A Practical Introduction to the LIBOR Transition

Martin Burke (with Claude)

Draft

Abstract

For half a century the world's interest-rate contracts were written on LIBOR: a *term*, *forward-looking*, *credit-bearing* rate you knew at the start of every period. The benchmarks that replaced it—SOFR, SONIA, €STR and their peers—are the opposite in every one of those respects: *overnight*, *backward-looking*, *nearly risk-free* rates known only at the *end* of a period. This note explains, for a reader who already knows risk-neutral pricing and the basics of exotic derivatives, what that inversion forces on every curve and every product. The organising idea is that the whole transition is, at heart, a *change of projection index*. We build the multi-curve picture that makes that statement precise—one discount curve, one projection curve per index—and show that swapping LIBOR for a compounded overnight rate is exactly a change of the projection curve, while the discount side had already moved years earlier, when collateral agreements pushed discounting onto the OIS curve. Along the way we meet the one genuinely new piece of mathematics the exotics desk has to absorb: a backward-looking compounded rate has a *clean* forward with *no* convexity adjustment when it is paid in the natural place, because compounding the overnight rate simply rebuilds the bank account—so the interesting optionality lives not in the vanilla coupon but in timing, in-advance conventions, and options observed *during* the accrual period. We close with the machinery that makes the switch fair and legally clean: the ISDA fixed spread adjustment, the fallback protocol, and the discount-curve "big bang".

Contents

1 Why LIBOR had to go

Start with what LIBOR actually was, because every design choice in its replacement is a reaction to one of its properties. Each morning a panel of banks answered a single question, tenor by tenor and currency by currency: *at what rate could you borrow unsecured funds, in reasonable market size, just before 11 a.m.?* The answers were trimmed and averaged, and the result was published as the London Interbank Offered Rate for that currency and tenor—an overnight rate, a one-month rate, a three-month rate, and so on. Four features of that construction matter.

- **It was a *term* rate.** A single number covered a whole three- or six-month period.
- **It was *forward-looking***—set at the *start* of the period it applied to. Borrow for three months today and you know your funding cost today.
- **It was *credit-bearing.*** It was the cost of *unsecured* lending to a bank, so it carried bank default risk plus a term-liquidity premium. When banks looked shakier, LIBOR rose relative to genuinely safe rates.
- **It was *survey-based.*** It rested on submitted quotes—expert judgement—not necessarily on transactions.

Two problems, one notorious and one fatal, followed from the last feature. The notorious one was manipulation: because submissions were judgemental and traders held positions that moved with the fixing, several banks were found to have nudged their quotes, and the 2012 scandal that followed cost billions in fines and all remaining trust in a self-reported number. The fatal problem was quieter. After 2008 the unsecured interbank term-lending market—the very thing LIBOR was meant to measure—withered to almost nothing. Banks stopped lending to each other unsecured for three months, so panel submissions became *expert judgement about a market that barely traded.* A benchmark referenced by hundreds of trillions of dollars of contracts was resting on the opinions of a shrinking panel about transactions that were no longer happening.

Regulators concluded the number could not be rehabilitated and had to be replaced by something anchored in actual, high-volume transactions. The UK Financial Conduct Authority announced in 2017 that it would no longer compel panel banks to submit past the end of 2021, setting a deadline for the market to move [1]; the Financial Stability Board had already laid out the direction of travel [2]. Most sterling, euro, yen and Swiss LIBOR settings ceased at the end of 2021; the main US dollar settings followed on 30 June 2023.

One clarification before going on: the transition retired *LIBOR*, not every IBOR. EURIBOR—the euro area's own interbank benchmark—and Japan's TIBOR were *reformed* rather than replaced, moved onto hybrid methodologies anchored in transactions where they exist, and continue to be published; euro swaps still fix against EURIBOR today. What those survivors adopted from the transition is its safety net: new contracts carry RFR-based fallbacks (€STR for EURIBOR) of exactly the kind built in §11–§12. So the machinery of this note is not only LIBOR history; it is the template now wired into every benchmark that remains.

The replacement benchmarks invert all four properties at once (Table 1). The rest of this note is, in effect, a tour of the consequences of that inversion, in three movements. First, the *new rates and their plumbing*: what a risk-free rate is (§2), how an overnight rate is compounded into a term rate (§3), the swap that trades it (§4), and why that swap's curve now does all the discounting (§5). Second, the *new coupon and its mathematics*: what it means to know a coupon only in arrears (§6–§7), the multi-curve framework that separates projection from discounting (§8), the measure-theoretic payoff (§9), and the forward-looking term rates that partially restore the old convenience (§10). Third, the *machinery of the switch itself*: the spread adjustment that makes a replacement fair (§11), the fallbacks that installed it (§12), and the discount-curve big bang that completed it (§13). Appendices collect the supporting theory—swap mechanics, change of measure, the pricing measures, convexity adjustments—and a glossary.

	LIBOR	**Risk-free rates (RFRs)**
Tenor	term (1M, 3M, 6M, ...)	overnight
Timing	forward-looking (set in advance)	backward-looking (known in arrears)
Credit	unsecured bank credit + term premium	(nearly) risk-free
Basis	survey / expert judgement	transaction-based

Table 1: The four properties that flip. Everything mechanical in the transition follows from reading this table one row at a time.

2 What a risk-free rate actually is

Each currency area chose an overnight benchmark rooted in a market that genuinely trades in size every day (Table 2). The US dollar rate, SOFR, is worth unpacking in detail, because it is the flagship of the transition and because its construction is exactly what makes it hard to manipulate and nearly free of credit risk—the two things LIBOR failed at.

SOFR is the **Secured Overnight Financing Rate**: the rate on overnight loans of cash secured against US Treasury bonds, in the *repo* market. A repo (sale-and-repurchase agreement) is a collateralised loan dressed up as two trades. The party that needs cash sells Treasuries today and simultaneously agrees to buy them back tomorrow at a slightly higher price; the difference, annualised, *is* the overnight interest rate on the loan. Economically nothing changes hands but cash for a night—the bonds are collateral—but legally the lender of cash *holds* the securities until morning. So who is on each side? The *cash lenders* are institutions with large balances to place safely overnight: money-market funds above all, plus government-sponsored entities and other corporate cash pools. The *cash borrowers* are securities dealers and banks financing their inventory of Treasuries—leveraged holders of government bonds who fund those positions one night at a time.

Two features make the loan almost riskless, and they are the point. First, it is *secured* by US Treasuries—the most liquid, highest-quality collateral there is—and it is *over*-collateralised through a *haircut*: the borrower gets a little less cash than the bonds are worth, so if it fails to repay, the lender simply keeps and sells the bonds and is made whole even after a small price move. Second, it is *overnight*: the exposure lasts until tomorrow morning, and repo positions are re-struck daily. The rate is therefore "risk-free" not because the borrower is riskless but because the loan is backed by government bonds and unwinds within a day; what it prices is the cost of *secured funding*, not bank credit.

Finally, SOFR is transaction-based by construction. Each day the New York Fed computes it as the *volume-weighted median* of the rates on the deepest segments of that market—tri-party repo (where a third-party agent holds and values the collateral), including the General Collateral Financing service, together with bilateral repo cleared through the FICC—amounting to on the order of one to two trillion dollars of actual trades. A median over an enormous pool of real transactions cannot be nudged by a submitting bank the way a survey quote could, which is precisely why regulators favoured it. Its counterparts abroad rest on the same principle: SONIA on unsecured wholesale overnight deposits, €STR on the euro-area equivalent, and SARON on Swiss repo.

Being a real market price, however, daily SOFR is *spiky* in a way a survey never was. Repo demand bunches around quarter-ends and Treasury settlement dates, and in the squeeze of 17 September 2019 SOFR printed above 5% while the Fed's policy rate sat near 2%—an overnight dislocation of some 300 basis points. This does not make it a bad benchmark; it makes it a benchmark that must be *averaged*. A single day's print is noisy, but a coupon compounds $\sim$90 of them, so a one-day 300 bp spike moves a quarterly coupon by roughly $300/90 \approx 3$ bp. That dilution is a first reason contracts reference the compounded rate over a period rather than any single fixing—the construction the next section builds.

Currency	**RFR**	**Secured?**	**Administrator**
USD	SOFR	secured (Treasury repo)	Federal Reserve Bank of New York
GBP	SONIA	unsecured	Bank of England
EUR	€STR	unsecured	European Central Bank
JPY	TONA	unsecured	Bank of Japan
CHF	SARON	secured (repo)	SIX Swiss Exchange

Table 2: The overnight risk-free rates that replaced LIBOR. Secured rates (SOFR, SARON) reference collateralised lending; the unsecured ones are still nearly risk-free because the default risk in a *single overnight* loan to a sound bank is minuscule. Either way, none carries a *term* bank-credit premium.

The common thread is that all of them are *overnight* and all are *nearly free of credit risk*. A secured rate like SOFR is protected by collateral; an unsecured overnight rate like SONIA is safe simply because lending to a solid bank until tomorrow morning carries almost no default risk. What none of them contains is the *term* credit-and-liquidity premium that a three-month unsecured loan does—and that premium was a large part of what LIBOR measured. Holding that difference in mind is the whole of §11.

The immediate practical consequence is structural. A coupon still has to cover a three-month period, but the only rate on offer is an overnight one. So a term rate now has to be *manufactured* out of an overnight rate by compounding it day by day over the period—and, because the daily fixings are not known until they happen, the manufactured rate is not known until the period is over. Those two facts—*compounding* an overnight rate into a term rate, and *knowing* the result only *in arrears*—are the mechanical heart of the transition, and the next sections build them up in turn: §3 constructs the compounded rate, §4 the swap that trades it, §5 why that swap's curve is now what everything discounts on, and §6 what it means to know a coupon only at the end.

3 Compounding an overnight rate

Fix an accrual period $[T_s, T_e]$ containing business days with overnight fixings $r_1, \dots, r_N$ and day-count fractions $\delta_1, \dots, \delta_N$ (usually ACT/360, so $\delta_i \approx 1/360$ per calendar day). Compounding *in arrears* grows one unit of currency along the realised overnight path,

$$1 + \delta_{s,e}\, R(T_s, T_e) \;=\; \prod_{i=1}^{N} \big(1 + \delta_i\, r_i\big), \qquad \delta_{s,e} = \sum_{i=1}^{N} \delta_i, \tag{1}$$

and $R(T_s, T_e)$ is the resulting annualised compounded rate. This is the ISDA and market-standard convention for SOFR and SONIA swaps and floating-rate notes. A simpler cousin, the *simple average* $\frac{1}{\delta_{s,e}} \sum_i \delta_i r_i$, appears in some cash products (and in the New York Fed's published SOFR averages); it omits the interest-on-interest term and so differs slightly.

A three-fixing example makes Equation (1) concrete, and shows the slightly non-obvious treatment of the day-count over a weekend. Take a Thursday fixing of 5.00%, a Friday fixing of 5.02%, and a Monday fixing of 4.98%, ACT/360. The Thursday and Monday rates each accrue for one day ($\delta = 1/360$), but the Friday rate accrues for *three*—Saturday and Sunday have no fixing, so Friday's carries across the weekend ($\delta = 3/360$). The compounded growth is

$$\Big(1 + \tfrac{0.0500}{360}\Big)\Big(1 + \tfrac{0.0502 \times 3}{360}\Big)\Big(1 + \tfrac{0.0498}{360}\Big) \;=\; 1.00069569,$$

and annualising over the five calendar days, $R = (1.00069569 - 1) \times 360/5 = 5.0090\%$—a day-count-weighted average of the fixings (Friday counts triple), plus an interest-on-interest sliver far below a basis point at this horizon. Nothing deeper is ever going on inside Equation (1); a real coupon just runs sixty-odd such factors instead of three.

Figure 1 works a single three-month coupon through Equation (1). The left panel is a plausible overnight path with a rate cut partway through; the right panel tracks the annualised compounded rate "known so far". Two features stand out. First, the running rate is volatile early—one or two fixings tell you little—and settles only as the period fills in, reaching its final value exactly at T_e: the picture of a coupon known in arrears. Second, the final compounded rate sits a touch above the simple average of the same fixings, the interest-on-interest that Equation (1) keeps and the average discards.

