The Magpie Yield Curve Model

Julian D. A. Wiseman*

November 2026

Abstract

The Magpie Yield Curve Model is a parsimonious model for fitting a parameterised smooth curve through prices of government bonds and similar instruments. Its desirable properties include the ability to share information across related curves, arbitrary differentiability with respect to the parameters, and differentiation through the optimisation—the change in parameters and hence theoretical prices per change in market prices.

The author and colleagues have used the Magpie Yield Curve model in multiple currencies, developed and emerging, large and small, both for relative value and for hedging.

Index terms: yield curve, model, fitting, parametric, relative value, multi-curve, differentiation

Contents

1 Introduction **1**
1.1 Bonds 1
1.2 Policy interest rates 2
1.3 Yield curve models 2
1.4 Literature 2

2 The Magpie Model **3**
2.1 The 'Magpie' name 3
2.2 Examples used 3
2.3 Measuring time 4
2.4 Rates 4
2.5 The error function and optimisation 5
2.6 Re-fitting 6
2.7 Loneliness 7
2.8 Curve connections 7
2.9 Connecting governments to swaps 9
2.10 Generalising linkages 9
2.11 Curve minus curve 9
2.12 Fitting speed 10
2.13 Trades, relative-value versus macro 10
2.14 Nomenclature 12
2.15 Multiple-curve equation changes 12
2.16 Choosing the horizontal parameters 12
2.17 Other error functions 13

3 Problems and features **14**
3.1 'Problem': approximate fitting . 15
3.2 Problem: meeting dates 15
3.3 Problem: forward guidance . . . 15
3.4 Problem: yield curve control . . . 16
3.5 Problem: imperfect implementation 16
3.6 Problem: weekends 16
3.7 Problem: the possibility of default 17
3.8 Problem: tax 17
3.9 Problem: repo specialness 17
3.10 Problem: pricing inefficiency . . 17
3.11 Problem: non-optimised choices . 18
3.12 Problem: plotting par curves . . 18

4 Good qualities **19**

1 Introduction

1.1 Bonds

Governments borrow money, in large part, by selling conventional bonds. Bonds are sold for cash—which governments spend on education, police and justice, health, defence, roads, economic research, and even repaying previous borrowings. In return for this the holder receives periodic coupons, and at maturity, with the last coupon, the principal of 100. The periodic coupons are semi-annual in the USA, the UK, Japan, Australia, Italy and some others; and annual in most of continental Europe. E.g., £100 'nominal' of the UK gilt 1⅝% of 22nd October 2071 pays coupons of £1⅝ ÷ 2 = £0.8125 every 22nd April and 22nd October (or on the following business day), until 22nd October 2071 when its final payment will be £100.8125.

*Contact via `www.jdawiseman.com/author.html`

1.2 Policy interest rates

In their currency, central banks set a short-term interest rate. (The mechanism by which this is attempted is not the subject of this paper.) This short-term policy rate has varied in the past, and is expected to vary in the future. These anticipated changes should be somehow embedded in bonds' prices. It can be useful to market practitioners to reverse out the implied path of expected central-bank policy, and to show which bonds' prices are cheap or expensive to that implied path.

1.3 Yield curve models

Broadly, two types of model yield curve model are used in fixed-income markets:

- Precise-fitting (sometimes called 'interpolated') used for pricing; and
- Best-fitting ('parametric') for estimation of relative value.

If there are known prices for a 5-year and 6-year swap, and it is wanted to make a tradable price for the 5-year forward 1-year, that price should be exactly consistent with the spot 5-year and 6-year swaps.

But one cannot trade a 5-year forward 1-year bond, and attempting to simulate that with 5- and 6-year bonds has a high repo cost. That allows the market's implied forward bond yields to deviate far from a natural value, because the relative value is often difficult or impossible to arbitrage. Hence the type of model more useful for bonds is parsimonious, smooth, and best-fitted to bonds' prices, but not necessarily precisely fitting any. It is broadly consistent with the pricing of a set of bonds, and shows the extent to which each bond is cheap or expensive to this.

Magpie is a yield curve model of the latter type: parsimoniously parametric, best-fitted, not necessarily fitting anything exactly, non-trivial deviations from exact being candidate relative-value trades.

1.4 Literature

One of the most famous early parametric yield curves models was by Nelson and Siegel [1987]:

> ... the instantaneous forward rate at maturity m, denoted $r(m)$, is ...
>
> $$r(m) = \beta_0 + \beta_1 \cdot \exp(-m/\tau_1) + \beta_2 \cdot \exp(-m/\tau_2)$$
>
> where τ_1 and τ_2 are time constants ... and β_0, β_1, and β_2 are determined by initial conditions.
>
> ... Experimentation with fitting this model to bill yields suggested that it is overparameterized. As the values of τ_1 and τ_2 were varied, it was possible to find values of the β's that gave nearly the same fit. Standard software for estimating nonlinear models failed to converge, giving another indication of overparameterization.

Yes, if "τ_1 and τ_2 were varied", so not constant, then this is overparameterised. But for fixed τ_1 and τ_2, with them very different and within the horizontal scale of the bill maturities, the fitting of the β's is extremely stable.

So instead Nelson and Siegel chose that the "forward rate at maturity m, denoted $r(m)$", be

$$r(m) = \beta_0 + \beta_1 \cdot \exp(-m/\tau) + \beta_2[(m/\tau) \cdot \exp(-m/\tau)]$$

There are several reasons this was not used in Magpie.

- First, it is underparameterised. It can have at most one turning point, which is enough for a bill curve (out to one year), but yield curves in multiple currencies have often had two turning points: negative to a nearby low in monetary policy; positive from there until about 15 or 20 years, then negative again.
- Second, the actual process by which interest rates are set doesn't have anything linear about it, so the linear term is unnatural.
- Third, the real world has decay-like processes working at very different speeds. When the economy has been overheating, there is a significant per-month probability of a higher policy interest rate. But the effect of demography on savings burns very slowly.

Despite the fame of the Nelson and Siegel paper, it wasn't the first with a parsimonious model. Dobbie and Wilkie [1978] describe a model for the gilt market, with equation

$$\text{Yield}(t) = A + Be^{-Ct} + De^{-Ft}$$

"where t is the term to redemption", and all five parameters are optimised. This modelled yield as a function of bond maturity, conceptually different from valuing each cashflow from a path of a prospective short-term rate. Further, the parsimony was reduced by having three curves, based on splitting the stocks into tertiles of low-, medium- and high-coupon.

Svensson [1994], extended Nelson and Siegel by one term, to "increase the flexibility and improve the fit":

$$r(m) = \beta_0 + \beta_1 \exp\left(-\frac{m}{\tau_1}\right) + \beta_2 \frac{m}{\tau_1} \exp\left(-\frac{m}{\tau_1}\right) + \beta_3 \frac{m}{\tau_2} \exp\left(-\frac{m}{\tau_2}\right)$$

But, as well as the disliked linear term, Cairns' reasoning (next paragraph) persuaded me that this is also unstable, in the sense of having multiple local minima.

The strongest effect on Magpie was reasoning by Andrew J. G. Cairns about the instability of its predecessor, my 1994 Exponential Yield Curve Model (Wiseman [1994]). My memory dates Cairns' reasoning to 1997, consistent with Magpie being published in February 1998 (Wiseman [1998]) after a few months of development and coding. So my source might have been Cairns [1997], or a draft of it, or a conference.

Then Cairns and Pritchard [2001] introduced a restricted-exponential family of models. These use an equation equivalent to Magpie's equation (1). And, as with Magpie, "we fix the exponential parameters ... at the outset and at no future point do we estimate their values."

But none of these models specify an error function that lies within the model, which denies them the power to differentiate through the optimisation (§2.6). And none of these models allow sharing of information across curves that are related but not identical (§2.8).

2 The Magpie Model

2.1 The 'Magpie' name

The first draft of the Magpie model was devised at J. P. Morgan (Wiseman [1998]), and then reused at other financial institutions including Société Générale (Wiseman [2012]). Magpie was subsequently extended to include the differentiation through the optimisation.

British folklore has it that magpie birds, *Pica pica* of the crow family, gather shiny objects. And in the Magpie model, when a curve lacks some relevant information, it borrows bright and shiny pieces of information from elsewhere. The name might not be perfect—is yours?—but it was the name chosen at its birth, and has been used since.

2.2 Examples used

In this paper the main examples used will be £ and €, but the model and reasoning are generally applicable in other currencies.

UK government bonds are called 'gilts', and are denominated in £. All gilts are of equal credit, the only distinguishing features are the maturity date (as of time of writing, the longest being 22nd Oct 2073), and the coupon (the largest being the 6% of 7th December 2028, issued in 1998; and the smallest being the 0⅛% of 31st Jan 2028, issued in 2020).

