Chapter 7

Curve Construction and Linear Products

In these notes we build the interest rate curve out of the instruments quoted in the market, and price the products whose value depends on the curve alone and not on its volatility. This is the layer underneath every model in the chapters that follow: a term structure model is calibrated to a curve, and a curve is not handed down, it is constructed.

7.1 Discount Factors, Zero Rates and Forward Rates

The single object underneath all of fixed income is the discount factor.

Definition 7.1 (Discount factor)**.** $P(t,T)$ is the value at time t of a contract paying \$1 with certainty at time T. It is the price at t of a zero coupon bond maturing at T, and $P(T,T)=1$.

Everything else in this chapter is a re-expression of the function $T \mapsto P(0,T)$. The reason interest rates are usually harder to model than equities is because the price of a stock today is a number, whereas the price of borrowing today is a whole curve, one number for every maturity. A model of interest rates has to move a curve, and the constraints linking the points of that curve to each other are what make the subject what it is.

The first such constraint is the one that lets us talk about rates at all.

Definition 7.2 (Simply compounded forward rate)**.** For $T_1 < T_2$ with year fraction $\tau = \tau(T_1,T_2)$, the forward rate $F(t;T_1,T_2)$ is the interest rate, agreed at t, at which one may borrow over the future period $[T_1,T_2]$.

The forward rate is not free to be anything. It is pinned by today's discount factors, and it should be derived rather than quoted, because the derivation is the pattern every no-arbitrage statement in this chapter follows.

Calculation 7.3 (The forward rate from the curve)**.** Consider the following portfolio, set up at time 0 at no cost. Buy one zero coupon bond maturing at T_1, and sell $P(0,T_1)/P(0,T_2)$ zero coupon bonds maturing at T_2. The proceeds of the sale exactly fund the purchase, so

$$\pi_0 = P(0,T_1) - \frac{P(0,T_1)}{P(0,T_2)} \cdot P(0,T_2) = 0.$$

At T_1 the long bond pays \$1. At T_2 the short position costs $P(0,T_1)/P(0,T_2)$. So the portfolio is exactly a contract to receive \$1 at T_1 and repay $P(0,T_1)/P(0,T_2)$ at T_2: it is a loan over $[T_1,T_2]$, entered into today at no cost. Its simple interest rate F satisfies $1+\tau F = P(0,T_1)/P(0,T_2)$, giving

$$F(0;T_1,T_2) = \frac{1}{\tau}\left(\frac{P(0,T_1)}{P(0,T_2)} - 1\right). \tag{7.1}$$

Any other rate for the same period admits an arbitrage: if the market quoted $F' > F$, lend at F' and fund it with this portfolio, and the difference is locked in at T_2 with no position of any kind held in between.

Note what the argument used and what it did not. It used only that the two bonds trade, and it produced a rate for a period entirely in the future without any model of how rates move. This is the sense in which linear products are "just discounting": their prices are consequences of today's curve, not of any assumption about its dynamics.

The same quantity gets quoted in several conventions, and the conversions between them account for a surprising share of the discrepancies between two people who think they are pricing the same trade.

Definition 7.4 (Zero rate)**.** The continuously compounded zero rate, or spot rate, for maturity T is the single rate $z(T)$ that reproduces the discount factor:

$$P(0,T) = e^{-z(T)T} \qquad \Longleftrightarrow \qquad z(T) = -\frac{\ln P(0,T)}{T}.$$

Definition 7.5 (Instantaneous forward rate)**.** The instantaneous forward rate $f(t,T)$ is the rate for borrowing over $[T, T+dT]$, quoted at t: entering at t into a contract to borrow from T to $T+dT$ charges interest $f(t,T)dT$.

Remark. Taking $T_1 = T$, $T_2 = T+\delta$ in (7.1) and letting $\delta \to 0$,

$$f(t,T) = \lim_{\delta\to 0} \frac{1}{\delta}\left(\frac{P(t,T)}{P(t,T+\delta)} - 1\right) = -\frac{\partial}{\partial T} \ln P(t,T),$$

and integrating from t to T using $P(t,t) = 1$,

$$P(t,T) = e^{-\int_t^T f(t,s)\,ds}. \tag{7.2}$$

So the discount curve, the zero curve and the instantaneous forward curve carry exactly the same information. Which of them one works in is a matter of what one wants to be smooth — a point we return to when we interpolate.

Definition 7.6 (Short rate)**.** The short rate is the instantaneous forward rate for immediate borrowing, $r(t) = f(t,t)$.

7.2 The Year Fraction Is Not a Length of Time

Everything above was written with a year fraction $\tau(T_1, T_2)$ in it, and the natural reading of that symbol — elapsed time — is wrong. It is whatever the contract says it is. This section is about the arithmetic that follows from that.

Definition 7.7 (Day count convention)**.** A day count convention is a rule assigning a year fraction to a pair of dates. The three in common use are

- *Act/360*: actual days elapsed, divided by 360. The money market convention, used for USD and EUR floating legs.
- *Act/365*: actual days over 365, used in sterling and several Asian markets.
- *30/360*: every month is treated as thirty days and every year as 360, so a semiannual period is exactly one half. Used for US corporate and agency bonds and for USD swap fixed legs.

A single swap routinely uses two of them, one per leg.

Calculation 7.8 (What the conventions are worth)**.** Take the semiannual period from 15 January to 15 July 2026, which is 181 actual days.[1]

Convention	Year fraction	Interest at 5%
Act/360	0.502778	2.5139%
Act/365	0.495890	2.4795%
30/360	0.500000	2.5000%

Act/360 against 30/360 is a gap of 1.39 basis points on one period. A swap's fixed leg quoted on the wrong basis is wrong by that on every payment, compounding to something no bid-offer spread absorbs. Act/360 is systematically the largest.

The same arithmetic makes a leap day worth 1.39 basis points on an annual Act/360 period, for the identical reason — both numbers are one day out of 360. Which is the useful way to remember the size of all of this: a day is roughly a basis point and a half at these levels, and every convention question is a question about how many days.

Two further conventions are needed before an actual cashflow can be computed, and both are about dates rather than fractions.

Definition 7.9 (Settlement lag)**.** A trade agreed today does not settle today. A spot-starting swap typically begins two business days forward, a bond purchase settles one or two business days after the trade, and a deposit rate quoted today applies from spot. So the curve built from today's quotes has its origin at spot rather than at today, and a discount factor to a date is a discount factor from settlement.

The lag is small and it does not cancel. A two day gap at 5% is a basis point and a half of value on the whole notional, so a curve whose origin is misplaced by two days misprices every instrument on it by about that, in the same direction, in a way that no repricing check will reveal because the same error is in both the build and the valuation.

[1] `curve::DayCount`, with the year fractions and both gaps below pinned.

Definition 7.10 (Accrued interest, clean and dirty price)**.** A bond that has run part way through a coupon period has earned interest that has not yet been paid. With coupon rate c and τ the year fraction since the last payment, on the bond's own convention,

$$\text{accrued} = 100\, c\, \tau,$$

and the two prices are

$$\text{dirty} = \text{clean} + \text{accrued}.$$

The *dirty* price is what changes hands. The *clean* price is what is quoted.

Remark (Why the quote is the one that is not paid)**.** The dirty price of a bond sawtooths: it climbs by the accrual through each period and drops by the coupon on payment day, so a chart of it looks like the bond is repeatedly falling for reasons unconnected to rates. Removing the accrual removes the sawtooth, and what is left moves only when the market moves. The convention exists to make the quoted series interpretable.

It also produces a specific, common error. A position marked at the clean price is undervalued by up to a full coupon, and a hedge ratio computed from clean prices on both legs is right, while one computed from a clean price on one leg and a dirty on the other is wrong by the accrual — which is a number that grows steadily and then jumps, so the error looks like a slow drift in the P&L punctuated by a step. Chapter 24's explain report is where that pattern shows up, and recognising it as an accrual convention rather than a model failure saves a great deal of time.

7.3 The Instruments

The curve is built from the instruments the market quotes, so we need those first. All of them are linear in the sense that matters here: their payoffs are fixed multiples of rates or of \$1, with no optionality, so their prices are determined by today's curve.

Definition 7.11 (Deposit)**.** A deposit lends \$1 from T_1 to T_2 at a simply compounded rate R, returning $1 + \tau R$ at T_2. Its value today is $P(0, T_2)(1 + \tau R) - P(0, T_1)$, and it is at par when $R = F(0; T_1, T_2)$.