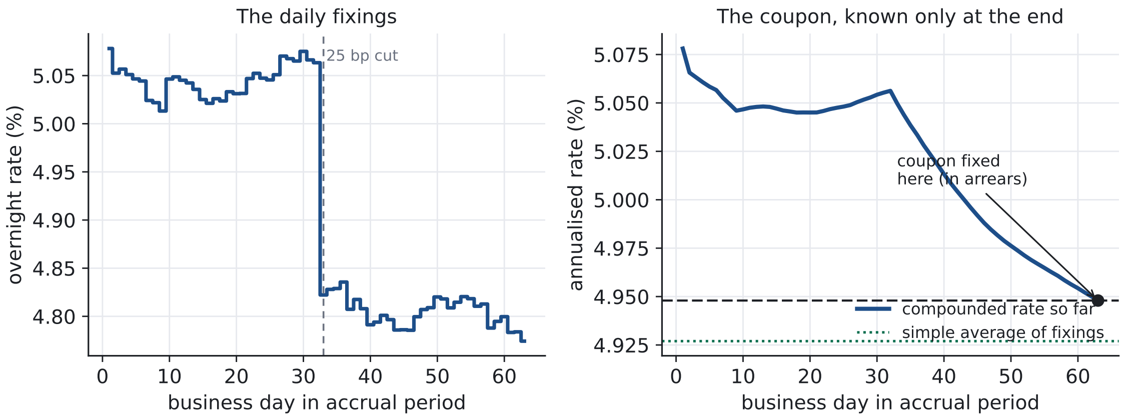

Figure 1: Compounding one coupon in arrears. *Left:* the overnight fixings over a three-month period, including a mid-period cut. *Right:* the annualised compounded rate of Equation (1) as the period fills in; it is fully determined only at T_e (the marked point), and finishes slightly above the simple average of the fixings. Synthetic path.

The key structural fact hides in Equation (1). The product on the right is exactly the growth of the RFR *money-market account* $B_t = \prod_{t_i \le t}(1 + \delta_i r_i)$ across the period:

$$1 + \delta_{s,e}\, R(T_s, T_e) \;=\; \frac{B_{T_e}}{B_{T_s}}. \tag{2}$$

Compounding the overnight rate does not *approximate* anything—it *reconstructs* the bank account, the numeraire of the risk-neutral measure. That identity—compounding rebuilds the bank account—is the hinge for two later results: the discount curve bootstraps straight out of OIS quotes (§8), and the backward-looking forward turns out to need no convexity adjustment (§9).

4 The overnight index swap

The instrument that trades the compounded overnight rate directly—and the one the rest of this note leans on—is the **overnight index swap** (OIS). It is an ordinary fixed-for-floating interest-rate swap whose floating index is the overnight RFR: over each period, one side pays a *fixed* rate agreed at the outset, the other pays the *compounded overnight rate* $R(T_s, T_e)$ of Equation (1), and only the net difference changes hands. A short OIS (a year or less) has a single period; a longer one pays annually, each floating coupon being the overnight rate compounded over that year.

What the market quotes is the *par OIS rate* K: the fixed rate that makes the swap worth nothing today—the rate at which the two legs have equal present value. To find it, value each leg. Write $P^d(0, T)$ for the discount factor—today's value of \$1 paid at T; the superscript d, for *discount*, anticipates §5–§8, where this curve becomes one of two and takes centre stage—and α_n for the accrual fraction of the n-th fixed period.

The fixed leg pays $K\alpha_n$ at each date T_n. Its value is K times the *annuity*, the discounted sum of the accruals:

$$V_{\text{fixed}} \;=\; K\,A, \qquad A \;=\; \sum_{n=1}^{N} \alpha_n\, P^d(0,T_n).$$

The floating leg pays, each period, exactly the interest a rolling overnight deposit would earn. So the whole leg replicates one simple trade: put \$1 on deposit today, roll it at the overnight rate, skim off the accrued interest at the end of every period, and keep the \$1 of principal until T_N. Valued today, that stream of interest is worth what you paid in (\$1) less the present value of the \$1 principal returned at the end:

$$V_{\text{float}} \;=\; 1 - P^d(0,T_N).$$

Setting the two equal, $K\,A = 1 - P^d(0,T_N)$, gives the par rate

$$K \;=\; \frac{1 - P^d(0,T_N)}{\displaystyle\sum_{n=1}^{N} \alpha_n\, P^d(0,T_n)}, \tag{3}$$

the classic "one minus the final discount factor, over the annuity." In words, K is a discounted average of the overnight rates the market expects over the tenor.

Quoted across maturities from a week to decades, the rates $K_1,\ldots,K_N$ are the RFR market's own term structure, and two things in this note are built directly on them. Read *forwards*, Equation (3) turns a curve of discount factors into a swap rate; read *backwards*, each quote K_n pins the next discount factor, which is how the whole discount curve is stripped from OIS quotes (§8). And forward-looking *term rates* are read off forward-starting OIS in the same way (§10).

Historically the same contract traded on the pre-reform overnight rates (Fed Funds in the US, EONIA in the euro area); today it references the RFRs of §2. Two loose ends now hang off this instrument, and they set the agenda. Why should the OIS curve—rather than anything else—be the discount curve at all? That is a question about how trades are *funded*, and it is where we go next (§5). And because its floating leg is compounded in arrears, an OIS coupon already shares the defining feature of the new world: it is not known until the period ends (§6).

5 OIS discounting: why collateral sets the discount rate

Pricing a derivative means discounting its future cashflows to today—but at *what* rate? For decades the answer was taken for granted; then the crisis turned it into one of the most consequential questions on a rates desk. This is the *other* half of the transition (the coupon was the first half), and it moved years earlier.

The old answer, and why it broke

Before 2007, desks discounted at LIBOR. The reason matters, because it explains both the old world and the shock that ended it. Discounting a future cashflow is, for a bank, nothing but the time value of *its own* money—its **funding cost**. A trading desk holds no free cash: to carry any position—to buy a bond, or merely to be owed money on a trade that has not yet paid—it must *fund* that position, drawing cash (in practice from the bank's central treasury) and paying interest on it for as long as the position is open. The bank, in turn, raises that cash in the market: taking deposits, issuing short-term paper, and borrowing from other banks. And the rate at which a bank borrowed *unsecured* from other banks for a few months *was*

LIBOR—exactly what the LIBOR survey measured (§1). So LIBOR was the marginal cost of a bank's cash, and discounting at LIBOR simply valued future cashflows at the desk's own cost of money. It also explains the internal coherence of the single-curve world: the *same* LIBOR was both the floating index a swap *paid* and the rate everything was *discounted* at.

That cash can be raised in two ways, and the difference is about to become the whole story. Borrowing *unsecured*—on the bank's name alone—costs more, because the lender bears the bank's default risk; for a bank that cost was about LIBOR. Borrowing *secured*—against collateral the borrower posts, as in a repo—costs less, close to the overnight risk-free rate, because the lender can seize the collateral if the borrower fails. Before the crisis these two were nearly equal, so it made no visible difference which notionally "funded" a trade. Collateralised derivatives, as we are about to see, are funded the *secured* way.

For pricing, meanwhile, it *did not matter* whether you discounted at LIBOR or at the overnight (OIS) rate: the two sat within a basis point or two of each other, and the choice was invisible.

The crisis made it visible. The spread between three-month LIBOR and OIS—flat near zero for years—blew out to tens and, at the Lehman peak, several hundred basis points (Figure 2, red). "Discount at LIBOR" and "discount at OIS" were now materially different answers, worth a great deal on a long-dated book. The desk had to decide which was *right*—and the answer turned out to be not a convention but a consequence of how trades are funded. (The same figure shows a second gap, the tenor basis, which forces the *projection* side apart in §8.)

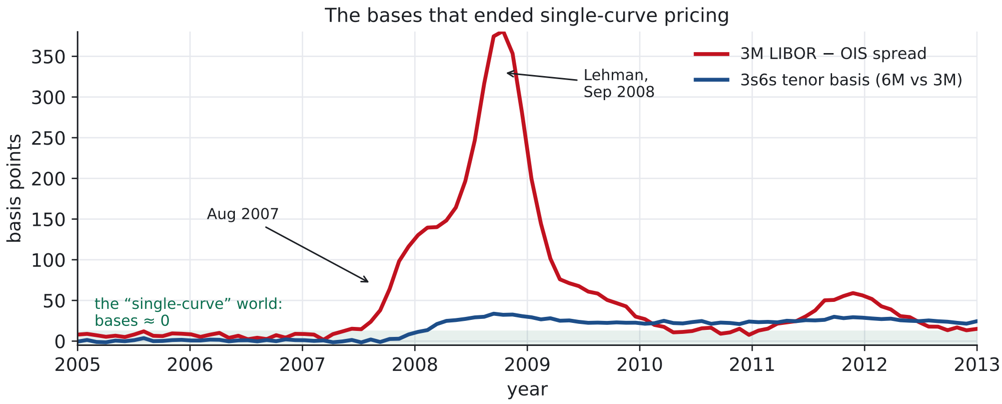

Figure 2: The bases that ended single-curve pricing. Before 2007 the LIBOR–OIS spread (red) and the 3s6s tenor basis (blue) sat near zero, and one curve served for both discounting and projection. After the crisis both opened up and never fully closed: the LIBOR–OIS spread forces the *discount* side onto OIS (this section), and the tenor basis forces the *projection* side apart (§8). (Synthetic series with the historical shape—the Lehman spike reached several hundred basis points—not market quotes.)

The collateral argument

Modern inter-dealer derivatives trade under a **Credit Support Annex** (CSA): the two parties mark the trade to market every day, and whichever is out-of-the-money posts *collateral*—variation margin—to the one that is in-the-money, cash covering the current value. The party *holding* that collateral pays interest on it, overnight, at the agreed *collateral rate c*.[1]

That single fact fixes the discount rate (Figure 3). Suppose a trade is worth V to you today. Your counterparty posts you V in cash; you hold it, but must pay the overnight rate c on it for

[1]The *effective federal funds rate* (Fed Funds, or EFFR) is the US overnight rate for *unsecured* interbank lending of reserve balances—the traditional US overnight rate, and what OIS and CSAs referenced before SOFR. It is *not* the LIBOR replacement: since the crisis, banks hold ample excess reserves and lend one another little unsecured, so its volume is thin, whereas SOFR rests on the vastly larger secured (repo) market of §2. Fed Funds appears here only in the *discounting* story—the collateral rate the 2020 "big bang" (§13) later switched to SOFR.

as long as you do. Economically, you are carrying the derivative on money that costs exactly c: the position is *financed at the collateral rate*. Replicating it therefore funds at c, so its value must discount future cashflows at c—not at LIBOR. In the language of §3, the correct numeraire is the bank account B that grows at the collateral rate, and a fully collateralised claim paying X at T is worth

$$V_0 \;=\; \mathbb{E}^{\mathbb{Q}}\!\left[e^{-\int_0^T c_s\,\mathrm{d}s}\,X\right] \;=\; P^d(0,T)\,\mathbb{E}^{\mathbb{Q}^T}[X], \tag{4}$$

with $P^d(0,T) = \mathbb{E}^{\mathbb{Q}}[1/B_T]$ the discount factor grown at the collateral rate—the same P^d that runs through the rest of the note. (Here $\mathbb{Q}$ is the risk-neutral measure and $\mathbb{Q}^T$ the T-forward measure; Appendices D–E review both, and the second equality is worked line by line there.)

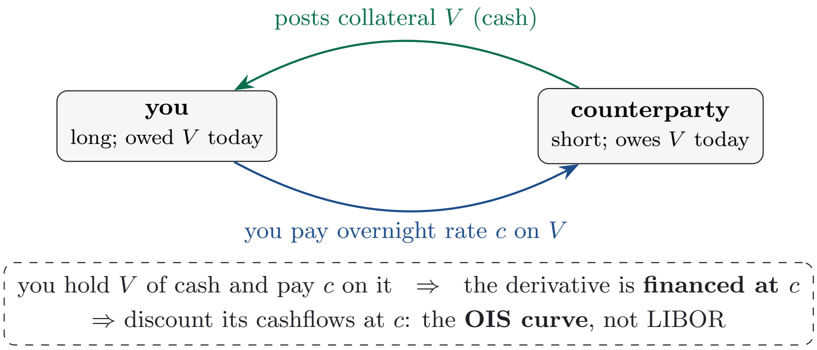

Figure 3: Why a collateralised trade discounts at the overnight rate. When a trade is worth V to you, your counterparty posts V of cash collateral, on which you pay the overnight collateral rate c. You are therefore carrying the position at cost c—it is financed at the collateral rate—so its fair value discounts future cashflows at c. Standard CSAs make c an overnight RFR, so the discount curve is the OIS curve.

The collateral rate is the OIS rate

Standard CSAs pay the *overnight* rate on cash collateral—Fed Funds historically, now SOFR in dollars, €STR in euros, and so on. The term structure of discount factors grown at that overnight rate is exactly the **OIS curve** of §4: "discount at the collateral rate" *is* "discount on the OIS curve," bootstrapped from OIS quotes (we build it explicitly in §8). Two riders are worth noting. The discount rate is *whatever the CSA actually pays*, so the discount curve is really a property of the collateral agreement—*CSA discounting*—and a book collateralised in several eligible currencies carries several discount curves, with a cheapest-to-deliver optionality over which collateral to post. And an *un*collateralised trade has no posted collateral to pin the rate; its funding and counterparty risk must be priced through separate valuation adjustments (FVA, CVA) that sit outside this note. For the standard single-currency, cash-collateralised trade, the discount curve is simply the domestic RFR, which is what we assume throughout.