The euro is a monetary union of twenty-one countries. About ten national governments have a properly active market in their debt (DE, FR, IT, ES, NL, BE, AT, PT, IE, FI), as do a few central entities (the EU itself, previously some special purpose vehicles). Obviously German bonds and Portuguese bonds are not identical: they are of different creditworthiness and different liquidity (meaning cost of trading in large size). But they are in the same currency, so there is some information in common between them.

Within these debt markets there are plenty of national quirks. Germany issues bonds with three names: Schatz (initial maturity about 2 years), Obl (5 years), and Bund (7, 10, 15, and 30 years), but despite the

different names, the credit is the same. Most eurozone government bonds pay coupons annually, but not Italy's, which pay semi-annually and have an unusual formula for conventional yield. Austria has issued two bonds with an original maturity of 100 years, maturing in Sept 2117 and June 2120. Germany's longest is just over 30 years, but several have issued 50-year bonds.

But despite GBP and EUR being the main examples, Magpie, or an earlier version of it, has been used in many other currencies including USD, JPY, CHF, CAD, AUD, NZD, SEK, NOK, ZAR, PLN, CZK, HUF, ILS, TRY, SGD, KRW, TWD, MXN, and also some of the pre-euro currencies. It has also been used in BRL, though with the required change in the measurement of time. So readers may assume that the Magpie Yield Curve Model has been and is usefully applicable in many currencies.

2.3 Measuring time

Interest rates are charged per calendar day, so for any given positive short-term interest rate, it will accumulate to more in a leap year than in a regular 365-day year. Hence within the model, time is measured in calendar days, not years. (Except Brazil, for which the convention is that interest is charged per business day, so that is the unit in which time is there measured.)

Day zero is the trade date, so bonds typically settle between day 1 (gilt settlement being 'T+1', and assuming no intervening bad days), and day 4 (assuming an intervening weekend and in euro for which 'T+2' is usual, at least until October 2027).

Yields, the y_i and x_i, are quoted per day rather than per year, so are 'small'. But to users output is more intuitive if quoted in units of yield per year rather than per day. Usual practice for such presentation has been, Julian-style, to multiply the continuously compounded rates by 365¼, and likewise to multiply implied changes in continuously-compounded bond yields—e.g., equation (26). Hence $y_0 = 0.0001$ is, quoted per year, a 3.6525% contribution to the continuously-compounded yield. Devout Gregorianists might prefer a multiplier of $365^{97}⁄_{400} = 365.2425$: it's a preference about presentation, not affecting valuation.

2.4 Rates

The model specifies the form for the instantaneous continuously-compounded forward rate curve, $r(t)$:

$$r(t) \equiv y_0 + y_1\, e^{-t/z_1} + \cdots + y_n\, e^{-t/z_n} \tag{1}$$

The n is fixed. The z_i are fixed, and have units of time. The y_i have units of yields, so time^{-1}. In large currencies $n = 4$ gives enough flexibility to capture the set of plausible smooth yield curves. In £ and €, $z_4 = 91.3125$ (3 months of a year of 365¼ days); $z_3 = 547.875$ (1½ years); $z_2 = 2191.5$ (6 years); $z_1 = 8766$ (24 years); and, implicitly, 'z_0' $= +\infty$.

For US dollars, in which the longest government bond is only 30 years, $z_1 = 7305$ (20 years); the other z_i being the same as for £ and €.

The form of $r(t)$ resembles a natural process: y_0 is the value at infinity, there currently being deviations away from this of size y_i, these deviations decaying with mean lives, with time constants, of the z_i.

Integrating $r(t)$ is easy, but it is useful to define intermediate terms, dm_i, being derivative multipliers, for a particular **s**ettlement day s and **p**ayment day p. The dm_i are independent of the y_i:

$$dm_0(s,p) \equiv s - p \tag{2a}$$

$$dm_i(s,p) \equiv z_i \cdot (e^{-p/z_i} - e^{-s/z_i}) \quad \forall i \geq 1 \tag{2b}$$

The discount between days s and p, that is, the value at day s of of one currency unit at day p, is

$$disc(s,p) \equiv e^{(-\int_s^p r(t)\,\mathrm{d}t)} = \exp\Big[\sum\nolimits_{i=0}^{n} y_i \cdot dm_i(s,p)\Big] \tag{3}$$

Derivatives and repeated derivatives then become mere multiplication.

$$\frac{\partial disc(s,p)}{\partial y_i} = disc(s,p) \cdot dm_i(s,p) \tag{4a}$$

$$\frac{\partial^2 disc(s,p)}{\partial y_i \partial y_j} = disc(s,p) \cdot dm_i(s,p) \cdot dm_j(s,p) \tag{4b}$$

$$\frac{\partial^3 disc(s,p)}{\partial y_i \partial y_j \partial y_0} = disc(s,p) \cdot dm_i(s,p) \cdot dm_j(s,p) \cdot dm_0(s,p) \tag{4c}$$

Instantaneous slide for a cashflow, being the change in price caused by the passage of time with the yield curve unchanged, is

$$-\frac{\partial disc(s,p)}{\partial p} = disc(s,p) \cdot r(p) \tag{5}$$

(though it would be usual to subtract from this the cost of funding).

A bond is then the sum of its cashflows. Consider bond b, the purchase of which would settle on day s_b (expressed as an integer), it paying amounts $a_{b,h}$ on days $p_{b,h}$ (if necessary moved to the next business day). Its theoretical price is

$$thry_b \equiv \sum\nolimits_h a_{b,h} \cdot disc(s_b, p_{b,h}) \tag{6}$$

which, with equations (4) and (5), implies similar equations for the derivatives and the slide.

2.5 The error function and optimisation

An error term is needed, something to minimise. Bond markets often quote the price of a bond as an internal rate of return, as a yield, and any natural error term must closely resemble sum-square yield error. But there are powerful advantages to using an expression entirely within the model. With $market_b$ being the market 'dirty' price of bond b, define:

$$error_b \equiv \left(\ln\left[\frac{thry_b}{market_b}\right] \cdot thry_b \div \frac{\partial thry_b}{\partial y_0} \right)^2 \tag{7}$$

The log term is close to the price error, as a ratio. Inside the square, the log term is divided by the change in price, per price, per unit y_0. That is, $error_b$ is very close to the square of the parallel move in $r(t)$ that would move $thry_b$ to $market_b$. This square of a change in y_0 is analogous to the square of a yield error. Indeed, for a bond with only one non-zero cashflow, this expression simplifies to the square of the error in the continuously-compounded yield. (In §2.17 is more explanation of this choice of bond error function, and of a possible alternative suitable for swaps.)

We have already seen that it is easy to compute $thry_b$'s derivatives, of which we will need those implied by equations (4) and (6), from which can be computed analytic expressions for $error_b$'s first and second derivatives. There are quite intricate, but not prohibitively so, and because these derivatives are (or pre-AI were) inconvenient to transfer from mathematical typesetting to code, they are shown here in a format nearer that of computer code. Using suffixed `thry` for derivatives of $thry_b$, and assuming pre-assignment of `lnTM` $\coloneqq \ln[thry_b/market_b]$:

$\partial error_b/\partial y_i =$
```
(thry0 * thryI * (1 + lnTM) - thry * thry0I * lnTM)
                                                  * 2 * lnTM * thry / (thry0^3)
```
(8)

$\partial^2 error_b/\partial y_i \partial y_j =$
```
( thry0^2 * thryI * thryJ * (lnTM^2 + 3 * lnTM + 1)
 + thry * thry0 * (thry0 * thryIJ - 2 * (thryI * thry0J + thryJ * thry0I)) * (1 + lnTM) * lnTM
              + thry^2 * (3 * thry0I * thry0J - thry0 * thry0IJ) * lnTM^2 ) * 2 / (thry0^4)
```
(9)

Sometimes use is also made, post-optimisation (§2.6), of

$\partial^2 error_b/\partial y_i \partial market_b =$
```
(2 * lnTM * thry * thry0I - (1 + 2 * lnTM) * thry0 * thryI)
                                                * 2 * thry / (market * (thry0^3))
```
(10)

Define **H**, the Hessian, to be the square matrix of double derivatives:

$$\mathbf{H} \equiv \left(\sum_b \frac{\partial^2 error_b}{\partial y_i \partial y_j} \right) \tag{11}$$

Then choose the y_i that minimise $\sum_b error_b$. Start with any sensible guess for y_i, and by Newton-Raphson repeatedly improve the current guess using

$$\left(y_i^{\text{improved}} \right) \leftarrow \left(y_i \right) - \left(\mathbf{H} \right)^{-1} \times \left(\sum_b \frac{\partial error_b}{\partial y_i} \right) \tag{12}$$

This is a very stable and very fast optimisation, which reaches machine precision in about four steps, even with a rough starting guess (§2.12). And, as market prices change, the y_i and the $thry_b$ can be updated even faster (§2.6).