Definition 7.12 (Reference rate)**.** A reference rate $L(t, T, T + \delta)$ is the market's published rate for borrowing over $[T, T + \delta]$, as observed at t. Historically these were IBOR rates — panel estimates of unsecured interbank borrowing costs, of which LIBOR was the largest — and they are now overnight risk free rates such as SOFR, compounded over the period.

In a single curve world the reference rate is the curve's own forward rate,

$$L(t, T, T + \delta) = \frac{1}{\delta}\left(\frac{P(t, T)}{P(t, T + \delta)} - 1\right),$$

which is the assumption we make until the last section of this chapter, where we take it apart.

Definition 7.13 (Forward rate agreement)**.** A δ FRA maturing at T pays at T

$$\frac{\delta\big(L(T, T, T + \delta) - R\big)}{1 + \delta L(T, T, T + \delta)},$$

where R is the rate agreed in the contract. The forward rate is the R making this contract worth zero.

The peculiar looking discounting inside the payoff is the market convention of settling at the start of the period rather than the end: the natural payment $\delta(L - R)$ would fall at $T + \delta$, and dividing by $1 + \delta L$ moves it back to T at the rate that has just fixed. Chapter 8 shows that the fair R is the forward rate under the appropriate measure, and that the corresponding futures contract is not, which is the first place a model becomes necessary.

Interest rate swaps

Definition 7.14 (Swap)**.** A fixed-for-floating interest rate swap exchanges a stream of fixed coupons for a stream of floating coupons on the same notional. The payer of the swap pays fixed and receives floating; the receiver does the opposite. The two legs need not share a payment frequency.

Take the notional to be 1 throughout; prices scale with it. Let the fixed leg pay at dates $\tau_1^r < \cdots < \tau_{N_r}^r$ with accruals δ_i^r, and the floating leg at $\tau_1^p < \cdots < \tau_{N_p}^p$ with accruals δ_j^p, both legs starting at τ_0. The fixed rate is R.

The fixed leg is a known set of cashflows, so it prices by discounting:

$$\text{Fixed}(0) = \Sigma_{i=1}^{N_r} \delta_i^r R\, P(0, \tau_i^r) = R\, \text{Ann}(0),$$

where we have named the sum.

Definition 7.15 (Annuity)**.** The annuity, or level, or PV01, is the value of \$1 paid on each fixed leg date, weighted by the accruals:

$$\text{Ann}(t) = \Sigma_{i=1}^{N_r} \delta_i^r P(t, \tau_i^r).$$

The annuity will do a great deal of work later. It is a price — it is the price of a portfolio of zero coupon bonds — and it is strictly positive, so by chapter 6 it is a numeraire, and it is the numeraire in which swap rates are martingales.

The floating leg takes one more step. Each coupon pays the reference rate fixed at the start of its period, so

Calculation 7.16 (The floating leg telescopes)**.**

$$\begin{aligned}
\text{Float}(0) &= \Sigma_{j=1}^{N_p} \delta_j^p\, L(0, \tau_{j-1}^p, \tau_j^p)\, P(0, \tau_j^p) \\
&= \Sigma_{j=1}^{N_p} \delta_j^p \cdot \frac{1}{\delta_j^p} \left(\frac{P(0, \tau_{j-1}^p)}{P(0, \tau_j^p)} - 1 \right) P(0, \tau_j^p) \\
&= \Sigma_{j=1}^{N_p} \big(P(0, \tau_{j-1}^p) - P(0, \tau_j^p) \big) \\
&= P(0, \tau_0) - P(0, \tau_{N_p}^p).
\end{aligned}$$

Every intermediate term cancels, and a floating leg is worth notional in at the start minus notional out at the end. There is a reason beyond the algebra: a floating rate note that pays the market rate on its own notional is worth par at each reset, because at the reset date the coupon is exactly the rate the market would charge for the next period. Note carefully that the cancellation used L being the curve's forward rate — the same curve that produced the discount factors. In the multi-curve world it is false.

Definition 7.17 (Par swap rate)**.** The par swap rate S is the fixed rate making the swap worth zero.

Setting Fixed = Float,

$$S(t) = \frac{\text{Float}(t)}{\text{Ann}(t)} = \frac{P(t,\tau_0) - P(t,\tau^p_{N_p})}{\Sigma_i \delta^r_i P(t,\tau^r_i)}, \tag{7.3}$$

The numerator is a floating leg of many payments collapsed to a difference of two discount factors, and the collapse is exact rather than approximate: each payment's forward is computed from the same curve that discounts it, so every term's numerator cancels the next term's denominator. It also does not require the two legs to share a frequency, which is why one formula covers a swap paying quarterly floating against annual fixed. and the value of a receiver swap struck at R is $\text{Ann}(0)(R - S(0))$, of a payer swap $\text{Ann}(0)(S(0) - R)$. Two facts fall straight out of this and both matter later. First, the swap rate is a ratio of a price to the annuity, so it is a martingale in the annuity measure — this is the observation that makes swaptions tractable in chapter 8. Second, the sensitivity of a swap to a parallel shift in the fixed rate is the annuity, so the annuity is the swap's risk as well as its numeraire.

Bonds

Definition 7.18 (Bond)**.** A bond pays a specified notional at maturity and periodic coupons until then. A zero coupon bond pays only the notional.

A bond is a set of known cashflows c_i at dates t_i, so its value from the curve is $\Sigma_i c_i P(0,t_i)$. This is the *dirty* or invoice price, and is what changes hands. What is quoted is the *clean* price, the dirty price less the accrued interest earned since the last coupon.

The market also quotes a bond by its yield.

Definition 7.19 (Yield to maturity)**.** The yield y of a bond is the single rate that reprices it:

$$P_{\text{dirty}} = \Sigma_i \frac{c_i}{(1 + y/m)^{mt_i}}, \tag{7.4}$$

with m the coupon frequency. Note which price appears: the yield is defined against the dirty price, since it is the amount actually paid that the discounted cashflows have to reproduce.

The yield is a summary statistic and not a valuation model. It collapses the entire curve onto one number, so two bonds of different maturities have yields that are not comparable in any useful way, and this is exactly why spreads to a curve exist. It is however the natural coordinate for risk.

Definition 7.20 (Duration, DV01 and convexity)**.**

$$D_{\text{mod}} = -\frac{1}{P}\frac{\partial P}{\partial y}, \qquad \text{DV01} = -\frac{\partial P}{\partial y} \times 10^{-4}, \qquad C = \frac{1}{P}\frac{\partial^2 P}{\partial y^2},$$

so that $\Delta P \approx -D_{\text{mod}} P\,\Delta y + \frac{1}{2} C P (\Delta y)^2$.

Exercise (Duration of a zero coupon bond)**.** Show that a zero coupon bond maturing at T, priced with continuous compounding, has $D_{\text{mod}} = T$, and that a coupon bond's duration is the coupon-weighted average of its cashflow times. Deduce that duration is bounded above by maturity, with equality only for a zero coupon bond.

Remark (Two different DV01s, and they are not equal)**.** The definition above differentiates with respect to y, the bond's own yield to maturity, and that is one of two things the same name is used for.

A *yield* DV01 shifts the single number y and reprices with the bond's own discounting. A *curve* DV01 shifts the zero curve — every point of it, by a basis point — rebuilds the discount factors and reprices. They differ: y is defined by (7.4) as the flat rate reproducing the bond's price, so a bond whose cashflows sit on a sloped curve has a y that is a cashflow-weighted average of the curve, and shifting the curve by a basis point does not shift that average by exactly a basis point unless the weights are uniform. The gap grows with the coupon and with the slope, and it is zero for a zero coupon bond, where the two definitions coincide.

Which to use is decided by what the hedge is. A position hedged with another bond wants yield DV01 on both, because the ratio is what is traded. A position hedged with swaps, or held in a book whose risk is bucketed against curve instruments, wants curve DV01, because that is the sensitivity the hedging instruments actually deliver. Reporting one and hedging with the other leaves a residual that looks like a small unexplained P&L which grows when the curve steepens — and the curve risk section below is where that becomes systematic.

Convexity is subject to the same distinction and matters for a further reason. It is why the DV01 itself moves as rates move, so a hedge set today is wrong tomorrow by an amount proportional to C and to the size of the move. That is a rehedging cost rather than a mark, and it is the linear-product version of the gamma cost of chapter 5.