Why this matters for the transition

OIS discounting was the *first* structural break of the crisis, and it was largely complete years before LIBOR's end: by the early 2010s collateralised and cleared books discounted on OIS, not LIBOR. It decoupled discounting from projection—the discount curve now came from OIS, the coupon still from LIBOR—which is exactly the two-curve world §8 will formalise. And it means the LIBOR transition inherited a discount side that had *already* moved to the RFR; what remained was to switch the overnight rate itself (Fed Funds/EONIA → SOFR/€STR) in the coordinated "big bang" of §13.

6 Known in advance versus known in arrears

The deepest structural change is one of *timing of information*. A LIBOR coupon was fixed at the start of its accrual period and paid at the end: you knew the cashflow the moment the period began. A compounded-RFR coupon is assembled from overnight fixings *across* the period, so it is not known until the period is essentially over (Figure 4).

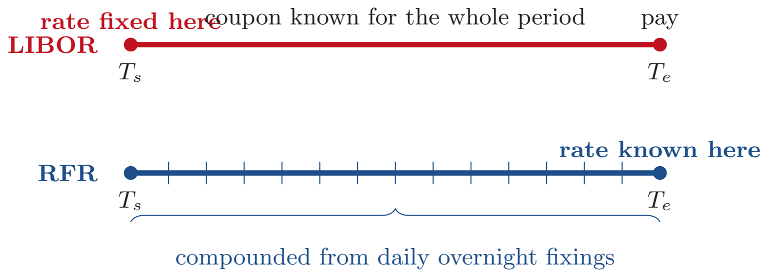

Figure 4: The timing inversion. A LIBOR coupon (red) is fixed at T_s and paid at T_e—known for the whole period in advance. A compounded-RFR coupon (blue) accumulates from the daily overnight fixings and is not known until T_e: known in arrears. Conventions (§7) claw back a few days' notice.

This sounds like an operational nuisance, and in part it is: a payer now learns the size of a payment a day or two before it is due, not three months ahead, so treasury and settlement functions need conventions that restore a little notice (§7). But it is also a genuine change in the mathematics. Anything that depends on the coupon *during* the accrual period—an early redemption, a cap observed mid-period, an option that expires before the period ends—now references a quantity that is only *partially* realised. The clean old picture, in which the floating rate was a single number revealed at T_s, is gone; in its place is a rate that congeals gradually. Making that precise—and showing, perhaps surprisingly, that the vanilla coupon costs no convexity for it—is the business of §9.

7 Buying back notice: the conventions

Knowing the coupon only at T_e is awkward for anyone who has to move cash on T_e. Three conventions restore a few days of notice, each trading off notice against faithfulness to the true compounded rate (Figure 5).

Lookback. Observe each day's rate as of a fixed number k of business days *earlier*, so the observation window is the accrual window shifted back by k days. The last rate you need is then known k days before T_e. Under a *lookback without observation shift*, the shifted rates are weighted by the day-count of the accrual day they stand in for; under a *lookback with observation shift*, both the rate *and* its day-count weight come from the shifted day, which is internally consistent and easier to hedge.

Lockout (suspension). Compound normally until k days before T_e, then *freeze* the rate at that last value for the final k days. Operationally the simplest—the tail is known early—but it distorts the true rate slightly and is awkward if the frozen days straddle a rate move. Used, for example, in some SOFR floating-rate notes.

Payment delay. Keep the true compounded rate but pay it k business days after T_e. This is the cleanest analytically—the coupon is undistorted—at the cost of a small timing (convexity) adjustment—made precise once we have the pricing measure (§9)—and a later cashflow.

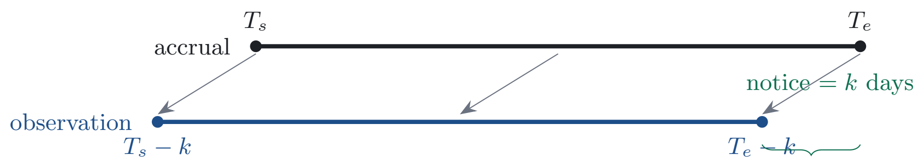

Figure 5: Lookback with an observation shift. The observation window (blue) is the accrual window (black) slid back by k business days, so the final rate needed is known k days before T_e—the notice the desk buys back. Lockout instead freezes the last k observations; payment delay leaves the window untouched and moves the cashflow later.

A numerical comparison makes the trade-off explicit. Take a 63-business-day coupon (ACT/360), fixings at 5.00%, and let the central bank cut twice near the end: 25 bp on day 56 and another 25 bp on day 61. Table 3 runs the *same* path through each convention with $k = 5$ days.

Convention	**What it sees**	**Coupon**	**vs. plain**
Plain in arrears	every fixing, days 1–63	4.9776%	—
Lookback ($k = 5$, obs. shift)	days −4–58: catches 3 days of cut 1 only	5.0096%	+3.2 bp
Lockout ($k = 5$)	days 1–58, then frozen at day-58 (4.75%)	4.9896%	+1.2 bp
Payment delay ($k = 5$)	every fixing; cashflow 5 days later	4.9776%	timing only

Table 3: One fixing path, four conventions (5.00%, cut 25 bp on day 56 and again on day 61). The lookback's shifted window misses most of both cuts; the lockout catches the first cut but freezes through the second; the payment delay reproduces the true coupon and pays it late. Notice is bought with staleness—each day of warning is a day of rate moves the coupon can no longer see.

The ordering is no accident: the more notice a convention buys, the more of the late rate move it fails to pass through, and which side of that trade you are on depends on which way rates moved. On this path (cuts) the borrower overpays under lookback and lockout by a few basis points; had the central bank hiked, the same conventions would have undercharged. Averaged over paths the conventions are close to fair—but trade by trade they differ, which is exactly the small basis the text above describes.

None of these is exotic; they are plumbing. But they matter for hedging, because two notes on the same index with different conventions carry a small basis, and because a lookback or lockout shifts *which* fixings land in a coupon—which is exactly what a hedge has to replicate. Market bodies such as the **ARRC** in the US—the Alternative Reference Rates Committee, the industry group the Federal Reserve convened to steer the dollar transition—and the sterling working group published recommended conventions to keep this basis contained [7].

8 Multi-curve pricing and the projection index

The single curve, and how it broke

In the world exotics desks grew up in, one zero-coupon curve did two jobs at once. From discount factors $P(0,T)$ you read off both the *forward rates* that set a swap's floating coupons and the *discount factors* that present-value them. The forward LIBOR for the period $[T_{i-1}, T_i]$, of length δ_i, was

$$F_i(0) \;=\; \frac{1}{\delta_i}\left(\frac{P(0,T_{i-1})}{P(0,T_i)} - 1\right), \tag{5}$$

and the fair rate of a forward-rate agreement was the very same number. One curve, read two ways, and no ambiguity about which was which.

Then, from August 2007, the identities that held that picture together came apart. §5 has already met the first symptom—the LIBOR–OIS spread—which forced *discounting* onto OIS. The *second* symptom splits the projection side: a *tenor basis* opened up. Earning three-month LIBOR twice in a row no longer matched earning six-month LIBOR once, so the "same" curve implied different forwards depending on which tenor you built it from. This basis is what the market quotes as a *basis swap*. The "3s6s" ("threes-sixes") basis of Figure 2 exchanges three-month LIBOR against six-month LIBOR on the same currency and notional; because the shorter tenor re-fixes more often and locks in less credit and liquidity exposure per reset, it trades below the longer one, and a positive spread—the 3s6s basis—is added to the *shorter* (3M) leg to make the swap fair. It sits at one end of a family (1s3s, 3s6s, 6s12s), and before 2007 every member was essentially zero, exactly as the single-curve no-arbitrage in Equation (5) demands. A single curve cannot reprice OIS, three-month-LIBOR and six-month-LIBOR instruments simultaneously once those bases are nonzero. The single-curve world was over.

One discount curve, many projection curves

The repair is to stop reading one curve two ways and instead *build two kinds of curve from two kinds of instrument*. Figure 6 shows the change of plumbing at a glance; the mechanics follow.

Before, in the single-curve world (Figure 6, top), a single zero-coupon curve $P(0,T)$ was bootstrapped from one set of instruments and then did *both* jobs. Its ratios gave the forward rates that projected the coupons (Equation (5)); the very same discount factors present-valued them—one curve in, two uses out. That is exactly why, in that world, a floating leg collapsed to the simple closed form $1 - P(0,T_N)$ derived below.

After (Figure 6, bottom), the two jobs are fed by *different* instruments and yield *different* curves. Discount factors are stripped from OIS quotes alone, giving the discount curve P^d; the forwards for each tenor are stripped from *that tenor's own* instruments—its par swaps, FRAs and basis swaps—*taking the discount curve as an input*, giving one projection curve per tenor. Nothing now forces the projected forward to equal the discount-implied forward, and their gap is the tenor basis. (The split was forced by the 2007–08 crisis, well before the LIBOR transition itself; the transition later changed *which* index each projection curve carries—the point this section closes with.)

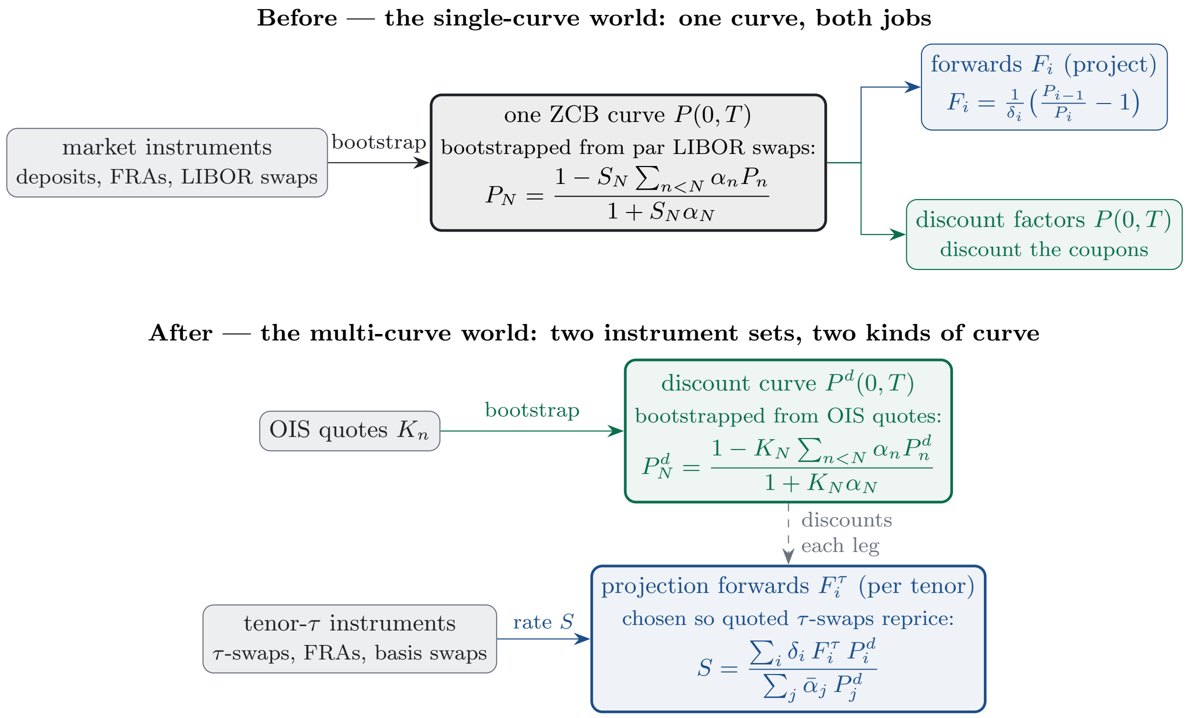

Figure 6: Curve construction, before and after the split (writing $P_n \equiv P(0,T_n)$ and $P^d_n \equiv P^d(0,T_n)$). *Top (single-curve):* one zero-coupon curve P is bootstrapped from a single set of instruments—each par LIBOR swap rate S_N pinning the next factor P_N—and then serves *both* jobs: a forward is read straight off it, $F_i = \frac{1}{\delta_i}(P_{i-1}/P_i - 1)$, and a floating leg is worth just $1 - P_N$. *Bottom (multi-curve):* the jobs separate. The discount curve P^d is bootstrapped from OIS quotes K_n by the *same recursion* (compare the two curve boxes—only the input rate, S versus K, differs). The projection forwards F^τ_i are then the *unknowns*: they are chosen so the tenor's quoted par swaps reprice, and that repricing (Equation (10)) *discounts on the separate* P^d. That is where the two inputs enter—the tenor instruments set the quoted rate S the forwards must match, and P^d supplies the discounting. The projected forwards need not equal those P^d alone would imply; their gap is the tenor basis of Figure 2.