Most bonds have multiple cashflows. But some fixed-income instruments have only one post-purchase cashflow, including T-Bills and futures settling against 'IBOR or against compounded overnight. For one-cashflow instruments equations (7) to (10) can be much simplified. Say that the instrument costs $market_b$ and later pays a single cashflow a_b. Using the derivative multipliers of equation (2), define:

$$ye_b = \left(\ln[a_b / market_b] + \sum_i y_i dm_i \right) \div dm_0 \tag{13}$$

Then for one-cashflow instruments, equation (7) is equivalent to $error_b = (ye_b)^2$. Given which:

$$\partial error_b / \partial y_i = 2 \cdot ye_b \cdot dm_i \div dm_0 \tag{14}$$

$$\partial^2 error_b / \partial y_i \partial y_j = 2 \cdot dm_i \cdot dm_j \div (dm_0)^2 \tag{15}$$

$$\partial^2 error_b / \partial y_i \partial market_b = -2 \cdot dm_i \div (market_b \cdot (dm_0)^2) \tag{16}$$

Equation (15) shows that the second derivatives are independent of the y_i, and hence the error function is precisely parabolic. This is Newton-Raphson's best case: if all the instruments are single-cashflow then, subject to machine precision, a single step should suffice to reach the minimum error. A zero-coupon bond is a single cashflow bond; a low-coupon bond is similar to single-cashflow bond; its error is nearly parabolic; hence the fitting is fastest when most of the long-dated instruments are low-coupon.

2.6 Re-fitting

Assume that the §2.5 fitting has been done. That is, $\forall i$, $\sum_b \frac{\partial error_b}{\partial y_i} = 0$. Then, after that initial fitting, the price of bond b moves by a small amount, ϵ_b. That sum of derivatives would now be non-zero by $\epsilon_b \frac{\partial^2 error_b}{\partial y_i \partial market_b}$, which can be computed analytically by equation (10) or (16). Correcting that error by Newton-Raphson implies

$$\left(\frac{\partial y_i}{\partial market_b} \right) = - \left(\mathbf{H} \right)^{-1} \times \left(\frac{\partial^2 error_b}{\partial y_i \partial market_b} \right) \tag{17}$$

This is the partial derivative, through the optimisation, of the y_i with respect to $market_b$. This is written as $\frac{\partial y_i}{\partial market_b}$, it being understood that any derivative with respect to a market price entails an implied re-optimisation.

Hence the sensitivity of the theoretical price of bond b to small moves in bond c, through an intervening optimisation, is

$$\frac{\partial thry_b}{\partial market_c} = \sum_{i=0}^{n} \frac{\partial y_i}{\partial market_c} \frac{\partial thry_b}{\partial y_i} \tag{18}$$

Actual optimisation could be done as infrequently as once per day, conceptually at the start, using the current cashflows (typically which are all one business day shorter than the previous day's, and from which have been removed payments made). That initial optimisation done, re-optimisation can then be a matrix multiplication: the input being changes in market prices; the output being changes in theoretical prices.

This has another use. Some yield curves have a few bonds, benchmarks, for which real-time pricing is faster and more reliable (typically the bonds most closely associated with futures contracts). These benchmarks can be used to reprice the whole curve. This requires two fittings. As before, fit using all bonds. Also fit a separate curve to just the benchmarks, with prices $bench_c$. (If there are fewer than $n+1$ benchmarks, then: either drop some of the z_i, usually the smallest-z largest-i; or add some instruments such as the interest-rate futures, per §2.9.) Then equation-(18) re-pricing uses $\frac{\partial y_i}{\partial bench_c}$ from the benchmark curve, and $\frac{\partial thry_b}{\partial y_i}$ from the all-bonds curve.

Separately, if plotting such bond-to-bond sensitivities, whether fixing the driver or the driven, it is more natural to display in something like relative yield terms:

$$\frac{\partial thry_b}{\partial market_c} \cdot \frac{\partial thry_c}{\partial y_0} \div \frac{\partial thry_b}{\partial y_0} \tag{19}$$

2.7 Loneliness

To illuminate the case when $b = c$, briefly consider a very simple model, in which a straight line is to be fitted through eleven $\{u, v\}$ points. For ten of the points, $u{\approx}2$; for the other, $u{\approx}50$. Self-evidently, moving the v of one of the ten $u{\approx}2$ points by 1 unit would move its fitted value by about 0.1 units. But moving the v of the only $u{\approx}50$ point by 1 unit would move its fitted value by almost 1 unit. So define:

$$loneliness_b \equiv \frac{\partial thry_b}{\partial market_b} \tag{20}$$

in which the derivative is through the intervening re-optimisation, per equations (17) and (18). This is the 'loneliness' of a bond, the reciprocal of the number of bonds in its 'neighbourhood'.

Numerically, $\sum_b loneliness_b$ seems always to be very close the number of y_i to be fitted. Indeed, again numerically, it seems that the equivalence is exact if $thry_b = market_b \, \forall b$. (A proof would be welcome, even if a conditional proof.) Heuristically, $loneliness_b$ is the number of parameters absorbed by bond b.

Then, in the single-curve case of the model, order the bonds by maturity, or, more pedantically (Wiseman [2008]), by $-\frac{\partial^2 thry_b}{\partial y_0 \partial y_0} / \frac{\partial thry_b}{\partial y_0}$. Take the shortest, the longest, and those the sum to which of $loneliness_b$ crosses integers. Heuristically, one parameter is 'used' in each of the ranges.

2.8 Curve connections

Figure 1 lists eleven euro-denominated government and pseudo-government yield curves. These eleven curves are not identical, but nor are they wholly different. If, for example, it were widely anticipated that in two months the European Central Bank will raise its policy interest rate, then the very front end of all of these curves would be upward sloping. One might observe the difference between the theoretical 6-month and 1-month interest rates, and in all of these curves it should be positive.

As a motivating example consider a small country, such as Portugal. Not that many bonds exist: about one or two bonds per year out to twelve years, and less dense in longer maturities, the size of each being about €10 billion. There are shorter-dated T-Bills, about one every other month, of average size about €2.5 billion, so trading action is thin and there is not always a usefully tradable price. Indeed, because Portuguese Bills are little traded, many non-giant financial firms do not bother with the expense and hassle of sourcing prices. Hence the Portuguese optimisation does not have many short-dated instruments to which to fit. The fitting's backwards extrapolation to time zero might have results that seem nonsensical: $r^{\mathrm{PT}}(0) - r^{\mathrm{DE}}(0)$ could be be negative despite Germany being the better credit; or could be unrealistically large positive despite there being negligible chance of immediate Portuguese default. But ask a bond trader "if there were a one-month and a six-month Portuguese bond, what would the yield difference be?"; the answer would use and be close to the equivalent spread in another country.

And indeed, Italy is a big country with a big national debt: many large bonds, many large T-Bills. The shape of the yield curve is readily observable in Italy.

This suggests a strategy: set $y_4^{\mathrm{PT}} = y_4^{\mathrm{IT}}$ (reminder: z_4=3 months) and set y_3^{PT}=y_3^{IT} (z_3=18m). And because Italy has a 50-year bond whereas Portugal's longest is 27 years, also set $y_1^{\mathrm{PT}} = y_1^{\mathrm{IT}}$ (z_1=24 years), and then jointly fit the two curves.

This is **not** setting the Portuguese curve to equal the Italian curve. Despite the other linkages, y_0^{PT} (flat, 'z_0'=+∞) and y_2^{PT} (z_2=6y) remain free to accommodate Portuguese prices. The Portuguese curve could be above or below the Italian curve; it could be flatter or steeper.

If the Italian data were to force y_1^{PT} or y_3^{PT} or y_4^{PT} to be 'wrong', then y_0^{PT} and y_2^{PT} would fit in a manner that best corrected this error. But the Italian data cannot force y_1^{PT} or y_3^{PT} or y_4^{PT} to be 'wrong', as there do not exist enough Portuguese instruments for a 'correct' value for these parameters to be observed or contradicted—hence the linkages.

This method generalises across the many euro-zone issuers. Figure 1's first six columns shows a set of linkages, similar to those previously used by this author. The rightmost three columns have explanatory information.