Exercise (A swap's DV01 is its annuity)**.** Show from (7.3) that the change in value of a payer swap for a one basis point rise in the par swap rate is $\mathrm{Ann}(0) \times 10^{-4}$. This is why traders quote swap risk in annuities and why the annuity is called the PV01.

7.4 Bootstrapping the Curve

We now have the two halves of the problem. The instruments above price off the curve; the market quotes their prices. Curve construction is the inverse problem: find the curve that reprices the quotes.

The instruments are chosen by liquidity, and different parts of the curve are pinned by different ones. The front end comes from overnight and short deposits and from overnight indexed swaps. The belly comes from short term interest rate futures, or from FRAs, which fix forward rates out to two or three years. Beyond that the quotes are par swap rates, out to thirty years and more.

The mental model is to solve outward. The shortest instrument involves only the nearest discount factor, so it determines it. The next instrument, given everything shorter, involves one new discount factor, so it determines that. And so on.

Example 7.1 (Bootstrapping by hand)**.** Suppose the market quotes a one year deposit at 4.00%, a two year annual par swap at 3.80%, and a three year annual par swap at 3.70%, all with $\tau = 1$ and annual fixed payments.

The deposit gives the first discount factor directly:

$$P(0,1) = \frac{1}{1 + 0.0400} = 0.961538.$$

The two year swap is at par, so by (7.3) with $S = 0.0380$,

$$\begin{aligned} 1 - P(0,2) &= 0.0380\,\big(P(0,1) + P(0,2)\big) \\ P(0,2)\big(1 + 0.0380\big) &= 1 - 0.0380 \cdot 0.961538 \\ P(0,2) &= \frac{1 - 0.036538}{1.0380} = 0.928190. \end{aligned}$$

The three year swap then gives, with $S = 0.0370$,

$$\begin{aligned} 1 - P(0,3) &= 0.0370\,\big(P(0,1) + P(0,2) + P(0,3)\big) \\ P(0,3) &= \frac{1 - 0.0370 \cdot (0.961538 + 0.928190)}{1.0370} = 0.896895. \end{aligned}$$

Reading off the continuously compounded zero rates and the simply compounded one year forward rates between the nodes,

	1y	2y	3y
$z(T)$	3.922%	3.726%	3.627%
forward over the year to T	4.000%	3.593%	3.489%

Notice how much further the forwards travel. Over three years the zero rate falls by 30 basis points and the forward rate by 51, nearly twice as far. This is not an accident of these numbers. The zero rate is an average of the forwards up to T, and an average always moves less than the thing being averaged, so a curve that slopes gently in zero space slopes sharply in forward space. Since the forwards are what a product actually fixes on, and the zeros are what one is tempted to draw and interpolate, this is a warning about what comes next.

Exercise (Forwards as the derivative of the zero curve)**.** Show that $f(0,T) = z(T) + Tz'(T)$, and deduce that a zero curve with a maximum at T^* has $f(0,T^*) = z(T^*)$, with the forward curve crossing the zero curve there. Confirm on the example above that a falling zero curve forces the forwards below it.

In practice one does not literally peel the instruments off one at a time. The instruments overlap, their dates do not line up, and the front end has effects a strict ordering cannot express. What is done instead is to set up a global solve: choose a set of curve nodes and an interpolation scheme, and find the node values that reprice every quoted instrument simultaneously. The bootstrap above is the special case in which the system happens to be triangular.

Interpolation is a modelling choice

Whatever the instruments are, there are finitely many of them and infinitely many dates, so the curve between the nodes has to be interpolated, and the choice is not innocent.

Suppose we interpolate linearly in the continuously compounded zero rate. Then on a segment between nodes at T_a and T_b,

$$z(T) = z(T_a) + \frac{T - T_a}{T_b - T_a}\big(z(T_b) - z(T_a)\big),$$

and since $\ln P(0,T) = -z(T)T$, the instantaneous forward rate is

$$f(0,T) = -\frac{\partial}{\partial T}\ln P(0,T) = \frac{\partial}{\partial T}\big(z(T)T\big) = z(T) + Tz'(T)$$
$$= z(T) + T \cdot \frac{z(T_b) - z(T_a)}{T_b - T_a}.$$

This is piecewise linear in T but *discontinuous at every node*, because z' jumps there while T does not. So a perfectly reasonable looking zero curve, smooth to the eye, implies a forward curve with a jump at every instrument in the bootstrap. Every product that fixes between nodes prices off those forwards.

The general lesson is the one already visible in the worked example: the forward rate is a derivative of the curve, so it is one degree less smooth than whatever we chose to interpolate. Interpolating linearly in $\ln P$ makes f piecewise constant, hence a step function. Interpolating in the instantaneous forwards themselves makes them as smooth as we like but gives up exact locality — moving one quote moves the curve everywhere. Monotone convex schemes, which the next remark describes, are built to keep forwards continuous while preventing the overshoot a naive interpolation through the forwards produces.

Remark (What a monotone convex scheme is)**.** It is Hagan and West's scheme, and it is the standard answer to the failure the figure below shows.

Start from what the bootstrap actually determines. It does not pin the instantaneous forward at any single date; it pins the *average* forward across each bucket between adjacent nodes, since that average is what discounts the cashflows in that bucket. Call those the discrete forwards f_i^d on $[T_{i-1}, T_i]$. Any interpolation that reprices the quotes must satisfy

$$\int_{T_{i-1}}^{T_i} f(t)\,dt = f_i^d\,(T_i - T_{i-1})$$

for every bucket — the area under the forward curve over each bucket is fixed, however the curve is shaped inside it. A scheme with this property is called *conserving*, and it is the reason the construction reprices by design rather than by iteration.

The scheme then proceeds in two passes:

Input: the discrete forwards f_i^d on each bucket $[T_{i-1}, T_i]$, as the bootstrap left them.

1. *Node values.* For each interior node T_i, interpolate between the two buckets that meet there,

$$f_i = \frac{T_i - T_{i-1}}{T_{i+1} - T_{i-1}} f_{i+1}^d + \frac{T_{i+1} - T_i}{T_{i+1} - T_{i-1}} f_i^d,$$

   weighting each bucket by the length of the *other* one, so the nearer bucket counts for more. At the two ends there is only one bucket, and f_0 and f_N are collared to it.

2. *A quadratic on each bucket.* On $[T_{i-1}, T_i]$ fit the quadratic through f_{i-1} and f_i whose integral is $f_i^d(T_i - T_{i-1})$. Three conditions, three coefficients, one answer — and the area condition is what makes the scheme conserving, so the quotes are repriced by construction rather than by a solver.

3. *Check the range.* The quadratic overshoots when a node value sits too far from the bucket average it belongs to. Written in terms of

$$g_{i-1} = f_{i-1} - f_i^d, \qquad g_i = f_i - f_i^d,$$

the curve stays inside the range its own inputs set exactly when (g_{i-1}, g_i) falls in a particular region of the plane, and leaves it otherwise.

4. *Repair, by case.* Where it leaves, replace the quadratic on that bucket with a piecewise quadratic that flattens: it holds one node value across part of the bucket and moves over the rest, with the split point chosen to preserve the area. Which of a handful of cases applies is decided by the signs and relative sizes of g_{i-1} and g_i, and every case preserves the area and stays monotone between the node values.

   Output: a forward curve that is continuous, reprices every quote exactly, stays inside the range its neighbours set, and is positive wherever the inputs are.

The area condition in step two is the whole reason no iteration is needed to reprice — a conserving scheme reprices by algebra. And the case selection in step four is where the scheme stops being linear, since which case applies depends on the data.

Remark (What the sequence shows that the result does not)**.** Two measurements from the same construction, taken as the three year point is passed and then left behind.[2]

Adding the next three quotes after the three year one moves the three year *zero* by nine ten-thousandths of a basis point. Expecting zero is reasonable and the argument for where it fails is the interesting part.

The argument runs: the scheme is conserving, so whatever happens to the shape inside a bucket, its area is fixed; a zero rate at a node is the sum of the bucket areas up to it; therefore the three year zero cannot move. Every step of that is correct. What it misses is that the three year par bond does not only pay at nodes. It pays semiannually, and with nodes at one, two and three years the coupons at half, one and a half and two and a half years land *inside* buckets — and a partial area is not conserved. Only whole ones are.