The rest of this section makes the lower panel precise: *how* each curve is stripped. Discounting itself is already settled—§5 showed that a collateralised trade discounts on the OIS curve P^d—so we take that as given, build P^d from OIS quotes below, and then build the projection curves on top of it.

To see what breaks, value a floating leg the *old* way first, on the single curve P. With each forward $F_i(0)$ read off that same curve through Equation (5), every coupon's contribution collapses, $\delta_i F_i(0) P(0,T_i) = P(0,T_{i-1}) - P(0,T_i)$, and the leg *telescopes* to a single pair of discount factors:

$$\begin{aligned} V^{\text{single}}_{\text{float}}(0) &= \sum_i \delta_i \, F_i(0) \, P(0,T_i) \\ &= \sum_i \big[P(0,T_{i-1}) - P(0,T_i)\big] \\ &= 1 - P(0,T_N). \end{aligned} \tag{6}$$

This is the same "par floater" identity as the OIS leg of §4: when one curve both projects and discounts, a floating leg is worth just $1 - P(0,T_N)$, independent of the tenor.

Projection in the new world is handled by a **projection curve for each index**. There is one curve that reproduces three-month-LIBOR instruments, another for six-month LIBOR, another still for each tenor, each stripped so that Equation (5)—now read off that index's *own* projection curve—reprices its market. The forward $F_i(0)$ that sets a coupon now comes from the projection curve of the coupon's index, while the discount factor $P^d(0,T_i)$ comes from the *separate* OIS discount curve, so the same floating leg is

$$V_{\text{float}}(0) \;=\; \sum_i \delta_i \, F_i(0) \, P^d(0, T_i). \tag{7}$$

Compare Equation (7) with Equation (6): because F_i and P^d are now drawn from *different* curves, the cancellation is gone—the sum no longer telescopes and the leg is *not* simply $1 - P^d(0, T_N)$. That broken telescoping is the entire content of "multi-curve" pricing: a coupon is *forecast* by one curve and *discounted* by another [3, 4].

Building the discount curve from OIS

Saying the discount curve comes from "the overnight market" is precise: it is stripped from the OIS quotes of §4. Recall that the par OIS rates $K_1, \dots, K_N$ across maturities $T_1 < \cdots < T_N$ (from a week out to thirty years) are the overnight market's own term structure; the discount factors $P^d(0, T_n)$ are backed out of them.

The stripping is unusually clean because an OIS is *self-discounting.* Its floating leg compounds the very overnight rate that defines the discount numeraire—the RFR bank account B of Equation (2), with $P^d(0, T) = \mathbb{E}^{\mathbb{Q}}[1/B_T]$. The floating payment over $[T_{n-1}, T_n]$ is exactly the compounded interest $B_{T_n}/B_{T_{n-1}} - 1$, so receiving it at T_n is worth $\mathbb{E}^{\mathbb{Q}}[1/B_{T_{n-1}}] - \mathbb{E}^{\mathbb{Q}}[1/B_{T_n}] = P^d(0, T_{n-1}) - P^d(0, T_n)$, and the whole floating leg *telescopes*:

$$\begin{aligned} V^{\text{OIS}}_{\text{float}}(0) &= \sum_{n=1}^{N} \big[P^d(0, T_{n-1}) - P^d(0, T_n) \big] \\ &= 1 - P^d(0, T_N), \end{aligned} \tag{8}$$

using $T_0 = 0$ and $P^d(0, 0) = 1$—the replication asserted in §4, now shown term by term. Setting it equal to the fixed leg is exactly the par condition behind Equation (3); solved for the last, still-unknown discount factor it becomes a *bootstrap*, each new quote K_N pinning the next factor from the ones already known,

$$P^d(0, T_N) \;=\; \frac{1 - K_N \sum_{n=1}^{N-1} \alpha_n \, P^d(0, T_n)}{1 + K_N \, \alpha_N}. \tag{9}$$

Running up the maturity ladder (Figure 7) turns the quoted OIS curve into the discount curve, one factor at a time. Concretely, with annual payments a one-year OIS quoted at $K_1 = 3.00\%$ gives $P^d(0, 1) = 1/1.03 = 0.9709$; a two-year quote $K_2 = 3.20\%$ then gives $P^d(0, 2) = (1 - 0.032 \cdot 0.9709)/1.032 = 0.9389$, and so on. (In practice one interpolates for dates between quotes, but this recursion is the skeleton.)

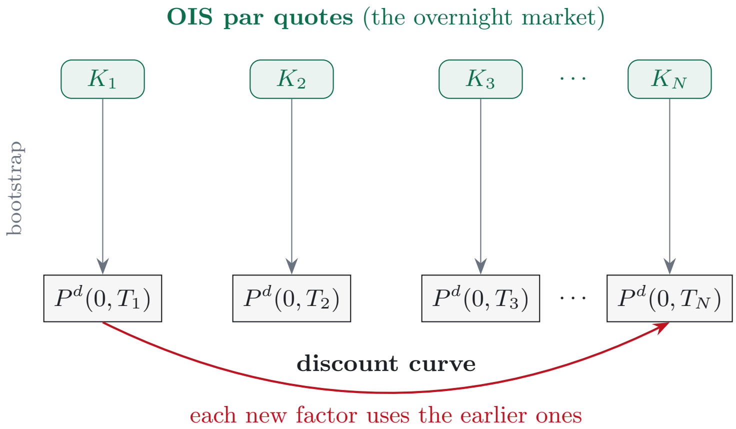

Figure 7: Bootstrapping the discount curve. Each quoted OIS par rate K_n pins one discount factor $P^d(0,T_n)$ through Equation (9); because the OIS floating leg telescopes (Equation (8)), the recursion is explicit and each factor depends only on the shorter-maturity ones already solved.

That the OIS is self-discounting—projected *and* discounted by the same overnight rate—is exactly the sense in which, for the overnight index itself, the multi-curve split collapses back to a single curve. It is the same coincidence that §13 returns to when the whole market re-anchors on RFR discounting.

Building the projection curves

The discount curve settles how to *present-value*; a projection curve settles the *forwards* $F_i^\tau(0)$ that fix the coupons of a given tenor τ (three months, six months, ...). With the OIS discount curve P^d already in hand, each projection curve is stripped from the instruments that actually reference its tenor: par vanilla swaps on τ, forward-rate agreements, and the tenor basis swaps of §8 that tie one tenor to another.

The calibration target is the quoted par swap rate. A fixed-for-floating swap on tenor τ, with fixed accruals $\bar{\alpha}_j$ and floating accruals δ_i, is at par when its fixed rate S makes the two legs equal in value. Discounting both on P^d and projecting the floating leg on the τ curve,

$$S \;=\; \frac{\displaystyle\sum_i \delta_i\, F_i^\tau(0)\, P^d(0,T_i)}{\displaystyle\sum_j \bar{\alpha}_j\, P^d(0,T_j)}. \tag{10}$$

This is the multi-curve twin of the OIS par rate of Equation (3): the denominator is again a P^d-annuity, but the numerator now carries the *projected* forwards F_i^τ instead of collapsing by telescoping. Given P^d, one chooses the τ curve so that Equation (10) reproduces every quoted S up the maturity ladder—the same bootstrap idea as for the discount curve, but now solving for the projected forwards.

The information the projection curve adds over the discount curve is precisely a basis. Writing $f_i^d(0)$ for the forward the *discount* curve itself would imply over $[T_{i-1}, T_i]$ (Equation (5) evaluated on P^d), the projected forward splits as

$$F_i^\tau(0) \;=\; f_i^d(0) \;+\; b_i^\tau, \tag{11}$$

with b_i^τ the tenor-τ forward basis—small, positive, and larger for longer tenors, closely related to the spreads Figure 2 plots. A projection curve is, concretely, “the discount curve plus the observed basis for this tenor.” In practice the discount curve and every projection curve are not stripped in isolation but calibrated *jointly*, in one solve, to reprice OIS, single-tenor swaps and tenor basis swaps at once.

For the new RFR indices this apparatus almost vanishes. A SOFR or SONIA coupon is projected by compounding the very overnight rate the discount curve is built on, so its projection curve *is* the discount curve, $b_i^\tau = 0$, and Equation (10) collapses back to the OIS par rate of Equation (3). The tenor-basis machinery was a LIBOR-era necessity; the transition largely retires it for a single currency (§13).

A basis swap is the cleanest instrument in which to see this decoupling at work: Appendix C prices a 3s6s swap in this framework and shows that its fair spread is nonzero *precisely* because the two tenors need separate projection curves—were one curve to project both, the spread would be forced to zero.

Why this is the right language for the transition

Once the projection curve and the discount curve are separate, the LIBOR transition has a one-line description: **it is a change of projection index**. LIBOR was merely one family of projection curves—one per tenor. Retiring it means replacing those curves with a single new prescription: the projected coupon is now a *compounded overnight RFR* rather than a term LIBOR fixing. The discount side, meanwhile, had *already* moved to OIS/RFR years earlier, when collateralisation forced the discount side onto OIS (§5); the residual tidying of the discount curve (the "big bang" of §13) was a separate, largely completed step. Holding these two apart—projection changing now, discounting changed already—keeps the whole transition legible. The compounded overnight rate is that new projection index in all its aspects: how it is built (§3), the instrument that trades it (§4), when it is known (§6), how notice is restored (§7), how it behaves under a pricing measure (§9), and its forward-looking cousin (§10).

9 The measure subtlety: a backward-looking forward is clean

This section is the mathematical core of the note. It relies on numeraires and forward measures, for which Appendices D and E provide a self-contained review; Appendix F develops the convexity adjustments met at its end.

Recall the fact every exotics desk relies on. In the single-curve world the forward LIBOR $L(t; T_s, T_e)$ is a martingale under the T_e-forward measure $\mathbb{Q}^{T_e}$, whose numeraire is the zero-coupon bond $P(\cdot, T_e)$:

$$L(t; T_s, T_e) \;=\; \mathbb{E}^{\mathbb{Q}^{T_e}}\big[\, L(T_s; T_s, T_e) \,\big|\, \mathcal{F}_t \,\big], \qquad t \le T_s. \tag{12}$$

Because it is set at T_s, its terminal value is fixed before the T_e payment, and a caplet is priced by Black's formula under $\mathbb{Q}^{T_e}$. Everything is known in advance.

Now take the backward-looking compounded rate $R(T_s, T_e)$. It is not known until T_e, so one might expect a messy convexity correction to project it. The opposite is true. Use the discount-curve version of Equation (2): since the RFR *is* the collateral/discount rate, the bank account B is the numeraire of the discount curve, and for $t \le T_s$ the fair forward of R is

$$R(t; T_s, T_e) \;=\; \frac{1}{\delta_{s,e}}\left(\frac{P^d(t, T_s)}{P^d(t, T_e)} - 1\right), \tag{13}$$

the *same* functional form as the single-curve forward of Equation (5)—a ratio of discount factors, now read off the discount curve itself. Better still, this forward is a martingale under the same $\mathbb{Q}^{T_e}$, and when the coupon is paid at T_e there is **no convexity adjustment**: the fair fixed rate on a forward contract on R is just the OIS forward rate for the period. The reason is Equation (2). Compounding the overnight rate rebuilds B_{T_e}/B_{T_s}, and a ratio of the numeraire to a bond is a martingale under the corresponding forward measure by construction [5, 6].

What changes is not the expectation but the *information schedule* (Figure 8). For $t \le T_s$, before the period starts, the backward-looking forward and the old LIBOR forward carry the *same* uncertainty: both are forward views of the same span. At T_s they diverge. LIBOR resolves all at once—its remaining uncertainty drops to zero. The compounded rate instead sheds its uncertainty gradually: with each passing day one more fixing is locked in and B_t/B_{T_s} becomes known, so the variance of the still-unknown remainder bleeds to zero, hitting zero exactly at T_e. Same fair forward, same terminal expectation, but one is revealed in a single step and the other over the whole period.