Figure 1: Euro-zone y equivalences, expressed as x indices

i:	0	1	2	3	4	Issuer	≈Longest	Bills
z_i:	∞	24y	6y	18m	3m		orig. mat.	
EU	0	11	17	28	31	European Union	30 year	·
DE	1	12	18	28	31	Germany	30 year	✓
AT	2	12	19	28	31	Austria	100 year	·
NL	3	12	20	28	31	Netherlands	30 year	·
FI	4	12	21	28	31	Finland	30 year	·
FR	5	13	22	29	31	France	50 year	✓
BE	6	14	23	29	31	Belgium	50 year	·
IT	7	15	24	30	31	Italy	50 year	✓
ES	8	16	25	30	31	Spain	50 year	·
PT	9	16	26	30	31	Portugal	30 year	·
IE	10	16	27	30	31	Ireland	30 year	·

There are multiple sets of y parameters. For some purposes it is convenient to keep these as y; but for the optimisation it is convenient to relabel as one series of yield parameters called x. The table shows that

$$x_{12} \;=\; y_1^{\mathrm{DE}} = y_1^{\mathrm{AT}} = y_1^{\mathrm{NL}} = y_1^{\mathrm{FI}} \tag{21}$$

Investors perceive Germany, Austria, The Netherlands and Finland to be the top rank of European sovereign issuers. But only Austria has very long government bonds: as well as the 94-year June 2120 listed in figure 1, also a Sep 2117, a Nov 2086 and a Jan 2062. The Netherlands has a 30-year, Finland a 31-year, and Germany a 30-year, so the three non-Austrian countries lack information about the steepness materially beyond 30s. Austria can provide this missing information, done by these four countries sharing their y_1.

There are other y_1 linkages: France, Italy, and Belgium have issued a 50-year. So has Spain, which shares its information with two of the smaller countries: Portugal and Ireland.

$$x_{16} \;=\; y_1^{\mathrm{ES}} = y_1^{\mathrm{PT}} = y_1^{\mathrm{IE}} \tag{22}$$

The fastest-decaying parameter, y_4, is effectively a monetary-policy parameter about forthcoming meetings. It is shared across the whole currency zone, as they share a central bank.

The y_3 ought to fit to short-dated instruments, mostly sub-2-year. But old bonds are generally quite sparse: typically only one to four per issuer per year. Germany, France, and Italy have active trading in their Bills, and most investors serious about fixed income collect prices for these instruments. Most other euro zone issuers have much less active markets in these instruments, and prices are not generally available. The 'Bills' column in figure 1 is thus a little ambiguous: it could be interpreted as a statement about the instruments, or about whether prices are available. Whichever, the y_3 are collected into three groups of issuers: Germany shaping the best credits; France shaping itself and Belgium; and Italy shaping the others.

$$x_{28} \;=\; y_3^{\mathrm{DE}} = y_3^{\mathrm{AT}} = y_3^{\mathrm{NL}} = y_3^{\mathrm{FI}} = y_3^{\mathrm{EU}} \tag{23}$$

$$x_{29} \;=\; y_3^{\mathrm{FR}} = y_3^{\mathrm{BE}} \tag{24}$$

$$x_{30} \;=\; y_3^{\mathrm{IT}} = y_3^{\mathrm{ES}} = y_3^{\mathrm{PT}} = y_3^{\mathrm{IE}} \tag{25}$$

2.9 Connecting governments to swaps

The eurozone example discussed what happens when there aren't prices for T-Bills, either because there aren't at all, or because they just aren't available.

The same can happen in other currencies. Faced with this problem in gilts, or in Swiss Confederation bonds, it is natural to import the shape of front part of the curve from the futures settling against compounded overnight rates, each of which is treated as a forward-settling three-month Bill. In Swiss francs the futures are active out to about two years, so there is little overlap between the curves: thus in Swiss it is natural to set $y_i^{\text{SARON}} = y_i^{\text{Gov't}}\ \forall i \geq 1$, so only the y_0's are independent.

Swaps are easy if they are all single-period, such as the futures settling against $ SOFR, £ SONIA, CHF SARON, or CAD CORRA. Swaps are also easy if the fixing index is also the discounting rate, in most currencies meaning overnight-index swaps ('OIS'). Treat each such swap as a bond, costing 100, with a 100 principal at maturity, each coupon being the swap rate multiplied by that period's daycount-convention year fraction. That set of swap pseudo-bonds allows there to be y_i^{Swap}, which, appropriately connected to $y_i^{\text{Gov't}}$, can be used to shape the early part of the government curve—and if the swap curve is shorter than the government curve, to be shaped by the longest bonds.

An alternative handling of swaps is discussed in §2.17.

2.10 Generalising linkages

There is a possible extension to the idea of linkages, never implemented, which perhaps could be allowed in future Magpie re-codes. However it should be said that, though other Magpie features have been used and found to be sensible and useful, this extension has not been tested, so the author has not seen its behaviour in practice; the author cannot promise that it will always be good.

Recall from §2.8 that Portugal's longest bond is about 30 years, and hence y_1^{PT} is connected to, set equal to, y_1^{ES}. But it might be desirable to impose something of the form $y_1^{\text{PT}} = 0.6\, y_1^{\text{ES}} + 0.4\, y_1^{\text{IT}}$, because it is quite possible that a trader might believe that a Portuguese 50-year, should it exist, would be thus priced. Obviously the imposed constraints must be mathematically consistent, which can be done by requiring: that each y be defined as a weighted sum of x; that all the y_i affected by any particular x have matching i; and that no x be redundant.

It is easier to understand the simpler connections in §2.8–2.9, in which, for each y, one x-weight is +1 and the others are 0. In the generalised version it might assist user understanding of the model's configuration if this were true for as many y as possible, even if not for all y.

2.11 Curve minus curve

Return to Germany, Austria, The Netherlands and Finland. The difference between any two of these countries, the spread, has only two degrees of freedom: one flat term, y_0; and one slope, y_2 with an e-life of 6 years. Hence this can be phrased differently: the model is configured to assume that it would be a tradable mis-pricing for Austria to be steeper than Germany in one part of the curve and flatter in another. There might be a genuine macro-economic reason why Austrian yields might be above or below German yields. And there might be a genuine macro-economic reason why Austria might be flatter or steeper. But this configuration of the model implies that any difference more complicated than that would not be a market statement about reality; it would be a market mis-pricing of a simpler statement. It could happen in the market, but that would be a relative-value mis-pricing and the model 'should' recommend it as a trade rather than fit to it.

Thus embedded in any particular instance of the model are assumptions about what could and could not be priced by an efficient market. Indeed, all parameterised yield curve models contain such assumptions, whether explicitly or implicitly. The Magpie model assumes and enforces that yields are parsimonious, with only five variables per issuer, and that differences between related curves are typically even more parsimonious.

The connections (§2.8) and the values of the z_i (§2.16) are equivalent to assumptions about what a truly efficient market could and could not price. Such assumptions should be consistent with the users' priors, and ideally would also be true, at least mostly.

2.12 Fitting speed

The set of eleven yield curves in §2.8 is described by thirty-two parameters, $x_0 \dots x_{31}$. Nonetheless, it typically takes about four steps to reach machine precision, even from a rough starting guess.

A standard starting guess is to have the curves flat at the yield of that curve's long bond, so each $y_0 \leftarrow \ln[1 + yield_{annual}] \div 365.25$ and $y_i \leftarrow 0 \ \forall i \geq 1$. Or, simpler to code, start with the yield curve flat at a continuously-compounded annual 3.6525% $\Leftrightarrow$ $y_0 \leftarrow 0.0001$, and, $\forall i \geq 1$, $y_i \leftarrow 0$. For the former, the sequence of root-mean-square average of changes of theoretical yields, has resembled ±150bp, ±6bp, ±0.05bp, $\pm 10^{-6}$bp, ≈0 (where 1bp = 0.01% of yield quoted per year).

These steps are occasionally useful for monitoring or debugging the optimisation. With Δy_i being the change in y_i caused by each lap of the Newton-Raphson equation (12), these steps are

$$10^4 \cdot 365.25 \cdot \sqrt{\frac{1}{\text{num. bonds}} \sum_b \left(\sum\nolimits_i \Delta y_i . \frac{\partial thry_b}{\partial y_i} \div \frac{\partial thry_b}{\partial y_0} \right)^2} \tag{26}$$

2.13 Trades, relative-value versus macro

Consider two gilts, the 2½% of 22nd July 2065 ('2H65'), and the 3½% of 22nd July 2068 ('3H68'). These are very similar instruments, having 38.5 and 41.5 years to maturity. And the shorter bond has a smaller coupon, so they have similar values of $\frac{\partial thry}{\partial y_0} \cdot \frac{1}{thry}$. A trade between the two, weighted by $1/\frac{\partial thry}{\partial y_0}$, is therefore a relative-value trade. Magpie can automate this distinction.

Assume two portfolios, p and q, each containing positive amounts of one or some bonds, or of one or some cashflows. For example, p might be one 40-year bond, and q might some 30-year plus some 50-year. Are these two collections of cashflows alike? If they are, then long one and short the other would be a relative-value trade, mostly immune to macro curve movements.

Let the theoretical prices of these collections of cashflows be $port_p$ and $port_q$. Next we need a definition of 'alike', which must be in terms of price risks.