The chain is then short. The five year quote creates a bucket beyond three, so the node forward at three years is rebuilt from the buckets on both sides of it and moves by 6.3 basis points. The quadratic on $[2, 3]$ is refitted through that new node value while preserving its area, so the shape inside the bucket changes and the total does not. The coupon at two and a half years is discounted through $\int_2^{2.5} f$, which is half a bucket, and that moves by 1.26×10^{-5} — fifty times the residual left in the whole bucket's area. The three year bond therefore no longer prices to par at the old node value, and the solver moves it until it does.

So the effect is real rather than numerical, and its size is set by how much of the instrument pays between the nodes. It is negligible here and matters in principle: it says the node values are not quite settled even in the region the quotes cover, and that a conserving scheme conserves less than the word suggests once the instruments pay off-node.

Over the same three quotes the *forward* at three years moves by six and a third basis points — seven thousand times as much. Which is §7.4's comparison of the four schemes arriving inside a

[2] `curve::bootstrap_par_stages`.

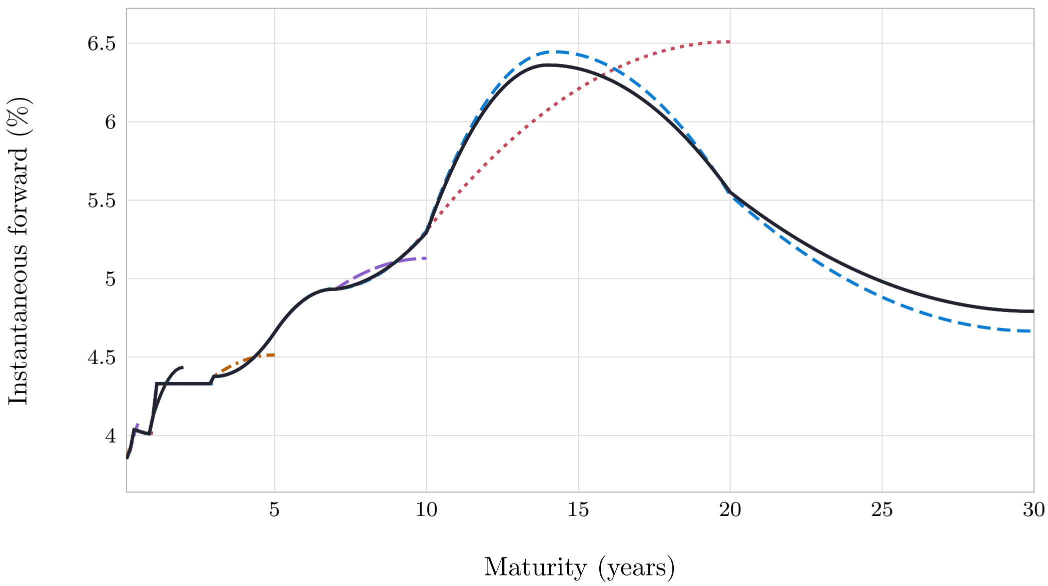

Figure 7.1: The construction as a sequence rather than as a result: the forward curve after each quote is bootstrapped, and then after each sweep of the re-solve. The stage before is drawn faintly behind the current one, because the movement is the point. Two things are visible here that the finished curve hides. Adding a quote does not only extend the curve rightwards — it moves the part already built, since a par bond pays coupons between the nodes and a forward below a node is built partly from the bucket above it. And the stages past the last quote are sweeps, which this scheme needs and the other three do not.

Drawn by `bootstrap_par_stages`. US Department of the Treasury, daily par yield curve rates, as of 2026-08-07 (par yield, semiannual coupon, actual/actual). Retrieved from https://home.treasury.gov/interest-rates-data-csv-archive.

single one: the zeros agree to a rounding error and the forwards do not agree at all, and anything fixing between the nodes is paid on the second.

Set against the other three schemes the effect is not subtle: on the same quotes the monotone convex curve stays inside four per cent to six and a half, where linear interpolation of the forwards runs from under two to nearly eight. The cost is that the case selection depends on the data, so the curve is not a linear function of its inputs. Risk cannot be obtained by applying a fixed matrix to the quotes and has to be bumped numerically instead — and two curves built from the same quotes on different days may not even be in the same case.

None of these is right. They are different priors about what the market would have quoted at maturities it did not quote, and two desks with identical market data and different interpolation will print different prices for anything that fixes off-node.

One might expect the schemes to agree at the nodes and differ only in between, on the grounds that the nodes are pinned by the quotes. That is true only if every calibrating instrument pays exclusively on nodes. It does not: a ten year par bond pays twenty coupons, and only a handful of them land on a node, so the value of the ten year node itself is solved for using interpolated discount factors. The schemes therefore do not produce one curve read three ways. They produce three curves.

7.5 The Curve as a Posterior

The bootstrap finds *the* curve that reprices the quotes exactly. Two things are wrong with that as a statement of the problem, and both were visible earlier in this chapter without being named.

The quotes are mids. A mid is the midpoint of a range in which no trade need have occurred, so repricing it exactly is fitting to a precision the data does not have — and where instruments overlap, as futures and swaps do, exact repricing forces the curve to contort to satisfy quotes that are not mutually consistent to begin with.

And between the quoted maturities the curve is not determined. §7.4 said as much and then chose a scheme; what it did not do is say how undetermined, which is a question with an answer.

Remark (Curve construction is a regression)**.** Put a prior on the forward curve — a Gaussian process, whose kernel says how smooth it is believed to be and over what scale — and treat each quote as a noisy observation of a functional of it. Since a zero rate is an average of forwards,

$$z(T) = \frac{1}{T}\int_0^T f(u)\,du, \tag{7.5}$$

the observation is *linear* in the curve, so the posterior is Gaussian and available in closed form. What comes out is not a curve but a distribution over curves: a mean, which is the fitted curve, and a standard deviation at every maturity.

Calculation 7.21 (The posterior, in full)**.** Represent the forward curve by its values on a grid $u_1, \ldots, u_n$ and give it a Gaussian prior,

$$f \sim N(\mu\mathbf{1}, K), \qquad K_{jk} = \sigma_p^2 \exp\left(-\frac{(u_j-u_k)^2}{2\ell^2}\right), \tag{7.6}$$

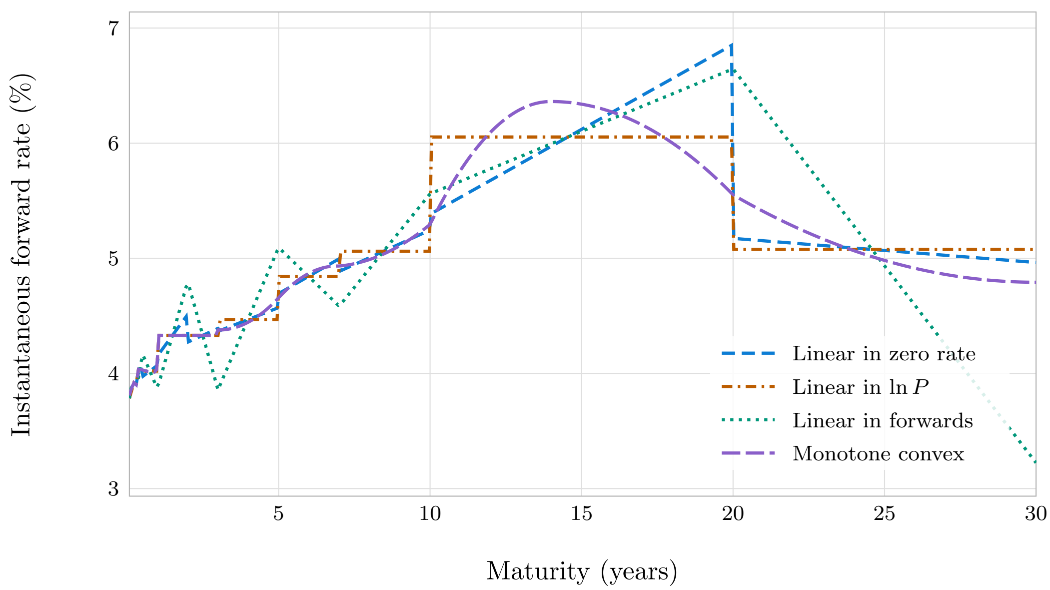

Figure 7.2: The US Treasury par yield curve bootstrapped under four interpolation schemes. All four reprice every quoted instrument exactly, so no examination of the fit distinguishes them, and their zero curves lie within a basis point or two of one another over most of the curve. The instantaneous forward curves, which are what anything fixing between the nodes actually pays, do not agree anywhere. The failures are clearest at the long end, where the quotes are twenty years apart: linear interpolation of the zero rate jumps at every node, linear interpolation of $\ln P$ is a step function, and linear interpolation of the forwards, though continuous, overshoots wildly — forced through the node values, it swings from 3.2% to 6.6% on a curve whose quotes never leave 3.8% to 5.2%. The monotone convex scheme is the repair: continuous like the third, and confined to 3.8% to 6.4% because it is not allowed past the range its neighbours set. It reprices exactly like the others, and it is the only one of the four that cannot be written as a fixed matrix applied to the quotes.