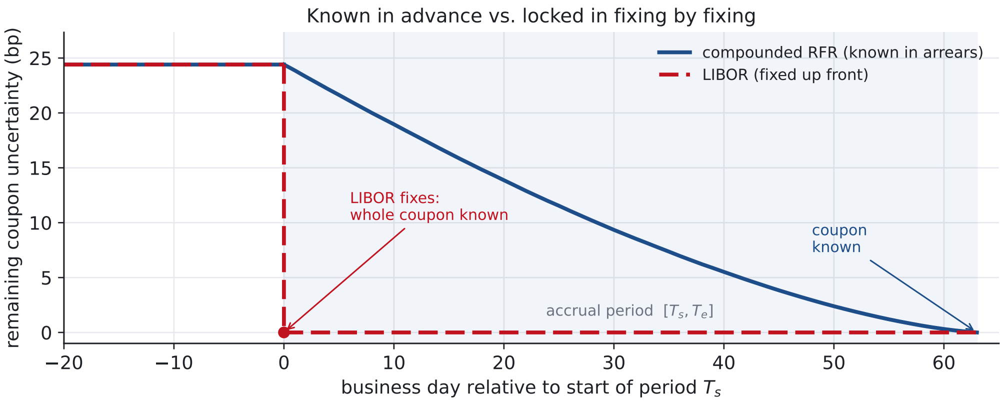

Figure 8: Why the backward-looking forward is clean. Before the period starts, the compounded-RFR forward (blue) and the LIBOR forward (red) share the same uncertainty. At T_s LIBOR fixes and its coupon is fully known; the RFR coupon instead locks in fixing by fixing, its remaining uncertainty decaying to zero at T_e. The forward is a martingale that simply accumulates realised fixings—hence no convexity adjustment for the natural (paid-at-T_e) coupon. Toy Gaussian short-rate model.

Where, then, does convexity live? Not in the plain in-arrears coupon, but in every departure from its natural timing:

- **Payment delay.** Pay the compounded rate at $T_p > T_e$ and the exact cancellation in Equation (13) no longer holds against $\mathbb{Q}^{T_e}$; a small timing (convexity) adjustment appears, of the usual sign and size.
- **Setting in advance.** Some conventions pay, over $[T_s, T_e]$, the compounded rate observed over the *previous* period. That is a genuinely forward-set quantity and carries its own adjustment.
- **Term-rate approximation.** Substituting a forward-looking term rate (§10) for the realised compounded rate introduces a term-versus-compounded basis whenever rates move within the period.
- **Options observed during the period.** A cap on the backward-looking rate is, once inside $[T_s, T_e]$, an option on a quantity whose remaining volatility is *shrinking* toward zero (Figure 8). The effect has a simple closed form. Model the forward's volatility as constant, σ, before the period starts, then decaying linearly as fixings lock in—$\sigma(t) = \sigma\,(T_e - t)/\tau$ for $t \in [T_s, T_e]$, with $\tau = T_e - T_s$ the unfixed fraction. The Black variance to expiry is then

$$\int_0^{T_e} \sigma(t)^2\,\mathrm{d}t \;=\; \sigma^2 T_s \;+\; \sigma^2 \int_{T_s}^{T_e} \Big(\frac{T_e - t}{\tau}\Big)^2 \mathrm{d}t \;=\; \sigma^2\Big(T_s + \frac{\tau}{3}\Big),$$

whereas the corresponding forward-looking (LIBOR-style) caplet, whose rate fixes at T_s, accrues only $\sigma^2 T_s$. A backward-looking caplet therefore prices like Black with its expiry extended by *one third of the accrual period*—the "one-third rule"—slightly more option value than its forward-looking twin, the premium for the extra fixing risk that trickles in through the window [5].

This is the substantive news for an exotics book: the vanilla coupon got *simpler* to project, and the optionality migrated into timing and into the accrual window itself.

10 The term-rate question

Compounding in arrears is natural for a swaps desk, but genuinely awkward for a corporate treasurer who wants to know next quarter's interest bill *today*, to budget against. That demand kept alive a forward-looking option: a **term RFR**. Term SOFR (published by CME) and Term SONIA are set at the *start* of a period, like LIBOR—but they are not surveys. They are read off the RFR *derivatives* market (SOFR futures and OIS) and are, in effect, that market's forward-looking expectation of what the overnight rate will compound to over the coming one, three or six months. In the language of §9, a term rate is simply $R(t; T_s, T_e)$ from Equation (13) read at $t = T_s$: the fair fixed rate of a forward-starting OIS, credit-free and known in advance—LIBOR's convenience without LIBOR's flaws.

So which rate does a given contract actually use? The market drew a deliberate line (Figure 9). **Cash products**—business loans, trade finance, some securitisations—may use Term SOFR, because the borrowers on the other side value knowing their coupon up front and trade in modest size. **Derivatives**—swaps, futures, cleared trades and most floating-rate notes—may *not*. They reference the *compounded RFR in arrears* of §3, the same backward-looking rate this note has been building throughout. A plain vanilla swap, with LIBOR retired and term rates excluded, therefore references compounded SOFR (or SONIA, €STR, . . .), set in arrears.

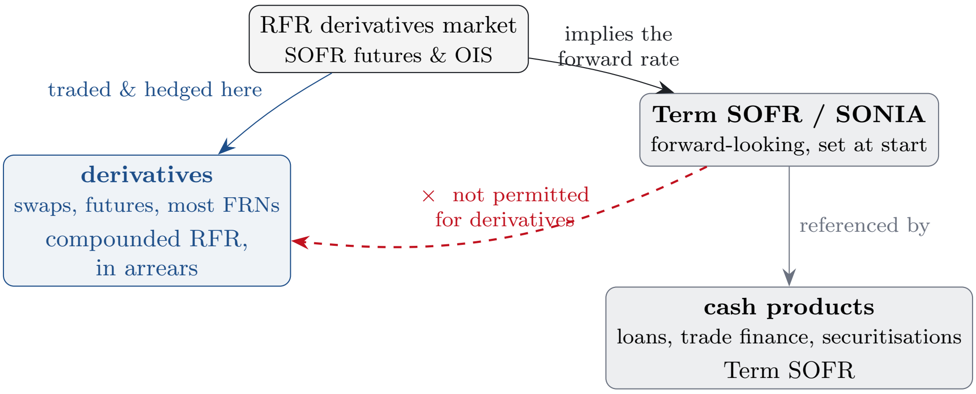

Figure 9: Which rate does each market use? The RFR *derivatives* market (SOFR futures and OIS) is the deep, liquid foundation. *Derivatives* reference the *compounded RFR in arrears* of §3 directly. A forward-looking *term rate* (Term SOFR/SONIA) is *read off* that same market and published at the start of each period; it may be used by *cash products*, whose borrowers want their coupon known in advance. It is deliberately *not* permitted for derivatives: because the term rate is derived from the derivatives market, letting derivatives reference it would let the tail wag the dog and re-thin the underlying—the very failure that killed LIBOR.

That prohibition—the red arrow in Figure 9—is not bureaucratic. If the mass of contracts *referencing* a term rate were allowed to dwarf the futures market it is *read from*, hedging them would swamp and distort that market—re-creating the thin-underlying fragility the transition set out to escape. Confining term rates to the (smaller) cash markets keeps the huge, liquid derivatives market as the robust foundation and the term rate as a thin, well-anchored read-off on top. This is why the **ARRC** and its peers recommended term rates for loans but reserved derivatives for compounded RFR [7].

There is also a price for the convenience. A forward-looking term rate and the realised compounded rate for the same period are not the same number: the term rate is the market's *expectation* at T_s, the compounded rate is the *realisation* at T_e, and the two diverge whenever rates move within the period. That gap—the *term-versus-compounded basis*—is the convexity flagged in §9, and it is why a loan on Term SOFR and the swap hedging it (on compounded SOFR) do not cancel to the last basis point.

11 The spread adjustment

When a LIBOR setting is switched off, every contract still referencing it needs a rate to *fall back* to—a replacement that takes over automatically so the contract can keep running. That replacement is the contract's *fallback rate*. How it is installed across millions of legacy contracts is the subject of §12; what concerns us here is what the fallback rate should *be*, and in particular why it cannot be the bare compounded RFR.

Return to the credit row of Table 1. Because LIBOR carried a term bank-credit premium and an RFR does not, LIBOR sat *above* the compounded RFR almost all the time (Figure 10). A contract that simply fell back from "LIBOR" to a bare "compounded RFR" would therefore jump down in value on switch-over day, handing wealth from one side to the other for nothing. To make the change *fair*—value-neutral at the moment it happens—you add back the missing premium as a fixed *spread adjustment*, and the fallback rate becomes

$$\text{fallback rate} = \underbrace{R(T_s, T_e)}_{\text{compounded RFR}} + \underbrace{s^*}_{\text{fixed spread}}. \tag{14}$$

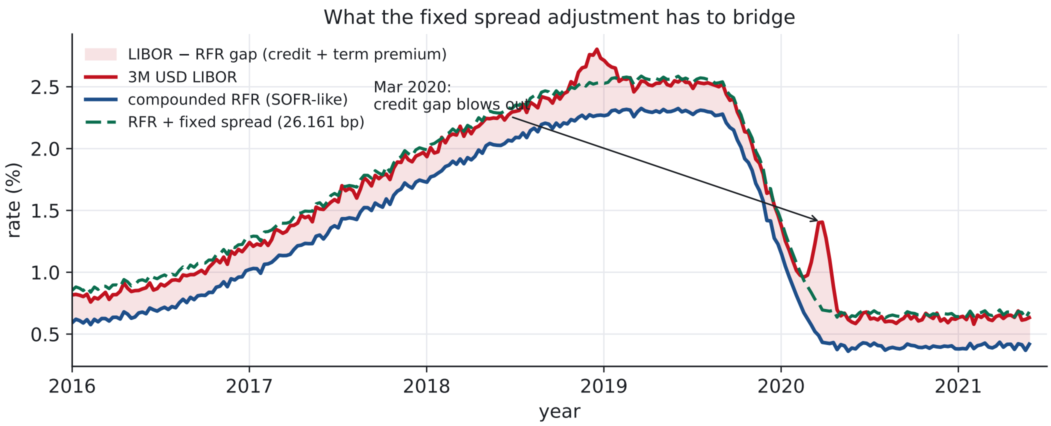

Figure 10: What the fixed spread has to bridge. Three-month LIBOR (red) sits above the compounded RFR (blue); the shaded gap is bank credit plus term premium, calm in normal times and violent in stress (the March 2020 spike). The fixed ISDA spread (green, here USD 3M = 26.161 bp) matches the gap *on average* but cannot track its swings—by design, since it is set once and for all at the historical median. Synthetic levels; the spread value is the published ISDA fixing.

The ISDA methodology fixes s^* per currency and tenor as the **five-year median** of the historical difference between LIBOR and the compounded RFR, computed over the five years up to the cessation trigger and then frozen permanently. Two design choices matter. The *median* (rather than the mean) blunts the influence of crisis-era spikes like March 2020, so one dislocation does not dominate a number that will govern contracts for decades. And *freezing* it removes any ongoing discretion—the residual, time-varying part of the credit gap in Figure 10 is deliberately *not* tracked; a fixed, fair, once-computed number is the price of certainty. The spreads were fixed on 5 March 2021 for all LIBOR settings, the day the FCA confirmed cessation [8, 9].

USD LIBOR tenor	1M	3M	6M	12M
Spread adjustment (bp)	11.448	26.161	42.826	71.513

Table 4: The ISDA fixed spread adjustments for USD LIBOR, fixed 5 March 2021 under the five-year-median methodology.[2] The spread grows with tenor because a longer unsecured term loan carries more credit and liquidity premium.

A fixed spread was not the only candidate bridge, and the road not taken is instructive. Banks' commercial lenders initially pushed for *credit-sensitive* replacement benchmarks—Bloomberg's BSBY and the American Financial Exchange's Ameribor—designed to keep coupons tracking bank funding costs in real time, as LIBOR had: when a lender's own cost of money rises in a stress, a credit-sensitive index raises its loan income with it. The objection is by now familiar. Those indices are built from unsecured bank-paper markets that are *thin*—the very fragility that killed LIBOR—and thinnest exactly in the stresses they are meant to price. Regulators said so with increasing force, IOSCO's 2023 review found the leading candidates fell short of its benchmark principles, and Bloomberg subsequently announced BSBY's cessation for late 2024. The fixed ISDA spread is therefore a deliberate design choice, not an oversight: it trades away credit *tracking* (Figure 10's time-varying gap) for robustness, and it concentrates what remains of the old credit sensitivity into a single, transparent, unchangeable number.

12 Fallbacks: fixing the legacy book

Equation (14) says what a fair replacement rate is; the *fallback* machinery is how it reaches trillions of dollars of contracts written long before anyone chose an RFR. Renegotiating each one bilaterally was never feasible, so ISDA amended the standard definitions once, centrally. The 2020 IBOR Fallbacks Supplement changed the meaning of "LIBOR" in new trades, and the accompanying Protocol let existing counterparties opt their legacy trades into the same language by both adhering [8]. For adhering parties, on a trigger, LIBOR is automatically replaced by the adjusted RFR of Equation (14): compounded in arrears, plus the fixed spread of Table 4.

Two details matter for anyone hedging a book across the change.