One possibility is that these are defined in terms of sensitivity, through the optimisation, to movements in the fitted $market_b$. But these derivatives are partial not total. Such movements are unnatural: in an actual market, if the yield of a 10-year bond moves higher by a non-trivial amount, then the yields of nearby bonds, perhaps of most bonds, would also move up, so it would be unfactual to assume that only one bond moves.

That suggests derivatives of the portfolios, wrt to the y or the x. There are two ways, neither entirely general. The first possibility is to define:

$$unlikeness(p,q) \ \equiv \ \max_j \left| \frac{\partial}{\partial x_j} (port_p - port_q) \right| \ \div \ \min[\sum_{\text{XX}} \left| \frac{\partial port_p}{\partial y_0^{\text{XX}}} \right| \ , \ \sum_{\text{XX}} \left| \frac{\partial port_q}{\partial y_0^{\text{XX}}} \right| \ ] \tag{27}$$

(Equations (27) and (28) are similar; the variations are highlighted grey.) This is the risk for a move in the most relevant y_i, per ∂_{y_0} risk. (The denominator is well-defined because the portfolios must contain non-negative amounts of bonds or cashflows, but as these bonds are not necessarily priced from the same curve, each portfolio's total y_0 risk is the sum over its risk wrt to the various y_0's.) But if the values within the max were similar at, say, i=2 and i=3, then it is likely that, for a different z between z_3 and z_2, the term would be even larger. Hence the j-discreteness of such a test could fail to catch some values, which would be an unwelcome sensitivity to the z_i (§2.16).

The second possibility is to temporarily increment n by 1, and label the two extra parameters in equation (1) as $z_\star$ and $y_\star$. Set $y_\star \leftarrow 0$, so that the curve is not changed by the increment to n. Then define:

$$unlikeness(p,q) \ \equiv \ \max_{z_\star > 0} \left| \frac{\partial}{\partial y_\star} (port_p - port_q) \right| \ \div \ \min[\sum_{\text{XX}} \left| \frac{\partial port_p}{\partial y_0^{\text{XX}}} \right| \ , \ \sum_{\text{XX}} \left| \frac{\partial port_q}{\partial y_0^{\text{XX}}} \right| \ ] \tag{28}$$

This works better if there is just one curve; but with multiple curves it fails to take any account of linkages.

Which to use? If all the bonds are priced by the same curve, it seems natural to use the maximisation over $z_\star$, equation (28). But if the bonds are priced by multiple curves, then to use the j-maximisation, equation (27). Whichever definition is used, p versus q, in the sense of long one and short the other, is a relative-value trade if $\mathit{unlikeness}(p, q) \leq \mathit{threshold}$, for some chosen *threshold*.

If using equation (28), in practice, for a barbell, the maximum of the numerator usually occurs with $z_\star$ within a moderate factor of the middle maturity. Nonetheless, there is some computational cost to finding these maxima for many trades. But this measure of 'Relative-Value-ness' is only slightly sensitive to the y_i, and only slightly sensitive to ageing of the bonds. Further, small errors would only slightly mis-rank putative trades. Hence it suffices to compute the *unlikeness*() only occasionally.

As guidance, figure 2 holds values of equation-(28) $z_\star$-type *unlikeness* for some cash-neutral ∂_{y_0}-neutral barbells of par semi-annual 4% bonds.

The object is to distinguish relative-value trades from those that are essentially macro in nature. This is not the same as a risk measure: it is about the trading-style nature of the trade. So a trader wishing to automate the RV/macro distinction could do as follows. Inspect the table in figure 2. Are the barbells approximately ordered, more RV at the top, more macro at the bottom? The ordering does not have to be precise, but it must be at least approximate. If not, then this methodology and the required trading style are too different: skip to §2.14. But if the barbells are approximately ordered by RV-ness, choose a preferred *threshold*, and have the trade-finding software select those for which $\mathit{unlikeness}() \leq \mathit{threshold}$.

There is an obvious variation to the process: rather than having a hard boundary on *unlikeness*(), instead use it as a penalty. This would judge by a tougher standard those putative trades that are more macro than those that are more relative-value.

To return to the initial example, with the details as for the barbells, a switch from a 3½% 41.5-year into both a ∂_{y_0}-neutral quantity of a 2½% 38.5-year and a little cash has an *unlikeness* of only 0.006 (again, yields 4% semi). That's less than the *unlikeness* of the fourth barbell in the list, so that switch is correctly labelled as a relative-value trade. As the bonds age shorter the unlikeness of a ∂_{y_0}-neutral switch increases, it becoming $\gtrapprox$0.04 when they have about 17 and 14 years remaining maturity. But a switch of the 3½% 41.5-year into both a ∂_{y_0}-neutral quantity of a 1⅝% 44.5-year and some cash has an *unlikeness* of 0.048. This eight-times-larger *unlikeness* might, or might not, be below the *threshold*: the trader to choose.

A reasonable value might be to set $\mathit{threshold} = 0.04$, and penalise putative trades by $\max[\mathit{unlikeness} - \mathit{threshold}, 0]$. But this is not necessarily 'correct', and a different choice might not be 'wrong', whatever those adjectives might mean.

Similar reasoning could be used by government debt managers. Long-term conventional bonds should be few and large (Wiseman [2006, 2026]), which is not the observed behaviour of most sovereign borrowers. (It is important that short-term debt, say sub-10-year, be approximately evenly distributed. But that doesn't apply to long-dated bonds, which should be few and large: to concentrate liquidity and thereby make more robust the orderliness of markets; to minimise outright and repo dealing costs for issuers and investors and intermediaries; to make prices more difficult to manipulate; and even to simplify repurchases in times of surplus.) So governments should not issue a new long bond unless it is materially different from existing bonds. This can be judged by the same standard: relative to existing bonds any new long bond should satisfy *unlikeness* $\gtrapprox$ 0.04. E.g., if 30 years is the longest bond a country issues, a new 30-year should not be issued until the old has aged down to about 25 years; and if 50 years, down to about 39 years.

Figure 2: *unlikeness* of some barbells

20s25s30s	0.004
30s40s50s	0.005
25s30s40s	0.005
15s20s25s	0.008
25s40s50s	0.008
25s30s50s	0.008
15s25s30s	0.010
20s30s40s	0.011
20s25s40s	0.011
15s20s30s	0.014
10s15s20s	0.015
20s25s50s	0.015
10s20s25s	0.017
15s30s40s	0.018
20s30s50s	0.018
5s7s10s	0.019
7s10s15s	0.021
15s20s40s	0.023
15s25s40s	0.024
10s15s25s	0.027
7s15s20s	0.027
3s5s7s	0.027
15s20s50s	0.029
15s30s50s	0.030
10s20s30s	0.031
15s25s50s	0.034
2s3s5s	0.035
10s15s30s	0.036
7s10s20s	0.036
5s15s20s	0.036
5s10s15s	0.039
10s25s40s	0.040
3s7s10s	0.043
7s20s30s	0.044
10s30s50s	0.044
2s5s7s	0.044
7s15s25s	0.047
10s15s40s	0.049
10s20s40s	0.051
7s25s40s	0.052
5s20s30s	0.054
7s30s50s	0.054
10s25s50s	0.056
3s5s10s	0.056
2s7s10s	0.057
3s10s15s	0.060
5s25s40s	0.061
5s30s50s	0.062
7s15s30s	0.062
5s15s25s	0.063
10s20s50s	0.064
5s10s20s	0.065
3s20s30s	0.065

2.14 Nomenclature

There is some standard nomenclature which users are encouraged to copy. Even if this naming is less than perfect, it has been used and consistency is hereby encouraged.

- There are 'portfolios' of instruments (such as the German Schatz, which have an original maturity of two years).
- One 'curve' might be fitted to several portfolios (e.g., Bubill ∪ Schatz ∪ Obl ∪ Bund).
- Several curves might be fitted together, those curves, the z, and the connections collectively comprising an 'academy'. (It used to be that portfolios were direct members of an academy; and I had envisaged the *Royal Academy of Arts* exhibiting the portfolios of several artists.)
- The horizontal time parameters are called z. The yield parameters at the level of a single curve are called y; at the level of an academy, x.
- The natural measure of a bond's average time, $-\frac{\partial^2 thry_b}{\partial y_0 \partial y_0} / \frac{\partial thry_b}{\partial y_0}$, is called 'tenure'.

2.15 Multiple-curve equation changes

The existence of several sets of y, collectively renumbered as x, necessitate some obvious and minor adjustments to the earlier equations.