Drawn by `bootstrap_par` and `Curve::inst_forward`. US Department of the Treasury, daily par yield curve rates, as of 2026-08-07 (par yield, semiannual coupon, actual/actual). Retrieved from https://home.treasury.gov/interest-rates-data-csv-archive.

the squared exponential kernel. Three numbers say everything the prior knows: μ, the level the curve reverts to where nothing is observed; σ_p, how far it may stray from it; and ℓ, the distance over which two forwards stop being related. The figure below uses μ equal to the average quoted rate, $\sigma_p = 1.5\%$ and $\ell = 8$ years.

Now the quotes. A zero rate is an average by (7.5), so observing m of them is

$$z = Hf + \epsilon, \qquad H_{ij} = \frac{\Delta u}{T_i}\mathbf{1}\{u_j \le T_i\}, \qquad \epsilon \sim N(0, \sigma_e^2 I), \tag{7.7}$$

with H the $m \times n$ averaging matrix — row i spreads weight $1/T_i$ evenly over the grid points up to T_i and puts zero beyond — and σ_e the quote error, two basis points here, standing for the bid-offer.

Everything is Gaussian and the observation is linear, so (f, z) is jointly Gaussian with

$$\mathrm{Cov}(f, z) = KH^\top, \qquad \mathrm{Var}(z) = HKH^\top + \sigma_e^2 I,$$

and conditioning a joint Gaussian is a formula rather than an argument:

$$\mathbb{E}[f \mid z] = \mu\mathbf{1} + KH^\top\big(HKH^\top + \sigma_e^2 I\big)^{-1}\big(z - \mu H\mathbf{1}\big), \tag{7.8}$$

$$\mathrm{Var}[f \mid z] = K - KH^\top\big(HKH^\top + \sigma_e^2 I\big)^{-1}HK. \tag{7.9}$$

The matrix inverted is $m \times m$ — eight by eight for the curve below — so the whole construction costs less than the bootstrap it replaces.[3]

Remark (Reading the two formulae)**.** In (7.8) the fitted curve is the prior mean plus a correction built from $KH^\top$, the covariance between a forward and each quote. A forward that no quote's averaging window touches has a zero column there and is left at the prior — which is the extrapolation beyond the last quote, arrived at by algebra rather than by convention.

In (7.9) the posterior variance does not mention z. The band is fixed by *which* instruments are quoted and how precisely, and not at all by what they say. So it can be computed before the market opens, and a desk can know where its curve is determined without knowing what the curve is.

And $\sigma_e \to 0$ recovers the bootstrap: the inverse becomes that of $HKH^\top$ alone, the mean reproduces every quote exactly, and the posterior variance collapses to zero in the directions the quotes span. Which is the precise sense in which an exact fit is a limiting case of this rather than a different thing.

Remark (Three constructions, one prior)**.** Writing it this way subsumes the three constructions this chapter has treated as different things. The bootstrap is the limit as the quote error goes to zero with as many degrees of freedom as instruments. A smoothing spline is the posterior mean under a particular prior — that equivalence is a theorem, not an analogy, and it is why the choice of spline and the choice of prior are the same choice. And a parametric family like Nelson-Siegel is an infinitely strong prior that the curve lies in a four-dimensional subspace.

So §7.4's remark that interpolation is a modelling choice can be sharpened: it is a *prior*, and writing it as one makes it something that can be stated, argued about, and checked against how curves actually move.

[3] `curve::fit_curve` is these two lines and the kernel above.

Calculation 7.22 (Where the market has spoken, and where it has not)**.** Fit the committed Treasury snapshot's quoted zero rates from one year out to thirty, with a two basis point error on each, under a prior that the forward curve sits near its own average and can stray one and a half per cent, with an eight year smoothness scale. Two standard deviations of the posterior forward rate:[4]

Maturity	**Forward**	**Band**	
1y	4.13%	±2bp	quoted
2y	4.31%	±4bp	quoted
5y	4.68%	±5bp	quoted
10y	5.17%	±15bp	quoted
20y	5.92%	±43bp	quoted
30y	4.49%	±110bp	quoted, and the last one
40y	4.39%	±287bp	extrapolated

Two features, and neither is what one would guess.

The band widens steadily through the quoted region, and is widest at the longest quote. That is not the picture a reader expects from a fit, which is a curve pinned at its knots and loose between them — and (7.5) is why the expectation is wrong. A zero rate constrains an *average* of forwards, not a forward. The front of the curve is covered by eight overlapping averages and is pinned to a couple of basis points; the stretch beyond twenty years enters only through two of them, so one number is asked to determine a decade of curve and cannot. Any intuition that a curve is best known where it is quoted is an intuition about pointwise observations, and these are not pointwise observations.

And past the last quote the band opens towards the prior. At forty years it is nearly three per cent wide, on its way to the four the prior allows. That is the honest statement about extrapolation: beyond the market's reach the curve is whatever was assumed. A construction that extends a flat forward past thirty years is asserting precisely this prior while presenting it as a curve — and pension and insurance liabilities are discounted off exactly that region.

The smoothness scale is doing visible work in both rows. At a three year scale the long end is a sequence of nearly independent points, one average cannot pin any of them, and the thirty year band comes out wider than any curve has ever been. Eight years is a statement that forward rates decades apart move together, and it is that statement, not the data, which makes the long end look determined at all.

Structure (What the band is for)**.** A distribution over curves is not merely a more honest object than a curve. Three things become computable that were not.

Cheapness becomes a z-score. A bond's spread to the curve is currently a number with no scale attached: is four basis points cheap, or is it inside what the construction cannot determine? The posterior answers it. Chapter 22 insists a wide spread is not a signal; this supplies the width against which to judge, and it is largest exactly where the curve is least constrained, which is where apparent mispricings congregate.

Bucketed risk acquires a covariance. §7.9 reports sensitivities to each calibrating instrument and then needs to know how those buckets move together. Equation (7.9) supplies one, and what

[4] `curve::fit_curve`.

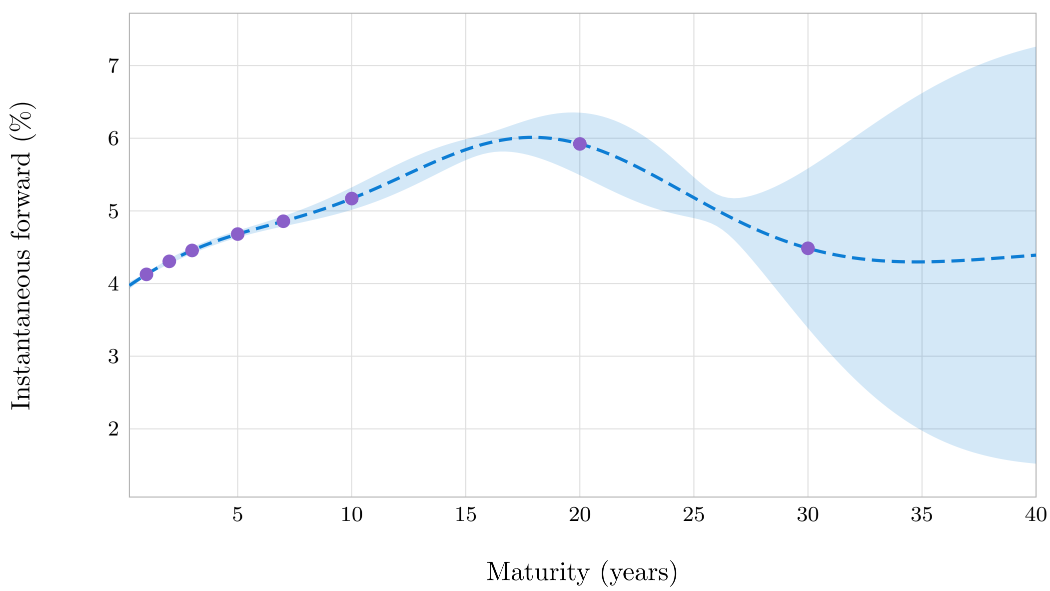

Figure 7.3: The forward curve fitted to the committed Treasury snapshot as a posterior. The shaded region is two standard deviations either side of the mean, and the marks are the eight quoted tenors, from one year to thirty, placed on the fitted curve to show where the information enters — not to claim the curve passes through them, since (7.5) makes each quote a constraint on an *average* of forwards rather than on the forward at its own maturity. The shape of the band is the content: a few basis points across the densely quoted front end, tens of basis points through the twenties where two quotes carry a decade of curve, and opening towards the prior past the last one. Nothing beyond thirty years is observed, and the figure says so rather than drawing a line through it.