The triggers include "pre-cessation". Fallbacks fire not only on a permanent cessation announcement but also when the regulator declares LIBOR no longer *representative* of its underlying market—so a LIBOR that is still *being published* but has been judged unrepresentative still flips contracts to the fallback. This closed the gap in which a zombie benchmark might have limped on.

Cash and derivatives fallbacks need not coincide. The ISDA derivatives fallback is compounded-in-arrears plus spread. Many cash products (loans, FRNs) instead follow the ARRC "waterfall", which prefers Term SOFR where available before dropping to compounded SOFR. When a swap hedges a loan, the two can fall back to *different* rates, leaving a residual basis exactly where a hedge was meant to be clean—a practical wrinkle worth checking contract by contract.

Finally, for "tough legacy" contracts that could be neither amended nor allowed to fail, regulators authorised a *synthetic* LIBOR: published under the old name but computed as a term RFR plus the same ISDA spread—a temporary bridge, since wound down, that let genuinely un-amendable contracts run off on an unambiguous number.

[2] Values as published by Bloomberg under the ISDA methodology, in basis points per annum. Reproduced here to three decimals; confirm against the source technical notice before any contractual use [9].

13 The discount-curve move

The projection index is the change this note has mostly been about. The discount curve is the *other* change—and, as §5 argued, it had largely already happened. This section closes that loop, completing the symmetry.

The first stage was *OIS discounting*—the collateral-driven move set out in §5. As standard CSAs made most inter-dealer trades fully collateralised, desks moved from discounting at LIBOR to discounting at the overnight collateral rate; by the early 2010s OIS discounting was standard for collateralised and cleared books.

The second stage was a coordinated one-day switch of the *overnight* rate itself. In October 2020 the major clearing houses (LCH, CME) moved the discounting of cleared USD swaps from the effective federal funds rate to SOFR—the "big bang". Changing the discount rate changes present values and DV01, so the switch could not be silent: the clearing houses paid cash compensation for the change in NPV and distributed offsetting basis swaps for the change in risk, so that no member gained or lost from the mechanics. The euro market made the analogous €STR-discounting switch (away from EONIA) the same year.

With both stages done, curve construction in the RFR world has a near-symmetry with the crisis-era proliferation it reversed. For a given currency, one overnight index now underlies *both* the discount curve and—through compounding—the projection of a SOFR or SONIA swap, so a single-currency RFR swap is close to *single-curve* again: the tenor-basis zoo that Figure 2 opened up is largely collapsed. What remains is a smaller set of genuine bases: cross-currency basis between one RFR and another, the term-versus-compounded basis of §10, and whatever residual credit-sensitive benchmarks survive in niches. The crisis split one curve into many; the transition changed the projection index and, for each currency, folded the projection and discount curves back onto a common overnight rate. The projection index changed—and the discounting had already moved to meet it.

A Day-count and compounding conventions

Day-count fractions. A period's length as a fraction of a year depends on the convention. ACT/360 (most USD and EUR money markets) counts actual calendar days over a 360-day year; ACT/365 (sterling) over 365. So a 91-day three-month period is $91/360 = 0.2528$ years ACT/360. Overnight fixings apply for the actual number of calendar days they cover—so a Friday fixing typically accrues for three days over a weekend.
Business-day conventions. Dates that fall on holidays are rolled by "modified following" (move to the next business day, unless that crosses month-end, in which case move back).
Compounded in arrears (Equation (1)):

$$R(T_s, T_e) \;=\; \frac{1}{\delta_{s,e}}\left[\prod_{i=1}^{N}(1+\delta_i r_i) - 1\right], \qquad \delta_{s,e} = \sum_i \delta_i.$$

Simple average (some cash products):

$$R^{\text{avg}}(T_s, T_e) \;=\; \frac{1}{\delta_{s,e}}\sum_{i=1}^{N}\delta_i r_i.$$

The two differ only by interest-on-interest, a few basis points at normal rate levels and small periods, more when rates are high or the period is long.

B Interest rate swaps: mechanics and pricing

An interest-rate swap exchanges two streams of interest on a common notional N over an agreed schedule; for a single-currency swap the notional itself is never exchanged, only the interest. Its purpose is to *transform* an exposure: a business paying a floating loan can enter a swap to pay fixed and receive floating, so the two floating legs cancel and it is left paying a known fixed rate; an investor wanting floating income does the reverse. Two kinds underlie everything in this note.

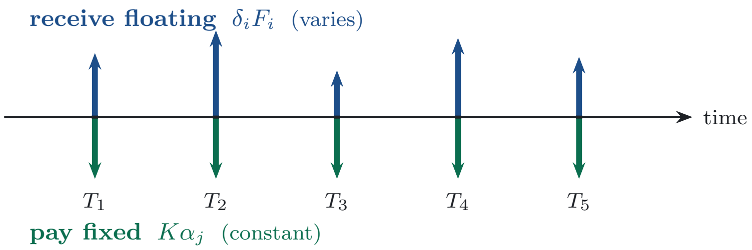

Figure 11: A payer swap's cashflows (notional $N = 1$, not exchanged). At each date the desk receives the floating coupon $\delta_i F_i$ (blue, up—varying, since F_i is not yet known) and pays the fixed coupon $K\alpha_j$ (green, down—constant); only the *net* changes hands. Each floating-versus-fixed period is a miniature FRA, so a swap is a *strip of FRAs*, and its value is their sum.

Vanilla (fixed-for-floating). One leg pays a *fixed* rate K, the other a *floating* index, and only the *net* changes hands each date (Figure 11). On the fixed leg, date T_j pays $K\,\alpha_j\,N$, with α_j the fixed accrual (day-count) fraction; on the floating leg, period $[T_{i-1}, T_i]$ pays $\delta_i\,F_i\,N$, where F_i is that period's floating rate—a LIBOR fixing, or (post-transition) the compounded RFR of §3. Taking $N = 1$ and discounting on the OIS curve P^d, a *payer* swap (pay fixed, receive floating) is the floating leg minus the fixed leg,

$$V_{\text{payer}}(0) \;=\; \sum_i \delta_i\,F_i(0)\,P^d(0,T_i) \;-\; K\sum_j \alpha_j\,P^d(0,T_j). \tag{15}$$

The fixed-leg sum $A = \sum_j \alpha_j\,P^d(0,T_j)$ is the *annuity* (or PV01/DV01): the present value of \$1 of fixed rate paid on the schedule, and hence the swap's sensitivity to the fixed rate—a one-basis-point move in K changes the swap's value by $A \times 1\,\text{bp}$. It is the risk number a swaps desk quotes and hedges. Setting $V_{\text{payer}}(0) = 0$ gives the *par swap rate*

$$S \;=\; \frac{\sum_i \delta_i\,F_i(0)\,P^d(0,T_i)}{\sum_j \alpha_j\,P^d(0,T_j)} \;=\; \frac{\text{floating-leg PV}}{A}, \tag{16}$$

the fixed rate at which the swap is free to enter—the market's fair fixed-for-floating exchange rate for that maturity, and the multi-curve par rate of Equation (10). In the single-curve world the floating leg telescopes to $1 - P(0,T_N)$ (Equation (6)), reducing this to the classic $S = (1 - P(0,T_N))/A$; the OIS par rate of Equation (3) is the same formula with the compounded overnight rate as the floating index. Away from par, a payer swap is worth $V_{\text{payer}} = (S - K)\,A$—the annuity times the rate difference, the cleanest way to see that all a swap's value away from inception lives in how far the market rate has moved from the struck one.

Basis (floating-for-floating). A basis swap exchanges two *floating* legs. The *tenor* basis swap (§8) pays one tenor's index plus a fixed spread against another tenor's: the 3s6s swap of Appendix C pays three-month LIBOR $+\,b$ against six-month LIBOR, and the fair spread b (Equation (18)) is nonzero precisely because the two tenors carry different projection curves. The *cross-currency* basis swap exchanges a floating leg in one currency against a floating leg (plus a basis) in another, *with* an exchange of notionals at start and maturity; its basis is one of the residuals that survive the single-currency collapse of §13.

Conventions. Schedules and day-counts vary by market: a USD vanilla swap historically paid quarterly floating (3M LIBOR, ACT/360) against semi-annual fixed (30/360); a SOFR swap pays annually on both legs, the floating leg compounded in arrears. None of this changes the structure of Equations (15)–(16); it only sets the dates T_j and the accruals α_j, δ_i.

C A basis swap, priced

The 3s6s basis of §8 is the cleanest thing to price in the multi-curve framework, and doing so shows exactly why the tenors need separate curves. A 3s6s basis swap exchanges two floating legs on one notional and maturity: the short-tenor leg pays three-month LIBOR plus a fixed spread b on a quarterly schedule, the long-tenor leg pays six-month LIBOR on a semi-annual schedule. Writing F_i^s and F_j^ℓ for the forwards read off the *short*- and *long*-tenor projection curves, and discounting *both* legs on the one OIS curve P^d, the swap is worth

$$V_{\text{basis}}(0) \;=\; \sum_i \delta_i^s \big(F_i^s + b\big) P^d(0, T_i^s) \;-\; \sum_j \delta_j^\ell \, F_j^\ell \, P^d(0, T_j^\ell). \tag{17}$$

Setting $V_{\text{basis}}(0) = 0$ gives the fair basis as the gap between the two projected legs, divided by the annuity of the spread leg,

$$b \;=\; \frac{\displaystyle\sum_j \delta_j^\ell F_j^\ell P^d(0, T_j^\ell) \;-\; \sum_i \delta_i^s F_i^s P^d(0, T_i^s)}{\displaystyle\sum_i \delta_i^s P^d(0, T_i^s)}. \tag{18}$$

The discount factors are the *same* P^d on both legs, so they do not create the basis—they only weight it. The basis lives entirely in the difference between the two *projection* curves, F^ℓ versus F^s. This is the multi-curve point in a single equation: if one curve projected both tenors, the telescoping no-arbitrage of Equation (5) would force F^ℓ to be the compounded roll of F^s, the numerator of Equation (18) would vanish, and b would be identically zero. A nonzero quoted b *is* the statement that the tenors need separate curves.

A worked number. Take a one-year 3s6s swap with flat curves: a three-month forward of $F^s = 3.00\%$, a six-month forward of $F^\ell = 3.10\%$ (the longer tenor trading 10 bp higher), and an OIS discount curve flat at 2.80% continuously compounded, so $P^d(0,T) = e^{-0.028\,T}$. The short leg pays at $0.25, 0.5, 0.75, 1.0$ ($\delta^s = 0.25$); the long leg at $0.5, 1.0$ ($\delta^\ell = 0.5$). The spread-leg annuity is $\sum_i \delta_i^s P^d(0, T_i^s) = 0.9827$, and the two projected legs are worth 0.02948 and 0.03036. Equation (18) then gives

$$b \;=\; \frac{0.03036 - 0.02948}{0.9827} \;\approx\; 8.9 \text{ bp}.$$

The fair basis is close to the raw 10 bp gap between the tenor forwards, but a shade below it: the six-month leg pays on a coarser, later-weighted schedule, so its discounted annuity is slightly smaller and drags the fair spread down. With no discounting ($P^d \equiv 1$) the two annuities coincide and b collapses to exactly $F^\ell - F^s = 10$ bp—the clean intuition, of which Equation (18) is the discount-weighted version.

D Change of measure: a primer

Risk-neutral pricing values a claim as an expectation of its *discounted* payoff. The freedom that makes rate modelling tractable is that the discounting may be done with *any* traded, strictly positive asset as the yardstick—the *numeraire*—so long as one uses the matching probability measure.

The numeraire is best understood as an *accounting unit*, a measuring stick for value. Cash in the bank account B, a zero-coupon bond, a swap's annuity: each is a self-financing portfolio worth something positive at all times, so any other asset's price can be quoted *in units of* it. The content of risk-neutral pricing is that, measured in the right unit, prices do not drift—they are martingales—so today's price is simply the expected future price *in those units*, with no further discounting. Change the unit and you change the probabilities (the measure), but never the price.

The numeraire theorem. Let $N_t > 0$ be the price of any traded asset. There is a measure $\mathbb{Q}^N$ under which the price of *every* traded asset, expressed in units of N, is a martingale:

$$\frac{V_t}{N_t} = \mathbb{E}^{\mathbb{Q}^N}\Big[\frac{V_T}{N_T}\,\Big|\,\mathcal{F}_t\Big], \qquad \text{equivalently} \qquad V_0 = N_0\,\mathbb{E}^{\mathbb{Q}^N}\Big[\frac{X_T}{N_T}\Big]. \tag{19}$$

The price V_0 is the same under every choice of numeraire (Figure 12); the freedom lies in picking the N under which the particular payoff is simplest.