- The definitions of $r(t)$ and of $disc()$, equations (1), (3), and (4), use the y appropriate for that curve. Likewise in equation (26), the inner sum is over the y appropriate for that b.
- The derivative multipliers, equation (2), are zero for an x that is not connected to a y in the curve under consideration; and use the 'y_0' definition for an x that is a y_0.
- In the *slide* equation, (5), the $disc()$ and the $r()$ are from the same curve.
- In equation (7) the $error_b$, and its derivatives in equations (8)–(10), use a derivative wrt the y_0 that is from the same curve as that used to price b. And likewise the bond-to-bond sensitivities in relative yield terms in equation (19).
- In equation (11), $\mathbf{H}$ is a matrix of double derivatives wrt to the x, which can be a matrix much larger than (n+1)×(n+1). E.g., with the shown linkages, the euro's $\mathbf{H}$ is 32×32.
- The Newton-Raphson updating rules, equations (12) and (17), are wrt the x rather than the y.
- In computing $\partial thry_b / \partial market_c$, equation (18), the sum can be over x, some terms being zero, or over the y of b's curve.
- In §2.7, $\sum_b loneliness_b$ seems always to be very close the number of x_i to be fitted, that is, to the number of parameters in the whole joint fitting.
- Of course, curves, portfolios, and even individual bonds could be weighted, by importance, by ease of trading, or otherwise. Simply insert into $error_b$ in equation (7) a $weight_b$, the product of the weighting for its portfolio, its curve, and the instrument_b. Perhaps also insert the same multiplier into equation (26), and replace its normalisation with a division by $\sum_b weight_b$. (Weights haven't been useful to me, but YMMV.)

2.16 Choosing the horizontal parameters

As a reminder, in £ and € the z_i were set to $+\infty$, 8766 days (= 24 years), 2191.5 (6 years), 547.875 (1½ years), and 91.3125 (3 months). Whence came these fixed numbers?

The z_i **must not** be fitted to minimise $\sum_b error_b$. Such a fitting would have multiple local minima, and, in practice as well as theory, would be unstable, as explained in Cairns [1997]. Indeed, even a constraint such as $z_i = const \times z_{i+1}$ $\forall i \geq 1$ would not be enough to stabilise the fitting. Such instability loses the power to compute the derivatives wrt market prices, per equations (12)–(20).

The z_i could be chosen to maximise the stability of the x fitting (as evaluated at the y- or x-minimum of the error), perhaps done most easily by minimising the Hessian's condition number, written $|||\mathbf{H}|||$. (Any

of the several definitions of the condition number would do: perhaps the sum of the squares of the elements of the matrix times the sum of the squares of the elements of the inverse of the matrix = the sum of the squares of the eigenvalues times the sum of the squares of the reciprocals of the eigenvalues.) But choosing the z_i on each day has the disadvantage that today's y_i cannot be compared to yesterday's, as they refer to different z_i. Minimising $|||\mathbf{H}|||$ also has the disadvantage that as a bond becomes close to maturity, the z_i are pulled shorter; and when the bond matures, or ceases to be included, then the z_i jump longer. This effect is particularly unwelcome as traders' interest in bonds seems to lessen as maturity approaches, so the effect is caused by a particularly unimportant bond.

Fixing well-chosen z_i allows dynamic historical analysis, for only negligible losses of fitting accuracy and stability. But the z_i must be well-chosen. First all of, this requires that $z_i > 0 \ \forall i$, as that is necessary for the yield curve to have a natural shape. Subject to which, well-chosen means that $\mathbf{H}$ must be very non-singular. That requires that the z_i be different from each other, and different from zero. Obviously, if any $z_i \approx 0$ then that y_i has little effect on anything, making the matrix near-singular. The shortest z at about three months is long enough for its y to have meaning, and seems to have worked well. (Only if there are many instruments shorter than a month can $n = 5$ and $z_5 = 30.4375 \approx 1$ month.) And, also obviously, if $z_i \approx z_j$ with $i \neq j$ then there is near-redundancy between y_i and y_j, again making the matrix near-singular.

Finally, if $z_i \approx +\infty$ with $i \neq 0$ then there is near-redundancy between y_i and y_0. But what does "$\approx +\infty$" mean? If the longest bond is about fifty years, then a 100-year z_1 would cause the effect of y_1 to be very like the effect of y_0. Consider an extreme example: in October 1999 the longest Turkish-lira government bond had a maturity of only two years, a coupon of 72%, and was issued at a yield of 88%. Even if there had been a long bond, its distant payments would have been worth so little that their value would not have been measurable. For such a market, three years was effectively infinite, forward rates that far forward being undetectable.

And WLOG one can impose that $z_i > z_{i+1}$, indeed $z_i \gg z_{i+1}$, the direction determined by 'z_0' $= +\infty$.

(Of course, these restrictions do not apply to the temporary $z_\star$ within equation (28) in §2.13.)

However, having said that the z must not be fitted, slightly, the choice of z can be informed by a partial fit that targets good stability. Assume that an issuer greatly extends its curve: the longest bond was 30 or 50 years, and has become 100 or 200 or 250 years. Keep the z values of 3 months, 1½ years, and 6 years. Increment n by +1, so there are two more z to choose, and test multiple values, perhaps starting near 20 years and 60 years. Test again, having bumped yields ±1% and ±4%; and having aged the curve by a decade using assumptions about future issue patterns. Choose the two extra z such that, across these possibilities, the $\mathbf{H}$ is reliably very non-singular, as measured by $|||\mathbf{H}|||$. Then fix these two new values of z, unless and until the world is much changed.

The z_i are horizontal parameters denominated in time, and should be chosen to span the horizontal structure of the observables. All sets of z satisfying the above non-redundancy constraints, after optimisation of the y (or x) using those z, give similar results, for both fitting and hedging. So any sensible z suffice, and all sensible sets of z are functionally similar.

2.17 Other error functions

On page 5, in §2.5, equation (7) defined $error_b$ to be $\ln[thry_b \div market_b]$, divided by $\frac{\partial thry_b}{\partial y_0} \frac{1}{thry_b}$, all squared, which is analogous to the square of the yield error. Other possible error functions were considered, discussed here.

Could the square of $(thry_b - market_b) \div \frac{\partial thry_b}{\partial y_0}$ have been used? For small errors, it is almost the same thing. But, in the easy case in which n=0, consider what happens if y_0 is very negative. The numerator would be dominated by the very large $thry_b$, and the denominator would be dominated by $-thry_b\times$ days until maturity. So, for $y_0 \ll 0$, such an error function would be approximately constant: machine precision could trap the optimisation in a huge-error state. This trap does not exist in the definition of *error* in equation (7).

In §2.9 there was discussion of fitting some curves to interest-rate swaps, each treated as a bond with a market price of 100. Though bonds have a fixed coupon and are quoted as a price, swaps have a fixed price and are quoted as a coupon rate. Particularly if there is to be fast matrix-multiply re-fitting in the style of

§2.6, it might be desirable to use the quoted 'price' of a swap, which is its conventionally-quoted coupon rate. Let cyf_h be the **c**onventional **y**ear **f**ractions for the fixed payments of an overnight-index swap, such that its fixed payments are $cyf_h \cdot rateMarket$, as paid on days p_h. Let s be the start day of the swap. Then the swap rate precisely consistent with the parameters is

$$rateThry \equiv (1 - disc(s, p_{\text{maturity}})) \div \sum\nolimits_h cyf_h \cdot disc(s, p_h) \tag{29}$$

An error function is needed for this. Echoing equation (7), and desiring units of $(y_0)^2$, could the square of $(rateThry - rateMarket) \div \frac{\partial rateThry}{\partial y_0}$ be used? Alas, for $y_0 \gg 0$ this tends to a constant.

Could $(rateThry - rateMarket)^2$ be used? Its units don't match equation (7), which would disallow use of both bonds and swaps in the same academy. And the units aren't even internally consistent: some of these rates might be quoted one-month, some annual.

This is remedied by converting the rates to daily continuously-compounded. For single-period swaps, the conversion produces, pleasingly, exactly the same error as does (7). But multi-period swaps typically have unequal period lengths. Adjusting for Easter, a year can be as short as 363 days or as long as 370 days, and the disparity can be far greater for a swap with a stub period. So there is not a purist way to convert a multi-period swap to continuously-compounded. Let d_b be the number of days in swap b's longest period, and let cyf_b be the **c**onventional **y**ear **f**raction of that period, or of the first such period. Then define:

$$errorSwap_b \equiv \left(\ln\left[\frac{1 + cyf_b \cdot rateThry_b}{1 + cyf_b \cdot rateMarket_b}\right] / d_b \right)^2 \tag{30}$$

For single-period swaps this precisely equals the value of *error*, yet uses a swap-style price rather than a bond-style price. Of course, use of equation (30) requires single and double derivatives of *rateThry*, and new versions of the *error* derivatives in equations (8) and (9).

Is it worth the effort of using *errorSwap*?