Drawn by `fit_curve`. US Department of the Treasury, daily par yield curve rates, as of 2026-08-07 (par yield, semiannual coupon, actual/actual). Retrieved from https://home.treasury.gov/interest-rates-data-csv-archive.

to do with it deserves care, because it is the same distinction chapter 24 draws between pricing and aggregation.

Shock the instruments one at a time to hedge. A bucketed sensitivity is the answer to "how much of instrument i do I trade". It has to be computed by moving instrument i and nothing else, because that is what a hedge does — you buy the ten year swap, not a correlated bundle. Independent bumps are therefore correct for hedging and would be wrong to replace with joint ones.

Shock them together to measure risk. The variance of the book's value is $\Delta^\top \Sigma \Delta$ with Δ the vector of bucketed sensitivities, and no single-instrument bump ever produces that number: a book flat to every bucket individually is flat outright, but a book with offsetting buckets can still be badly exposed if the two do not move together the way the offset assumed. So the covariance enters at aggregation, not at differentiation, and the two are different calculations from the same Δ.

And the eigenvectors are the scenarios worth running. Shocking along the leading eigenvectors of Σ gives level, slope and curvature moves — chapter 24's factors, arriving from the construction rather than from a principal component analysis of history. That is a change of basis on the same information, not extra information.

One caution about *which* covariance, since two are in play and they answer different questions. Equation (7.9) is the uncertainty about today's curve given today's quotes: interpolation risk, which does not go away by waiting and is invisible to a historical estimate. A covariance estimated from a history of quotes is how the market moves from day to day. A full risk figure wants both, and adding them is only correct if they are independent, which is a stronger assumption than it looks.

And the extrapolation risk has a number. A thirty year liability discounted off a curve whose last quote is at twenty carries an uncertainty the table above quantifies. Insurers and pension funds hold exactly this exposure, and reporting it as a point estimate is the same error chapter 21's Bayesian calibration argues against, committed at the one place where nobody thinks of the curve as a model.

The cost is a prior, which somebody must choose and defend. That is the same cost as choosing an interpolation scheme, and it is paid in the same currency — the difference is that this version says out loud what is being assumed, and reports what it is worth.

7.6 The Front End Is a Step Function

Everything above treated the curve as a smooth object to be interpolated. At the short end that is false in a way that costs money, because the overnight rate does not drift — it is held by a central bank and moves only when the committee meets.

The meeting dates are published a year ahead. So the forward overnight curve is a step function with known step locations and unknown step sizes, and the construction that says so is the one to use: piecewise constant between meetings, with a knot at each meeting date, solving for the jumps.

Calculation 7.23 (Solving for the steps)**.** It is a forward substitution, with every part determined by the calendar.

Let the overnight rate be r_0 until the first meeting and step to r_k after the kth, at known dates $m_1 < \cdots < m_K$. An overnight indexed swap pays the compounded overnight rate over its life, so to the accuracy that matters here its quote is the average of the path,

$$q_i = \frac{1}{T_i}\int_0^{T_i} r(u)\,du = \frac{1}{T_i}\left[\sum_{k<i} r_k\,(m_{k+1}-m_k) + r_i\,(T_i - m_i)\right], \tag{7.10}$$

writing $m_0 = 0$. Choose the quotes so that the ith matures just after the ith meeting and every term in that bracket but the last is already known, so

$$r_i = \frac{q_i T_i - \sum_{k<i} r_k(m_{k+1}-m_k)}{T_i - m_i}.$$

One division per meeting, taken in order. There is no optimiser and no residual: the system is triangular because each instrument spans exactly one more meeting than the last, and the calendar is what makes it so.[5]

On four meetings at a quarter's spacing with swaps maturing a fortnight after each, quotes generated from a path that hikes, holds, hikes and cuts return exactly that: +25, 0, +25, −25 basis points.

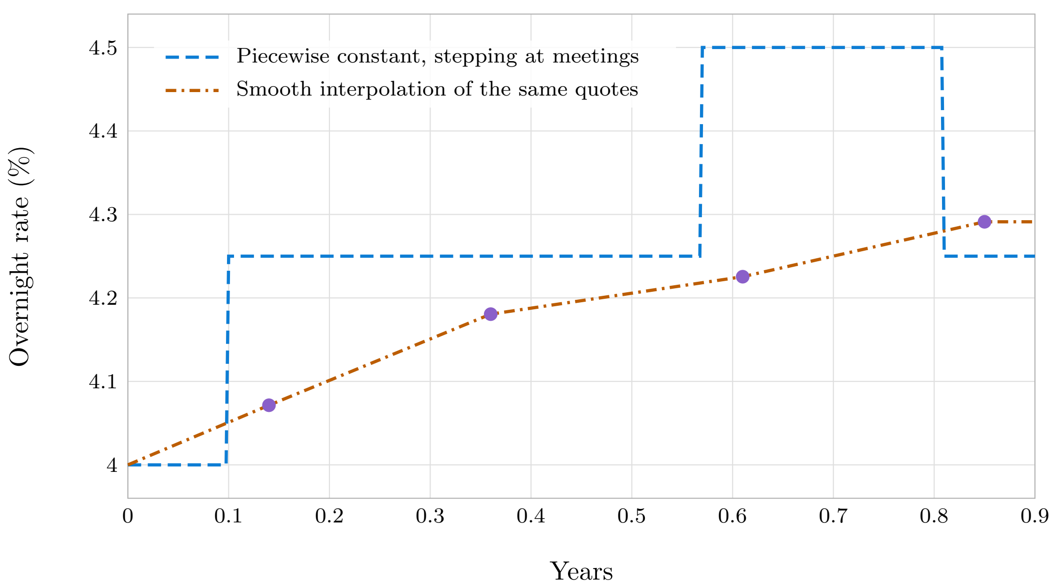

Figure 7.4: The same four quotes built two ways. Both reprice every swap exactly; they disagree about the overnight rate on almost every day in between. The step curve holds the rate flat between decisions because that is what the rate does, and the smooth one spreads each move across the weeks either side of it. The error is worst just *after* a decision — nineteen basis points a week after the first meeting against four and a half a week before — because by then the smooth curve is still catching up with a move that has already happened. Over the year its largest error is twenty-eight basis points, and its sign is set by where in the accrual period the meeting falls, which is public information.

Drawn by `build_meeting_curve`.

[5] `events::build_meeting_curve`, with the inversion checked against a known path.

Twenty-eight basis points is many times the bid-offer, and it is not noise: it is one-signed and determined by the calendar, so any instrument whose accrual straddles a meeting is mispriced systematically. This is §7.4's point at its sharpest — the interpolation scheme is a statement about how the underlying rate behaves, and here the correct statement is available for free from a published calendar.

Remark (The output is a schedule of decisions)**.** Building the curve this way changes what the calibration produces. Instead of a smooth function whose values have no individual meaning, the output is a set of numbers indexed by meeting, each of which is the market's expected policy move there — and since the outcome space is a short list of twenty-five basis point multiples, chapter 4 showed how to read a probability off it. The four steps above divide by the size of a move to give $+1$, 0, $+1$, -1: a hike fully priced, a hold, another hike, a cut.

Nothing in 7.23 consults an opinion. The quotes are market instruments and the steps are whatever reprices them, so the curve reports what has already been priced rather than what anyone believes. That is what makes the number useful: a mispricing acquires a location, at a named meeting, rather than being a general sense that the front end is rich.

Calculation 7.24 (Risk to a decision rather than to a tenor)**.** A front end book is not naturally described by tenor buckets, because what it is exposed to is decisions. The risk report that says so is indexed by meeting: move one step by a basis point, leave the others alone, reprice.

