately from availability under regulation 89, and it can point backwards, so yesterday’s interest figure was not final when it was printed.

Wrong: A booked credit is money you will keep. **Right:** Cheques can be returned, direct debits represented or refunded and credit transfers recalled, so a balance made entirely of booked entries can still overstate what you will keep.

Wrong: The narrative on a statement was written by the bank. **Right:** Almost none of it was; the end-to-end identifier and the remittance text come from the originating party, and a payment arrives on its account identifier whatever the reference says.

Wrong: Reconciliation means proving the two numbers agree. **Right:** Matching totals proves nothing, because two errors cancel out; the discipline is matching individual events and having a defensible explanation for every single difference.

Wrong: The bank credit should equal the day’s sales. **Right:** The processor reports gross and the bank receives a net settlement after fees, refunds, chargebacks and reserve movements, so logic that looks for that credit will never find one.

10.99 Chapter summary in 20 lines

1. A statement is not your money and not a picture of it; it is a report over an interval of a ledger, written at a moment by a system that was still processing.
2. Every entry carries at least three dates: when the transaction happened, when it was booked, and the date it counts from for interest.
3. ISO 20022 names them BookgDt, ValDt and the transaction’s own dates, they are three separate facts, and all three should be stored with the timezone.
4. Business date is not a derived quantity but a value the institution declares, and it changes at a cut-off rather than at midnight.
5. You have two balances at once: the booked balance the bank has recorded, and the available balance you are allowed to spend.

6. On the afternoon of 16 March Priya's booked balance overstated her position by £47.30 and her available balance understated it by £52.70, and every number on the screen was produced correctly.
7. The most useful property of any statement is its arithmetic proof: opening booked balance plus the signed entries in the period equals the closing booked balance.
8. The second most useful is the sequence number, because a gap means a statement was never delivered and a firm will otherwise reconcile cleanly to a false position.
9. There is no single statement: the app shows a merged view, the PDF and the machine-readable file show the ledger, and the file is generated at a cut-off and can be days behind reality.
10. camt.053 is the end-of-day booked statement and camt.052 is intraday and not final, so reconciling against the latter gives a position that changes underneath you.
11. MT940's field :61: carries an entry date as MMDD with no year, which makes the turn of the calendar a scheduled bug rather than a hypothetical one.
12. A pending card item is not a ledger entry but a hold against available funds, and the amounts need not match, the item can vanish silently, and no standardised field links the two.
13. Authorisation validity runs from three days to thirty depending on scheme and on whether the merchant flagged the authorisation final or estimated, and issuers release holds on their own shorter schedules.
14. The ISO 20022 balance type codes are precise, and XPCD — booked plus pending — is the number a mobile app calls "my balance" and the standard explicitly calls a projection.
15. Available balance is defined by no standard because it is a commercial construct, which is why it is never a safe input to an accounting process.
16. A balance can overstate and understate at once, so hold booked entries, outstanding holds and credits with finality dates as three separate things and know which one you are presenting.
17. Almost none of the narrative is written by the bank, and the tightest United Kingdom reference field is eighteen characters, so a payer reference is evidence of intent and never a key.
18. Reconciliation is not matching but explaining: timing, netting, granularity, sign and currency differences are all expected, and rules must eliminate every harmless one so the harmful one stands alone.
19. In a payments business it is three-way — internal ledger, processor report and bank statement — run daily, at item level, with breaks categorised, owned, aged and escalated.
20. Between the first quarter of 2018 and the second quarter of 2023 twelve safeguarding payments firms became insolvent with an average 65 per cent shortfall in client funds, which is why reconciliation is the control that turns a database of intentions into a defensible statement of obligation.

Sources used: The Payment Services Regulations 2017, SI 2017/752, regulation 89 (legislation.gov.uk). FCA Handbook, CASS 7.15.12R, 7.15.13R, 7.15.15R, 7.15.20R, 7.15.22R, 7.15.27R, 7.15.29R and 7.15.31R. Financial Conduct Authority, CP24/20 "Changes to the safeguarding regime for payments and e-money firms" (September 2024) and PS25/12 (7 August 2025), together with the FCA's safeguarding requirements page for payment and e-money institutions. ISO 20022 external code sets: BalanceType (OPBD, CLBD, OPAV, CLAV, ITBD, ITAV, FWAV, PRCD, XPCD, INFO) and EntryStatus (BOOK, PDNG, INFO, FUTR); camt.052, camt.053 and camt.054 message definitions, including the SWIFT camt.053.001.02 message definition report. SWIFT MT940 and MT942 specifications, field :61: format 6!n[4!n]2a[1!a]15d1!a3!c16x[/16x][34x]. Banking Administration Institute BAI2 cash management balance reporting format, record codes 01, 02, 03, 16, 88, 49, 98 and 99 and account type codes 010, 015, 040 and 045. Open Banking Limited, Read/Write API Standard v3.1.3, "Domestic Payment Message Formats", covering the Bacs Standard 18 mapping derived from Bacs "Electronic Funds Transfer, File Structures" (PN5011), the Faster Payments ISO 8583 field mapping and the CHAPS MT103 mapping. Visa Acceptance Solutions developer documentation on the final authorisation indicator. Worldline Acquiring documentation, "Authorization validity", scheme validity periods for the European

region. Pay.UK, Image Clearing System.

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