- Bonds, which have fixed coupons, must use *error*. Using *error* for both bonds and swaps, versus using *errorSwap* for swaps, would have almost zero affect on the fitted result. It also has the merit of the programming simplicity of just one error function.
- But if there is to be fast matrix-multiply re-fitting in the style of §2.6, that would need to be driven by the observed change in quoted 'prices': for bonds being prices; for swaps being notional coupons, which needs the presentation of *errorSwap*. Also needed will be an *errorSwap*-style version of equation (10), being $\partial^2 errorSwap_b / \partial y_i \partial rateMarket_b$.
- Equation (7)'s *error* refers to a derivative wrt y_0. For multi-period swaps settling against a fixing index which is not the discounting rate, a more complicated version of equation (29) is needed, referring to both the fixing and discounting curves. With two curves, there are two y_0's, confounding the equation (7) definition of *error*. Hence such swaps must use *errorSwap* from equation (30).

An aside: *errorSwap* is used to minimise, for swaps, the mis-fit of continuously-compounded yields. Could a different *error* do the same for bonds more directly than does equation (7)? Alas! Using notation similar to that of equation (6), the yield of a set of positive cashflows is very close to

$$pseudoYield \;\equiv\; \left(\sum\nolimits_h a_h \cdot disc(s, p_h) \cdot \ln[disc(s, p_h)]\right) \;\div\; \left(\sum\nolimits_h a_h \cdot disc(s, p_h) \cdot (s - p_h)\right) \tag{31}$$

Close, but not identical to: even if *thry* exactly equals *market*, usually *pseudoYield* won't exactly equal *market*'s continuously-compounded yield. Computing the stand-alone continuously-compounded yield for each instance of *thry* entails many inner Newton-Raphson loops; but a whole-curve *pseudoYield* cannot be computed for a stand-alone *market* price. So one definition of yield can be used only for *market*; the other only for *thry*; and they are incomparable. Hence bonds should be fitted to price not yield; for bonds, the equation-(7) *error* is strongly recommended.

3 Problems and features

The Magpie Yield Curve Model has some excellent features. It also has some features which might be thought undesirable. Their importance spans the scale from irrelevant to sometimes material.

3.1 'Problem': approximate fitting

The fitting is approximate, not exact, which, depending on the purpose to which a model is to be put, might be a problem. Therefore the no-arbitrage condition typically does not hold. If Magpie must be used to make no-arbitrage forward rates, as is required of most swap-pricing models, then: in the single-curve case, select precisely $n+1$ bonds; in the multi-curve case, pick a number of bonds precisely equal to the number of x parameters. In either case, choose bonds such that $|||\mathbf{H}|||$ is not much bigger than its lowest possible value.

3.2 Problem: meeting dates

Central banks announce in advance the dates of their meetings. Typically, there is a mid-year announcement of the meeting dates in the following calendar year. Intra-meeting moves are rare—though do happen, for example as done by several central banks in March 2020. Because such off-schedule moves are rare, the expected forward rates between known meeting dates should be almost flat, the bulk of the expected changes happening as discontinuous jumps on scheduled meeting dates. This is inconsistent with the definition of $r(t)$ in equation (1).

Reality is slightly more complicated. Some central banks have meetings with varying features: there are or have been meetings with press conferences, and without press conferences; meetings accompanying an inflation report, and meetings not so. So market participants can believe that some meetings are 'live' and some are not, even though the central bankers sometimes claim that all meetings are live. So one could argue that the true curve should be piecewise almost flat between live meetings—a stronger condition than that in the previous paragraph—and indeed some models for particular situations encode such assumptions.

Whichever, Magpie's forwards are not piecewise almost flat, not even for just the first year. This does not materially affect bonds $\gtrapprox$ 1 year, but does affect meeting-specific pricing.

3.3 Problem: forward guidance

If an interest rate were set very negative, rather than paying this tax on deposits, people and companies might withdraw money in paper form. This has costs—vaults, security, insurance—so policy interest rates can go slightly negative, but very negative would cause deposits to be converted to physical cash.

So a central bank with a very low policy rate might want alternative means of loosening monetary policy. Central banks have sometimes given, and might again give, 'forward guidance'. That can be conditional guidance: "the policy rate will remain $\leq r\%$ until inflation exceeds our target". That can be date guidance: "the policy rate will remain $\leq r\%$ until some particular date". The latter form can resemble there not being a 'live' meeting for some years.

E.g., US Federal Reserve, 16 Dec 2008, "the Committee anticipates ... low levels of the federal funds rate for some time" (Bernanke [2008]). E.g., ECB, 4 July 2013, "The Governing Council expects the key ECB interest rates to remain at present or lower levels for an extended period of time" (Draghi [2013]). E.g., BoE, 14 Aug 2013, "The Committee intends not to raise Bank Rate from its current level of 0.5% at least until the Labour Force Survey ... unemployment rate has fallen to a threshold of 7%" (Carney [2013]). E.g., US Federal Reserve, 29 Apr 2020, "we expect to maintain interest rates at this level until we are confident that the economy has weathered recent events and is on track to achieve our maximum employment and price stability goals" (Powell [2020]). E.g., Norges Bank, 6 May 2020, "In the Committee's current assessment of the outlook and balance of risks, the policy rate will most likely remain at today's level for some time ahead. We do not envisage making further policy rate cuts" (Olsen [2020]).

And, as with §3.2, these can be inconsistent with $r(t)$ in equation (1). Magpie cannot price a sharp corner in the forward rate path on the expiry date of the forward guidance. But Magpie can fit around this, so would be only slightly wrong.

Of course, Magpie is a model. Its output can be a useful set of trade ideas, and has been in many currencies over many years. But it is a model: models simplify; models sometimes over-simplify. It is a model, not an omniscient trading algorithm to be used unthinkingly.

3.4 Problem: yield curve control

With rates materially above zero, central banks have typically set a short-term rate, which motivated the structure of this and other yield curve models. But near the lower bound central banks have sometimes tried to cap longer rates, and this has also happened in time of war (Garbade [2020]). To be specific, a central bank might undertake to buy some, or an unlimited amount, of the $\leq m$-year bonds at a yield of $r\%$. E.g., the Bank of Japan (Kuroda [2016]), and the Reserve Bank of Australia (Lowe [2020]). If market yields are far below $r\%$, then this has little effect. But if market yields are above that, the central bank would buy much or most of the qualifying bonds.

There are multiple possible conclusions.

- Magpie shows the m-year bond as expensive, and should. The central bank is over-paying for this bond, and those who own it should sell it. So the model works. But the model will show as cheap some of the bonds shorter than m years, even though they are actually expensive. So the model works only in part.
- The central bank has bought and locked away almost the entirety of some bonds. Their 'prices' are therefore non-traded prices of instruments which are no longer accessible to the market. Such bonds should not be in the model at all.

  But, awkwardly, not all central banks reveal how much of each bond has been bought. Estimates can be substituted, but there can be material uncertainty around such estimates.
- Some governments, such as the USA until 1984 (Carfine [2009]) and the UK until 1994 (Plenderleith [1995]), issued some bonds that had embedded optionality. E.g., the 5½% 2008–12 gilt had a final maturity of 10th Sep 2012, but the government could call it—mature it early—on any date from 10th Sep 2008. Magpie is not an optionality model, so such bonds have always been excluded. And if the central bank gives away an option for some bonds, perhaps those bonds should likewise be excluded.

  But this is also awkward: a bond with $m+\epsilon$ years to maturity does not yet have any such optionality. It might soon have such optionality; but the possibility of a change in monetary policy means that it might not.

My preference is to exclude bonds that do not have a true market price. That is, if the central bank is known or estimated to have bought the bulk of a bond, then that bond's price is too much of an accounting fiction. But also, if bonds are being held mostly by those planning to soon sell them to the central bank, the bonds' prices are dominated by an imminent sale to an inefficient player. It must be sensible to exclude all such bonds—even though the maturity boundary of the latter type is a judgement call. If that means that there are no good prices for short-dated securities, then that part of the curve should be shaped by interest-rate futures and swaps (§2.9).

3.5 Problem: imperfect implementation

Central banks' implementation of monetary policy is not perfect. It should be easy—a narrow passive corridor—but the more complicated systems chosen are imperfectly reliable. In particular, in short-term rates there can be seasonality (weekly, monthly, meeting cycle, quarterly, annual) or other quirks, each of which is likely to be inconsistent with the $r(t)$ in equation (1).

3.6 Problem: weekends

Actual interest rates do not compound continuously. Indeed, actual interest rates do not compound on any precisely fixed frequency.

Assume the relevant interest rate is 5%, quoted £-style Act/365. But this is quoted 'simple', one business day at a time. So over a Friday-to-Monday weekend money would multiply by $1 + 5\%\times3\div365$, whereas over three good days days it would multiply by $(1+5\%\times1\div365)^3$. These differ by 56.30 per billion, or 7203 per billion per year. Or, for the simple weekend to equal compounded weekdays, the weekend rate would have to be $((1 + 5\%\times1\div365)^3 - 1)\times365\div3 \approx 5.00068496\%$.