That is not a bucket bump. Moving a step shifts the entire curve *after* that meeting and nothing before it, so what it measures is the exposure to everything the instrument accrues from that date onwards. Differentiating (7.10) gives the answer in closed form:

$$\frac{\partial q}{\partial(\text{step at } m)} = \frac{T - m}{T}, \tag{7.11}$$

the fraction of the instrument's life falling after the meeting.[6]

For a one year swap against the four meetings above, that is 90%, 68%, 43% and 19%. The risk is heavily front-loaded — the nearest meeting carries nearly five times the exposure of the furthest — which no tenor bucketing reveals, because in tenor coordinates a one year swap is simply a one year swap.

And the localisation is exact in a way a bucket cannot be. A three month swap has substantial exposure to the meeting inside its life and *identically zero* to every meeting after it, however close the next one falls. A key rate bucket, which shifts a region of the curve smoothly, always leaks some sensitivity across a date on the far side of maturity; this does not, because the curve genuinely does not move there.

Calculation 7.25 (Trading a view on one meeting)**.** The sensitivities suggest a use. A trader who thinks a meeting will deliver more than the curve has priced wants a position exposed to *that* step and to no other — otherwise the trade is also a bet on the neighbouring meetings, and being right about one while wrong about another is indistinguishable from being wrong.

Such a position exists, and the reason is the structure that made 7.23 a forward substitution. In profit and loss terms a step of size δ at m adds $\delta(T - m)$ to what a unit notional swap maturing at T accrues, so the exposure matrix

$$E_{ik} = (T_i - m_k)^+$$

[6] `events::meeting_sensitivities`, checked against a numerical bump.

is lower triangular, hence invertible, and the weights isolating any one meeting are a back substitution — the same calendar structure used the other way round.[7]

The trade is a butterfly. Isolating the third of four meetings gives weights of about $+4.6$, -8.7, $+4.3$ on the three instruments spanning it and nothing on the fourth: wings up, body down, roughly $1 : -2 : 1$. That is not a coincidence of these dates. Exposure $(T-m)$ is *linear* in the meeting date, killing a linear function is what a second difference does, and it is the same reason chapter 22's curve fly kills level and slope. A meeting view is a butterfly in meeting space.

And the choice of instruments decides whether it can be put on at all. The back substitution divides at each step by $T_i - m_i$, the exposure of an instrument to the meeting it barely spans. Swaps maturing a fortnight after each meeting make that divisor a fortnight, and the notionals compound: isolating the third meeting then needs a gross position of a thousand times the exposure it buys. Swaps maturing just *before the next* meeting make the divisor a whole segment, and the same trade needs eighteen. The instruments that are convenient for the bootstrap are the wrong ones for the trade, and the difference is a factor of sixty.

Remark (What the position is really short)**.** With unit exposure the profit is exactly the gap between what the committee does and what the curve had priced, which makes the trade a clean expression of the view and is the reason to normalise it that way.

A butterfly in meeting space is short the assumption that the meetings either side of the target are correctly priced — the construction removes exposure to their *steps*, not to being wrong about which meeting the market has mispriced. If the committee moves at the neighbouring meeting instead of the target one, the trade loses on both wings. So the view being expressed is sharper than "rates will rise": it is a view about timing, and timing is the thing the position cannot hedge.

Structure (Three constructions, three priors)**.** *Against monotone convex.* It is tempting to reach for §7.4's machinery and put its knots at the meetings, and it is the wrong tool. That scheme exists to stop a *continuous* forward curve overshooting between sparse quotes. Here the curve is meant to be discontinuous and the discontinuities are at known dates, so piecewise constant with knots at meetings is not a crude monotone convex — it is the right shape for a rate held constant by decision and moved by decision. The two therefore govern different regions and are stitched rather than blended: stepping out to the horizon over which meetings are scheduled and quoted, smooth beyond it. Where they join is a modelling choice.

Against the posterior. The step construction is a Gaussian process prior, written in the kernel rather than in prose. Saying the forward is constant between meetings and free to jump at them is saying that $K(u, v)$ of (7.6) is replaced by a block kernel: perfect correlation between two dates in the same inter-meeting segment, and none across a meeting. The smooth schemes are the opposite extreme — correlation decaying with distance and no knowledge of the calendar at all.

Two things follow. The block prior can be *softened* — correlation across a meeting near one rather than at zero — which is the honest description of a curve whose meeting dates are known but whose committee sometimes moves between them. And running the posterior rather than the exact solve puts error bars on the steps, so the implied probability of a hike arrives with a width: not that the market prices a hike, but that it prices one to within so many basis points given the bid-offer on the swaps that determine it.

[7] `events::isolate_meeting`, with the resulting exposures checked to be one at the target meeting and zero at every other.

7.7 Two Curves: Projection and Discounting

Everything so far assumed a single curve doing two jobs: producing the forward rates the floating leg pays, and discounting the cashflows. That assumption is what made the floating leg telescope, and it is false.

It was always slightly false and became visibly so in 2008, when the spread between overnight rates and term interbank rates — which had been a basis point or two — opened to hundreds of basis points and stayed open. Two rates that had been treated as the same rate were plainly not. The reason is credit and liquidity: lending to a bank for three months is not the same as lending overnight and rolling, because the three month loan cannot be withdrawn if the bank deteriorates.

The modern picture separates the two jobs.

- The *projection* curve produces the forward rates a floating leg pays. It is specific to the tenor and the index: a three month curve and a six month curve on the same currency are different curves.

- The *discount* curve present-values the cashflows. It is determined by how the trade is collateralised, and not at all by which index the coupons reference.

So the floating leg becomes

$$\mathrm{Float}(0) = \Sigma_j \delta_j^p \, F^{\mathrm{proj}}(0; \tau_{j-1}^p, \tau_j^p) \, P^{\mathrm{disc}}(0, \tau_j^p),$$

in which nothing cancels, and the par swap rate is the general ratio in (7.3) rather than its telescoped form. Both the projection curves and the discount curve are bootstrapped, from their own instruments, usually simultaneously.

The claim that collateral determines the discount curve deserves a derivation rather than an assertion.

Calculation 7.26 (Collateral sets the discount rate)**.** Consider a derivative, fully collateralised in cash under a standard agreement: at every instant the party that is out of the money posts collateral equal to the trade's value V_t, and the party holding it pays interest on it at the collateral rate $c(t)$, which for a standard cash agreement is the overnight rate.

Take the position of the party holding the trade as an asset. It owns the derivative, worth V_t, and it owes back the collateral it holds, also V_t, on which it pays $c(t)V_t\,dt$. Consider the self-financing portfolio consisting of the derivative, the collateral obligation, and a money market account funded at the risk free rate $r(t)$. Being fully collateralised, the position costs nothing to hold: the collateral received funds the trade exactly. Its total value is therefore zero at all times, and its change is

$$dV_t - c(t)V_t\,dt$$

— the change in the derivative's value, less the interest owed on the collateral against it.

This combination is a traded, self-financing position with zero cost. By the no-arbitrage argument of chapter 4 it must have zero drift in the risk neutral measure, so $\mathbb{E}^{\mathbb{P}}[dV_t] = c(t)V_t\,dt$, and hence the process

$$e^{-\int_0^t c(s)ds} V_t$$

is a martingale under $\mathbb{P}$, giving

$$V_t = \mathbb{E}^{\mathbb{P}}\left[e^{-\int_t^T c(s)ds} V_T \,\middle|\, \mathcal{F}_t \right]. \tag{7.12}$$

The discount rate is c, the rate paid on the collateral, and the risk free rate r has dropped out entirely. It has dropped out because a collateralised derivative is never funded — the collateral funds it — so the cost of funding never enters. This is why the market discounts collateralised trades on the overnight curve, and it is also why the answer changes when the collateral terms change: post collateral in a different currency and c becomes that currency's overnight rate adjusted by the cross currency basis; post no collateral and there is no single c at all, and the value acquires funding and credit adjustments instead.

Remark (The step that assumes rehypothecation)**.** One sentence in that derivation is carrying more than it looks: *the collateral received funds the trade exactly.* That is only true if the party holding the collateral may actually spend it.

Under a standard title transfer agreement it may. The cash is transferred outright rather than pledged, the receiver owns it, and it can be used — rehypothecated — to fund the hedge. Then no separate funding arises, r cancels, and (7.12) follows.