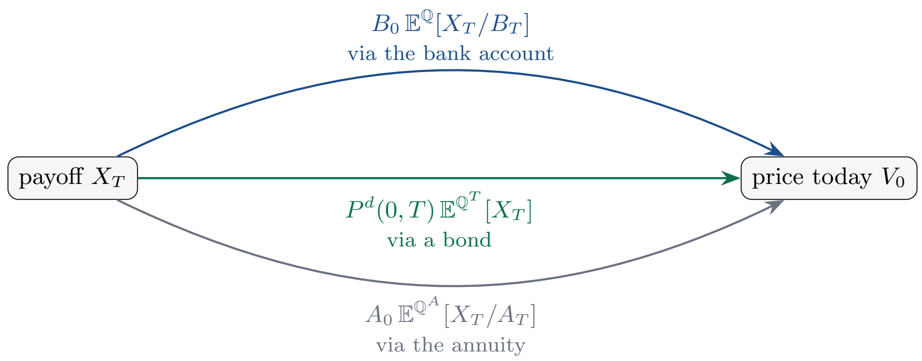

Figure 12: One price, many routes. The value today of a payoff X_T can be computed with *any* numeraire—the bank account (risk-neutral $\mathbb{Q}$), a bond (the T-forward $\mathbb{Q}^T$), the annuity (the swap measure $\mathbb{Q}^A$)—and all three give the *same* V_0. Each route discounts by its numeraire and takes an expectation under the matching measure; the route is chosen so that the payoff is simplest.

The canonical example. Take the simply-compounded forward rate for $[S,T]$, $F(t;S,T) = \frac{1}{\delta}\big(P(t,S)/P(t,T) - 1\big)$. Multiply by the bond: $\delta\,F(t;S,T)\,P(t,T) = P(t,S) - P(t,T)$, a difference of two traded bonds. Divided by the numeraire $P(\cdot,T)$ this is a *traded asset over the numeraire*, hence a $\mathbb{Q}^T$-martingale by Equation (19)—so $F(t;S,T)$ itself is a $\mathbb{Q}^T$-martingale, and $\mathbb{E}^{\mathbb{Q}^T}[F(S;S,T)] = F(0;S,T)$: its expected future value *is* today's forward, with no adjustment. That one line is behind the LIBOR forward of Equation (12), the clean backward-looking forward of Equation (13), and the caplet priced by Black's formula—and any *departure* from its natural numeraire is the convexity of Appendix F.

Switching numeraires. Moving from numeraire N to numeraire M is a change of measure whose Radon–Nikodym derivative is the ratio of the normalised numeraires,

$$\frac{\mathrm{d}\mathbb{Q}^M}{\mathrm{d}\mathbb{Q}^N}\Big|_{\mathcal{F}_t} = \frac{M_t/M_0}{N_t/N_0} = \frac{M_t\,N_0}{N_t\,M_0}. \tag{20}$$

In a diffusion model this reweighting is, by Girsanov's theorem, a *drift* change: a $\mathbb{Q}^N$-Brownian motion becomes a $\mathbb{Q}^M$-Brownian motion plus a drift equal to the volatility of $\ln(M/N)$. Volatilities are untouched; only drifts move. The recipe used throughout the note follows: choose the numeraire under which the quantity you care about is *driftless*—a martingale—so that its expected value is simply its forward.

The switch the note uses: $\mathbb{Q} \to \mathbb{Q}^T$. Take $N = B$ (the bank account, $B_0 = 1$) and $M = P(\cdot, T)$ (the T-bond, $M_0 = P(0,T)$). Equation (20) gives the density

$$\left.\frac{\mathrm{d}\mathbb{Q}^T}{\mathrm{d}\mathbb{Q}}\right|_t = \frac{P(t,T)/P(0,T)}{B_t/B_0} = \frac{P(t,T)}{P(0,T)\,B_t}, \qquad \text{so at } T: \quad \left.\frac{\mathrm{d}\mathbb{Q}^T}{\mathrm{d}\mathbb{Q}}\right|_T = \frac{1}{P(0,T)\,B_T}, \tag{21}$$

using $P(T,T) = 1$. Substituting it into the $\mathbb{Q}^T$ price recovers the identity the note leans on everywhere:

$$P(0,T)\,\mathbb{E}^{\mathbb{Q}^T}[X_T] = P(0,T)\,\mathbb{E}^{\mathbb{Q}}\Big[X_T \cdot \tfrac{1}{P(0,T)\,B_T}\Big] = \mathbb{E}^{\mathbb{Q}}\big[e^{-\int_0^T r_s\,\mathrm{d}s}\,X_T\big].$$

"Discount factor × plain expectation" and "expected discounted payoff" are the same number—you have merely changed the accounting unit. And, by Girsanov, the forward rate that carries a drift under $\mathbb{Q}$ (its expected value is *not* the forward) becomes the driftless $\mathbb{Q}^T$-martingale of the worked example above: switching numeraire is exactly what removes the convexity, or leaves it, of Appendix F.

E The pricing measures used in this note

Three numeraires, and their measures, recur (Table 5); they behave quite differently over the life of a trade (Figure 13).

Measure	Numeraire	Martingale	Used here for
Risk-neutral $\mathbb{Q}$	bank account B_t	V_t/B_t	discounting, compounding
T-forward $\mathbb{Q}^T$	bond $P(\cdot,T)$	T-forwards, forward rates	LIBOR & RFR forwards, caplets
Annuity $\mathbb{Q}^A$	annuity A_t	par swap rate S_t	swaptions, par swap rate

Table 5: The pricing measures used in this note, their numeraires, and what each makes a martingale.

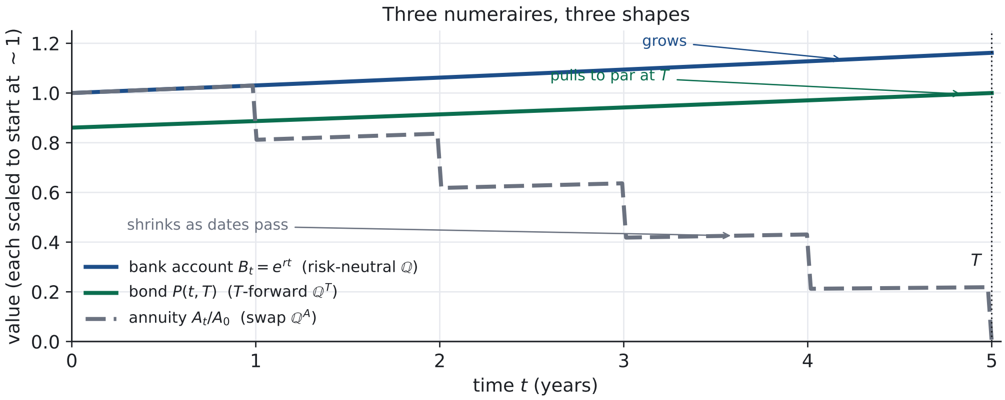

Figure 13: The three numeraires over the life of a trade (flat rate r, five annual dates). The bank account B_t (blue) *grows*; the zero-coupon bond $P(t,T)$ (green) *pulls to par* at T; the annuity A_t (grey, scaled to start at 1) *shrinks* to zero as fixed dates fall away. Three self-financing portfolios, three shapes—and quoting a price in units of each yields the risk-neutral, T-forward and swap measures respectively.

Risk-neutral measure $\mathbb{Q}$. Numeraire: the bank account $B_t = \exp\!\int_0^t r_s\,\mathrm{d}s$—for a collateralised trade, growing at the collateral rate (§5). Pricing is $V_0 = \mathbb{E}^{\mathbb{Q}}[X_T/B_T] = \mathbb{E}^{\mathbb{Q}}[e^{-\int_0^T r}\,X_T]$. This is the measure behind the compounded-RFR bank account of §3 and the discounting identity Equation (4); the discount factor is $P^d(0,T) = \mathbb{E}^{\mathbb{Q}}[1/B_T]$.

T-forward measure $\mathbb{Q}^T$. Numeraire: the zero-coupon bond $P(\cdot, T)$. Since $P(T,T) = 1$, pricing collapses to $V_0 = P^d(0,T)\,\mathbb{E}^{\mathbb{Q}^T}[X_T]$: discount at the bond, then take a *plain* expectation. Under $\mathbb{Q}^T$ the T-forward of any asset is a martingale—in particular the forward rate for a period ending at T. This is the measure that makes the LIBOR forward a martingale (Equation (12)) and gives the backward-looking forward its clean, convexity-free form (Equation (13)). It also prices a *caplet*, an option on that forward, in one line: since the rate is a driftless $\mathbb{Q}^{T_e}$-martingale, caplet $= P^d(0,T_e)\,\mathbb{E}^{\mathbb{Q}^{T_e}}\left[(F-K)^+\right]$, which is Black's formula on the forward F. The note uses the T_e-forward measure $\mathbb{Q}^{T_e}$ for a coupon paid at T_e.

Annuity (swap) measure $\mathbb{Q}^A$. Numeraire: the fixed-leg annuity $A_t = \sum_j \alpha_j\, P^d(t,T_j)$ of Appendix B. Under $\mathbb{Q}^A$ the par swap rate S_t of Equation (16) is a martingale, so a payer swaption is worth $A_0\,\mathbb{E}^{\mathbb{Q}^A}[(S_T-K)^+]$—the natural measure for swap-rate products, and the analogue for a whole swap of what $\mathbb{Q}^T$ is for a single forward.

All three are linked by Equation (20): moving from $\mathbb{Q}$ to $\mathbb{Q}^T$, for instance, reweights by $B_0\,P(t,T)/(B_t\,P(0,T))$, which is exactly what turns the discounted expectation $\mathbb{E}^{\mathbb{Q}}[e^{-\int r}X]$ into $P^d(0,T)\,\mathbb{E}^{\mathbb{Q}^T}[X]$.

What the transition changes

A measure is defined by its numeraire, so the three measures above are the *same* abstract objects before and after LIBOR. What the switch to risk-free rates changes is *which curve instantiates each numeraire*—and it brings them into alignment. Three points.

The short rate becomes real. The risk-neutral numeraire $B_t = \exp\!\int_0^t r_s\,\mathrm{d}s$ has always compounded an idealised, unobservable *risk-free* short rate r. Before, no traded rate matched it: the liquid term rate, LIBOR, carried bank credit and a term premium, so the model's r and the market's benchmark were different animals. An overnight RFR is, to a basis point, that risk-free short rate made *tradable*—B is now built by compounding a rate the market actually fixes each day (§3)—and, because the RFR is also the collateral rate (§5), the risk-neutral pricing measure and the discounting are one object, not two that merely nearly coincided.

The forward-measure martingale becomes *derived*, not *assumed*. The bond $P(\cdot,T)$ that defines $\mathbb{Q}^T$ is now unambiguously the OIS/RFR discount bond. Under it, the LIBOR forward was a martingale only by modelling convention—LIBOR is not the discount curve, so the clean behaviour of Equation (12) had to be *posited*. The compounded-RFR forward, by contrast, is a martingale by *construction*: compounding the overnight rate rebuilds the numeraire B, and a ratio of B to a bond is automatically a $\mathbb{Q}^{T_e}$-martingale (Equation (13)). The transition thus upgrades the forward-measure result from an assumption to a consequence—and, by Appendix F, removes the convexity any such mismatch would otherwise carry.

The three numeraires collapse onto one rate. For an RFR product the same overnight rate is the short rate (numeraire B), the discount rate (the bond P^d, hence the annuity A) *and*, through compounding, the projected index. So $\mathbb{Q}$, $\mathbb{Q}^T$ and $\mathbb{Q}^A$ are all built from a single curve—the near single-curve world of §13—rather than from the LIBOR-era tangle of a discount curve and a separate projection curve per tenor. The measures do not change in kind; they become *mutually consistent*. What stays genuinely multi-curve, keeping these measures distinct across markets, is the residue of §13: cross-currency basis and the term-versus-compounded basis.

F Convexity adjustments: a primer

A *convexity adjustment* is what you add to a forward rate when you pay or observe it somewhere other than its *natural* place. Every forward rate is a martingale under exactly one measure—the one whose numeraire matches its natural payment (Appendix E)—and under that measure its expected value is simply today's forward. Price it under any *other* measure and the expectation shifts; the shift is the convexity adjustment.