Segregate the collateral instead — hold it at a third party custodian, with no right to use it — and the cancellation fails. The receiver must now fund the position from its own balance sheet at its own rate while the collateral sits earning c somewhere it cannot reach. Two rates appear where one did, the difference between them does not cancel, and the value picks up a funding adjustment.

This is not a technicality about paperwork. It is exactly why variation margin and initial margin are treated differently: variation margin is rehypothecable and therefore discounts at c, while initial margin must be segregated by regulation, cannot be used, and so costs the posting party its funding spread for the life of the trade. That cost has its own name, a margin valuation adjustment, and it exists precisely because the sentence above is false for it.

Remark. Notice that (7.12) is the ordinary risk neutral pricing formula of chapter 4 with the accumulation factor $A_t = e^{\int_0^t r}$ replaced by $e^{\int_0^t c}$. In the language of chapter 6, the collateral account is the numeraire. Nothing about the theory changes; what changes is which account the theory is applied to, and the answer is the one the contract actually pays interest on.

The differences between all these curves are traded in their own right. Tenor basis swaps exchange one tenor's floating leg for another's, overnight indexed swaps exchange fixed for compounded overnight, and cross currency basis swaps exchange floating legs in two currencies. In a single curve world all three would be worthless by construction. That they are not is the measure of how wrong the single curve was.

7.8 The Rest of the Multicurve Stack

Separating projection from discounting is the first split. There are three more major splits, and they compose: a trade is priced off a projection curve chosen by its index, a discount curve chosen by its collateral, a currency adjustment if the two are not the same currency, and — if it is a bond with optionality — a spread reconciling the model to the quote. Each is a different object and they are frequently confused with one another.

Tenor basis

Two swaps on the same currency and maturity, one paying three month rates and one paying six month, do not have the same par rate. The difference is the *tenor basis*, and it is quoted directly as a basis swap: exchange three month payments plus a spread for six month payments.

Remark (What the tenor basis is the price of)**.** Lending for six months rather than rolling three month loans twice differs in two ways, and both are worth money.

The first is credit. The six month loan cannot be withdrawn if the borrower deteriorates at the three month point, so the lender is exposed for longer and charges for it. This is the same argument that separated the discount curve from the projection curve, applied between two term rates rather than between term and overnight.

The second is optionality, and it survives even when the credit component does not. Rolling gives the lender a decision at the three month point — to roll, or not — and a decision has non-negative value. So the rolled position is worth at least the term position, and the basis is the price of that flexibility.

Since 2008 the basis has been persistently non-zero and it widens in stress, which is what a credit and optionality premium should do. In an index transition to overnight-based rates much of the credit component goes away; the optionality does not, and neither does the fact that a curve must be built per index.

The consequence for construction is that there is no single curve. Each index and tenor has its own projection curve, they are built simultaneously with the discount curve because the basis instruments involve two of them at once, and a book of swaps on four indices carries four projection curves and one discount curve, all of which must be arbitrage-free jointly.

Cross-currency basis

The same problem across currencies. Borrowing dollars directly and borrowing euros and swapping them into dollars should cost the same, since chapter 17 derives the forward exchange rate as a ratio of numeraires with no freedom in it. They do not cost the same, and the wedge is the cross-currency basis.

Definition 7.27 (Cross-currency basis swap)**.** Two parties exchange floating payments in different currencies over the same schedule, and — unlike a single-currency swap — they exchange *notionals* as well, at the start and again at maturity, both at the spot rate fixed on day one. A euro-dollar basis swap therefore has one party pay dollar floating on a dollar notional and receive euro floating on a euro notional, having handed over the euros at the start and taken them back at the end.

The basis is a spread b added to the floating rate on the *non-dollar* leg, by convention for every pair quoted against the dollar:

$$\text{pay } \big(\text{EUR floating} + b\big) \quad \text{against} \quad \text{receive USD floating.}$$

A negative b, which is the usual sign, means the party receiving dollars pays for them.

Remark (Why the notional exchange is the whole point)**.** The instrument is often described as though it were an interest rate swap in two currencies and it is not.

It is a funding trade. The notional exchange makes it a loan of one currency against another for the life of the swap: euros in at the start, euros back at the end, dollars the other way. The floating legs are the interest on those loans. So the basis is the price of borrowing dollars against euro collateral rather than a view on either rate, and it moves with the scarcity of dollar funding rather than with the level of either curve.

And it is the reason a discount curve depends on the collateral currency. A dollar cashflow collateralised in euros earns the euro collateral rate on the margin, which has to be converted; the conversion is priced by this instrument, so the discount curve for that trade is the dollar curve adjusted by the basis. That is not a small correction applied afterwards — it is a different curve, and a book with agreements in several currencies carries one per agreement.

For curve construction the effect is mechanical and unavoidable: a dollar cashflow collateralised in euros is not discounted on the dollar curve. It is discounted on a dollar curve adjusted by the basis, because the collateral is earning a euro rate and the conversion is not free. So the discount curve depends on the collateral *currency* as well as the collateral rate, and a book with agreements in several currencies has a discount curve per agreement. Chapter 22 treats the basis as a trade; here it is an input, and treating it as a small correction is how a cross-currency book comes to be systematically mismarked.

Option-adjusted spread

The last object in the stack is different in kind from the others.

Definition 7.28 (Option-adjusted spread)**.** Let $V(s)$ be the model value of an instrument with all cashflows discounted on the curve shifted by a constant spread s, where a model has been used to value any optionality the instrument contains. The option-adjusted spread is the s solving

$$V(s) = \text{market price}.$$

For a bond with no optionality this is just a yield spread expressed against a curve rather than against a single yield. The interesting case is a callable bond or a mortgage pass-through, where the quoted price is low partly because the issuer or the homeowner holds an option. A plain yield spread attributes all of that cheapness to credit and liquidity. The OAS values the option separately and reports what is left.

Remark (An OAS is a property of a model, not of a bond)**.** This is the most misused number in fixed income, and the reason is visible in the definition: $V(s)$ requires a model for the optionality, so the OAS is whatever spread reconciles *that model* to the price. Two desks with different volatility assumptions compute different OAS for the same bond at the same price, and neither is wrong. The number is a residual, and a residual carries every error in everything that produced it.

Two consequences follow.

Comparing OAS across desks or across model versions is meaningless, whereas comparing it across bonds priced by one model on one day is exactly what it is for — the model errors are common and largely cancel in the comparison.

And the spread should be *uncorrelated with volatility*, which makes a diagnostic. Suppose the volatility market rises. The bond cheapens, because the option the issuer holds is worth more. A model given the new volatility reprices that option by the same amount, so at an unchanged s the model value falls with the market's and the spread does not move — which is exactly what "option-adjusted" is claiming to have achieved. If instead the model's volatility is stale, or its option is misvalued, the model price does not fall as far as the market's and s has to move to close the gap. So an OAS that tracks the volatility surface is the warning sign: the optionality has not been removed, only relabelled as spread. Regressing changes in OAS on changes in implied volatility costs nothing, and a significant loading means the number is not doing the job its name claims.

7.9 Curve Risk

Finally, how the risk of all this is expressed. A book of linear products has one exposure, the curve, but the curve is a function and not a number, so its risk has to be bucketed.

The crudest measure moves the whole curve. The parallel DV01 is the change in value for a one basis point shift in every rate, and for a single swap it is the annuity, by the exercise above.

The parallel number is not enough, because a book can be flat to a parallel shift and heavily exposed to the curve steepening. What is computed instead are key rate durations: shift the quote of one calibrating instrument by a basis point, rebuild the entire curve, and reprice. The result is the sensitivity to that instrument, in the units the trader can actually hedge in, because the hedge is a position in that instrument. Summing the buckets recovers the parallel number.

This is a general principle rather than a fixed income one, and it deserves stating plainly. The risk is reported against the instruments used to build the curve, not against the curve's own internal parametrisation, and the two differ. Change the interpolation scheme and the bucketed risk moves even though the curve reprices the same instruments, because a shift in one quote propagates differently. Risk is a property of the construction, not only of the prices.

References

- Hagan, P. S., & West, G. (2006). Interpolation methods for curve construction. *Applied Mathematical Finance*, 13(2), 89–129.

- Piterbarg, V. (2010). Funding beyond discounting: collateral agreements and derivatives pricing. *Risk*, 23(2), 97–102.

- Bianchetti, M. (2010). Two curves, one price. *Risk*, 23(8), 66–72.