The exact statement. Let X be a rate that is a martingale under its natural measure $\mathbb{Q}^N$, so $\mathbb{E}^{\mathbb{Q}^N}[X_T] = X_0$ (today's forward). Suppose the actual payoff instead discounts under a *different* measure $\mathbb{Q}^M$. With the change-of-measure density of Appendix D, $Z_T \equiv \mathrm{d}\mathbb{Q}^M/\mathrm{d}\mathbb{Q}^N\big|_T$ and $\mathbb{E}^{\mathbb{Q}^N}[Z_T] = 1$,

$$\underbrace{\mathbb{E}^{\mathbb{Q}^M}[X_T] - X_0}_{\text{convexity adjustment}} \;=\; \mathbb{E}^{\mathbb{Q}^N}[X_T Z_T] - X_0 \;=\; \mathrm{Cov}^{\mathbb{Q}^N}\big(X_T,\, Z_T\big). \tag{22}$$

The adjustment is *exactly* the covariance, under the natural measure, between the rate and the change-of-measure density—that is, between the rate and the ratio of the two numeraires. Three things read straight off: it vanishes if the rate and the numeraire ratio are uncorrelated or if either has no volatility; its *sign* follows that correlation; and its *size* is second order—roughly $X_0\,\rho\,\sigma_X\,\sigma_Z\,T$, a product of volatilities. That variance-like character is what the name "convexity" refers to.

Equation (22) has a simple picture (Figure 14). Under its natural measure the rate is centred on today's forward X_0. Pricing under $\mathbb{Q}^M$ *reweights* that distribution by the density Z_T: outcomes in which the numeraire ratio is large receive more probability. If the rate tends to be high in exactly those outcomes (positive correlation), the reweighted mean shifts *up*, and that shift is the adjustment; uncorrelated, the tilt is symmetric and the mean does not move.

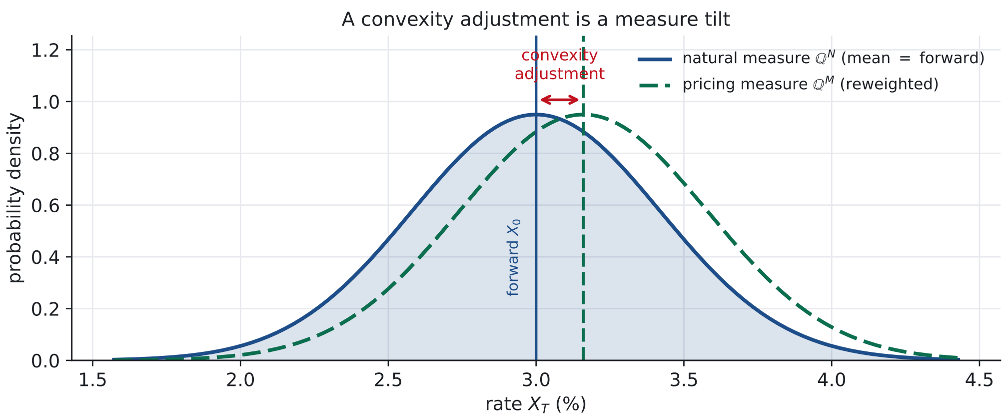

Figure 14: A convexity adjustment is a measure tilt. Under its natural measure $\mathbb{Q}^N$ the rate X_T is centred on the forward X_0 (blue). Pricing under a different measure $\mathbb{Q}^M$ reweights the distribution by the density Z_T (green); when the rate and Z_T are positively correlated the mean shifts up, and the shift—the covariance of Equation (22)—is the convexity adjustment (illustrative Gaussian).

Where it shows up. The canonical cases are all this one mechanism with different numeraires:

- **Payment delay / timing.** A rate natural to T_e but paid at $T_p \neq T_e$ is valued under $\mathbb{Q}^{T_p}$, not $\mathbb{Q}^{T_e}$; the density is a ratio of bonds, and the adjustment is the small timing correction of §9.

- **In-arrears / in-advance.** Paying a period's rate at its *start* instead of its end (LIBOR-in-arrears; the "setting in advance" of §9) is the same mismatch, and yields the classic in-arrears convexity.
- **CMS.** A constant-maturity-swap payment references a swap rate S paid at a single date rather than through its annuity. S is a martingale under the annuity measure $\mathbb{Q}^A$ (Appendix E); pay it under a $\mathbb{Q}^T$ and a (larger) CMS convexity adjustment appears.
- **Futures vs. forwards.** A futures contract is margined daily, making the futures rate a martingale under the risk-neutral $\mathbb{Q}$; the forward is a martingale under $\mathbb{Q}^T$. The gap is the futures/forward convexity—why a rate future and the matching FRA do not quote at the same rate.
- **Quanto.** Paying a rate in the "wrong" currency multiplies the numeraire by an FX factor; the adjustment is a covariance with the exchange rate.

A worked case: paying in arrears. The cleanest example is a rate paid at its own *reset* rather than at period end—LIBOR-in-arrears, the "setting in advance" of §9. The forward L for $[T_e, T]$ is a $\mathbb{Q}^T$-martingale, but paid at T_e it is valued under $\mathbb{Q}^{T_e}$. The density is $Z \propto P(T_e, T_e)/P(T_e, T) = 1 + \delta L_{T_e}$—it *rises* with the rate—so $\mathrm{Cov}^{\mathbb{Q}^T}(L, Z) = \delta\,\mathrm{Var}^{\mathbb{Q}^T}(L) > 0$, and to leading order

$$\mathbb{E}^{\mathbb{Q}^{T_e}}[L_{T_e}] \;\approx\; L_0\left(1 + \delta\, L_0\, \sigma_L^2\, T_e\right), \tag{23}$$

a *positive* adjustment growing with the rate's accumulated variance $\sigma_L^2 T_e$. The same template, with the annuity ratio for Z, gives the (larger) CMS adjustment; and a payment *later* than the natural date flips the correlation's sign, giving the small negative payment-delay adjustment of §9.

Application to the compounded RFR. The backward-looking compounded RFR paid at T_e is the rare case with *no* adjustment: its natural measure is $\mathbb{Q}^{T_e}$ and it is paid at T_e, so $M = N$, the density Z_T is 1, and the covariance in Equation (22) is zero (§9). Convexity reappears only when the timing is disturbed—payment delay, setting in advance, or approximating the realised rate by a forward-looking term rate—precisely the departures listed there.

G A minimal implementation

The whole coupon calculation is a few lines. Given arrays of overnight fixings and their day-count fractions, compound in arrears, then add a fixed spread to form the fallback rate of Equation (14).

```
import numpy as np

def compounded_in_arrears(rates, dcf):
    """Annualised compounded RFR over one accrual period.
    rates : overnight fixings (decimal, e.g. 0.0505)
    dcf   : day-count fraction per fixing (1/360, or 3/360 over a weekend)
    """
    growth = np.prod(1.0 + rates * dcf)      # = B_end / B_start
    accr   = dcf.sum()                        # delta_{s,e}
    return (growth - 1.0) / accr

def fallback_rate(rates, dcf, spread_bp):
    """Compounded RFR + fixed ISDA spread adjustment (in bp)."""
    return compounded_in_arrears(rates, dcf) + spread_bp * 1e-4
```

```python
def lookback(rates, dcf, k):
    """k-business-day lookback WITH observation shift: use the rate and
    the day-count from k days earlier (arrays pre-aligned by caller)."""
    return compounded_in_arrears(np.roll(rates, k), np.roll(dcf, k))
```

A three-month USD coupon on fixings around 5% returns a compounded rate near 5%; adding the 3M spread of Table 4 (`spread_bp = 26.161`) gives the fallback rate a legacy LIBOR contract would pay.

H An end-to-end example: from OIS quotes to a fallback coupon

A single short calculation draws the strands of the note together: start from two market quotes, build the discount curve (§4, §8), read off a forward (§9), apply the fallback spread (§11), and present-value the coupon a legacy contract will actually pay. The example is deliberately minimal: two quotes, annual and semi-annual simplifications, one coupon.

Step 1 — the quotes. The OIS market shows a 6-month par rate $K_{0.5} = 4.90\%$ and a 12-month par rate $K_1 = 5.00\%$ (each, for simplicity, a single payment at maturity with accruals $\alpha = 0.5$ and $\alpha = 1$).

Step 2 — bootstrap the discount curve. Equation (9) with one payment reduces to $P^d = 1/(1 + K\alpha)$:

$$P^d(0, \tfrac{1}{2}) \;=\; \frac{1}{1 + 0.5 \times 0.0490} \;=\; 0.97609, \qquad P^d(0, 1) \;=\; \frac{1}{1 + 0.0500} \;=\; 0.95238.$$

Step 3 — the forward compounded rate. The fair forward of the compounded RFR over $[\tfrac{1}{2}, 1]$ is the clean, convexity-free ratio of Equation (13):

$$R\big(0; \tfrac{1}{2}, 1\big) \;=\; \frac{1}{0.5}\left(\frac{0.97609}{0.95238} - 1\right) \;=\; 4.9780\%.$$

No model was needed—only the discount curve. This is the multi-curve collapse of §8 in action: for the RFR itself, projection *is* discounting.

Step 4 — apply the fallback spread. A legacy contract that paid 6-month USD LIBOR on this period falls back, under the ISDA protocol (§12), to the compounded RFR plus the fixed 6M spread of Table 4, $s^* = 42.826$ bp:

$$\text{fallback coupon rate} \;=\; 4.9780\% + 0.42826\% \;=\; 5.4063\%.$$

Step 5 — present value. The coupon pays $\delta \times \text{rate} = 0.5 \times 5.4063\%$ at $T = 1$, worth today

$$0.5 \times 0.054063 \times P^d(0, 1) \;=\; 0.5 \times 0.054063 \times 0.95238 \;=\; 0.025744$$

per unit notional—about 2.574% of notional. Each of the five steps rests on a different part of the note: the OIS par-rate identity, the bootstrap, the martingale property of the backward-looking forward, the fixed spread, and OIS discounting. A production system differs only in scale—hundreds of quotes, splined interpolation, joint calibration—not in kind.

I Glossary

IBOR Interbank Offered Rate—the family (LIBOR, EURIBOR, ...) of survey-based term rates being retired.

RFR Risk-free rate—the overnight, transaction-based replacement (SOFR, SONIA, €STR, TONA, SARON).

OIS Overnight-indexed swap—a swap of a fixed rate against compounded overnight RFR; its curve is the discount curve.

CSA Credit Support Annex—the collateral agreement whose collateral rate sets the discount curve.

OIS / CSA discounting Discounting a collateralised trade at the rate its collateral earns—an overnight (OIS) rate—so the discount curve is the OIS curve (§5).

In arrears / in advance A rate known at the *end* of its period (compounded RFR) versus at the *start* (LIBOR, term RFR).

Lookback / lockout / payment delay Conventions (§7) that restore notice on an in-arrears coupon.

Term SOFR / Term SONIA Forward-looking term rates built from RFR derivatives, set in advance, permitted in restricted uses.

Spread adjustment The fixed, five-year-median credit spread added to a compounded RFR to make a fallback fair (§11).

Fallback The contractual replacement of LIBOR by adjusted RFR on a trigger (§12).

Pre-cessation trigger A fallback triggered by loss of representativeness, not just by permanent cessation.

Synthetic LIBOR A temporary rate published under the LIBOR name but computed as term RFR plus spread, for tough legacy contracts.

Tenor basis The spread between forwards implied by different LIBOR tenors (3s6s, ...), a symptom of the multi-curve world.

The "big bang" The October 2020 switch of cleared-swap discounting from Fed Funds/EONIA to SOFR/€STR.

ARRC / ISDA / CCP The US Alternative Reference Rates Committee; the derivatives trade body that authored the fallbacks; central clearing counterparties.

References

[1] A. Bailey. The future of LIBOR. Speech, Financial Conduct Authority, 27 July 2017.

[2] Financial Stability Board. *Reforming Major Interest Rate Benchmarks*. FSB, 2014.

[3] M. Bianchetti. Two curves, one price: pricing and hedging interest rate derivatives decoupling forwarding and discounting yield curves. *Risk*, 2010.

[4] M. Henrard. *Interest Rate Modelling in the Multi-Curve Framework: Foundations, Evolution and Implementation*. Palgrave Macmillan, 2014.

[5] A. Lyashenko, F. Mercurio. Looking forward to backward-looking rates: a modeling framework for term rates replacing LIBOR. *Risk* / SSRN working paper, 2019.

[6] F. Mercurio. A simple multi-curve model for pricing SOFR futures and other derivatives. SSRN working paper, 2018.

[7] Alternative Reference Rates Committee. *SOFR Floating Rate Notes Conventions Matrix* and *Recommendations regarding Term SOFR*. ARRC, 2021.

[8] International Swaps and Derivatives Association. *IBOR Fallbacks Supplement and Protocol*. ISDA, 2020.

[9] Bloomberg Index Services / ISDA. *IBOR Fallback Rate Adjustments—Spread Fixing Technical Notice*, 5 March 2021.

[10] D. Brigo, F. Mercurio. *Interest Rate Models—Theory and Practice*. Springer, 2nd edition, 2006.

[11] L. B. G. Andersen, V. V. Piterbarg. *Interest Rate Modeling*. Atlantic Financial Press, 2010.