3.7 Problem: the possibility of default

Magpie is an interest-rate model. It assumes that currency is fungible. It assumes that one currency unit payable by a bond on a particular date is worth the same as one currency unit payable by that same bond on that same date. For interest-rate purposes, this is a good assumption.

But for bonds with credit risk this is not necessarily true.

Corporate defaults have a fairly well-defined process. Generally, payments are accelerated. All bonds become payable immediately, and future interest payments aren't payable. As there usually isn't enough money to pay everybody, bond holders get only a fraction of their principal, such as 60 cents on the dollar. So consider two bonds, with the same seniority and same maturity date, one with a coupon of 2%, the other with a coupon of 8%. In default, they pay the same reduced principal. In default, the extra 6% of coupon is worth nothing.

Sovereign defaults haven't always worked this way, and won't necessarily work this way in the future. Also sovereign defaults can be made more complicated by the existence of STRIPS, and even more complicated if the STRIPS have coupon-principal fungibility.

Magpie doesn't see any of these difficulties: it treats money from principal and money from coupon as fungible money, and payable on the scheduled date, even though in default neither is necessarily true.

Therefore if the market is pricing a material risk of default, high-coupon bonds should be priced cheap to their *thry*, and low-coupon should be priced expensive. This can be treated as a post-processing problem: fit assuming all bonds are money good; then regress rich-cheap money $= \ln[market_b/thry_b]$ or rich-cheap yield $= -\ln[market_b/thry_b] \cdot thry_b \div \frac{\partial thry_b}{\partial y_0}$ against high-coupon-ness = either $clean_b - 100$ or $nominalCoupon_b - yield_b$.

3.8 Problem: tax

In some jurisdictions, at some times in the past and perhaps in the future, market pricing has been affected by tax. When, for example, capital gains are taxed less than are income, then a currency unit from principal can be worth more than a currency unit from coupon. Magpie, for the same reasons as the §3.7 credit argument, does not see this: one unit of money is one unit of money.

There are two possible mitigations. One could pick a type of market participant, ideally the type believed to dominate price formation, and adjust each bond's cashflows to match that tax status. Or there could be separate curves for low-, medium- and and high-coupon, per Dobbie and Wilkie [1978]. But these are imperfect kludges.

3.9 Problem: repo specialness

It is possible to use a bond as collateral to borrow money. For most bonds, the price of such borrowing is determined by matters such as monetary policy. But some bonds can be 'special': market participants so wish to borrow the bonds that the lender of the bond, the owner, is remunerated by borrowing money more cheaply than if posting a 'general collateral' bond. A possible fix is to adjust the price of special bonds by subtracting the present value of an estimate of all such future specialness. Without that Magpie will typically show such bonds as expensive, as they are to those unable to use repo.

3.10 Problem: pricing inefficiency

Magpie, for all its subtleties, is about pricing sets of cashflows against a smooth forward rate curve. But sometimes bonds do not behave solely as bundles of cashflows.

For example, in Oct 2001 the US Treasury announced that it would stop issuing 30-year bonds (Fisher [2001]); resumption being announced in Aug 2005 (Bitsberger [2005]). Outside this period usual practice had been up to four 30-year bonds a year, maturing on the 15th of the middle month of each quarter. Except that there were none between the 5⅜% Feb 2031 and the 4½% Feb 2036. For the subsequent 1½ decades the Feb 2031 and Feb 2036 traded expensively; as did, to a lesser extent, nearby bonds. As well as being bundles of cashflows, these bonds were priced as if they had a rarity value. Magpie does not perceive rarity value, and reported the bonds as expensive.

Magpie has been correct: these have been expensive bundles of cashflows. But Magpie has also been wrong: for much of the time shorting these bonds would not have been profitable.

Again in the US, newly issued US Treasuries are often expensive to the 'old' and multiple-'old' securities. These new securities are special in repo, but often not sufficiently special to offset the expensiveness. They have a different benefit: they are extremely liquid, it being possible to buy or sell large quantities for minuscule dealing costs. For a long-term holder of bundles of cashflows, these bonds are expensive, so Magpie is correct in some sense, even though not necessarily correct for a very-short-term holder.

Indeed, in the author's judgement, because of these two persistent features, US Treasuries have been the market for which Magpie has been least useful.

In the UK, *circa* 2010–2022, gilts near 40 years traded expensively to 30s and 50s. Magpie, and also other measures of rich-cheap, all saw this, yet it persisted, probably because pension funds needed to match regulation-defined liabilities. Investors with very long-term horizons would have profited by selling their 40s into a mix of 30s and 50s; but for hedge funds this anomaly was too persistent to pay the repo costs.

The Magpie model values bundles of cashflows. But sometimes bonds also have other qualities: ease of trading; matching the regulatory measure of liabilities; being deliverable into a futures contract; tax advantageousness; etc. Usually those not needing the other qualities should sell the bonds deemed expensive by Magpie, and buy the cheap. But that does not promise that the marginal investor is indifferent to the other qualities, so does not promise that any expensiveness will quickly lessen. As explained at the end of §3.3, the purpose of models, and their flaw, is that they simplify.

3.11 Problem: non-optimised choices

The choice of the z_i (§2.16), and the choice of the connections (§2.8), are at least partly manual. They don't have to be wholly manual: they can be informed by numerical tests asking which putative connection would most shrink $|||\mathbf{H}|||$. But the connections are constraints on the model, and these constraints should map to reality (§2.11), which needs a judgement about how complicated reality really can be.

3.12 Problem: plotting par curves

The final problem is a fault of almost all yield curve models. Of course, bonds must be valued using their actual cashflows, on their actual holiday-adjusted dates of payment. But for plotting a par curve, quarterly or semi-annual payments can introduce the appearance of a wave, which would be an artefact of the arbitrarily assumed payment dates. E.g., 22nd July to 22nd January is 184 days; non-leap-year 22nd January to 22nd July is only 181 days. These artefacts are more obvious with yields higher rather than lower, 8% rather than 1%, but nonetheless are generally unwelcome. The cleanest solution, and it is emphasised that this is for the plotting of curves rather than the valuation of traded instruments, could be to show, converted to the relevant compounding convention, the par yield = coupon of continuously-paying bonds:

$$365.25 \times \left(1 - disc(s,t)\right) \div \left(\int_s^t disc(s,t')\ \mathrm{d}t'\right) \tag{32}$$

where s is the standard settlement day for the most important instruments in this academy.

The fault is that, alas, the integral lacks analytic form. Indeed, the integral lacks analytic form even in the special case of $n = 1$, implying that all the forward-rate models in §1.4 have this fault.

One possibility that suffices for plotting is to use weekly payments, and then post-process to remove points that are sufficiently close approximations to a straight line between remaining neighbours. This works, but seems ugly for a model in which so much is analytic.

A more elegant possibility is to estimate the integrals piecewise, by approximating $disc(s,t)$ with an odd-degree polynomial. Assume this piece spans $t_s \le t \le t_e$.

- Known are both $disc(s,t_s)$ and $disc(s,t_e)$, which allows a linear approximation to $disc(s,t)$.
- Known are both $r(t_s)$ and $r(t_e)$, and that $\frac{\partial disc(s,t)}{\partial t} = -r(t) \cdot disc(s,t)$, a total of four constraints which allows a cubic approximation.

- Known are both $\frac{\partial r(t_s)}{\partial t}$ and $\frac{\partial r(t_e)}{\partial t}$, which allows a quintic.

One could go further, but in practice the quintic is sufficient. Using reasonable example parameters, for an interval from 3.25 to 3.5 years, the quintic's integral error $\approx \pm 10^{-9}$.

4 Good qualities

Despite the 'problems' listed in §3, the Magpie Yield Curve Model has some very good qualities.

i: Equation (1) spans the full range of plausible smooth yield curves, and not much more than that.

ii: Magpie can share data across related curves: the early part of one issuer's curve being shaped by another issuer's T-Bills (§2.8), or by interest-rate futures or by swaps (§2.9).

iii: The discount between any two dates has a simple analytic form, and is arbitrarily differentiable (equations (2) to (5) in §2.4).

iv: The existence of analytic derivatives reveals much about the implied structure of the market, including the *unlikeness* measure of whether being long one portfolio of bonds versus short another portfolio is relative-value or macro (§2.13).

v: The fitting has a unique local minimum, and because the error function is nearly parabolic the fitting is very fast and very stable (§2.5 equation (7); discussion underneath §2.5 equation (15); and §2.12).

vi: The error function being within the model reveals the derivative, through the optimisation, of the parameters with respect to the market prices (§2.5, §2.6).

vii: These derivatives allow very fast re-optimisation by matrix multiplication (§2.6).

viii: These derivatives allow even faster repricing from the price changes of a few benchmarks (§2.6).

ix: And these derivatives through the optimisation allow computation of each input instrument's '*loneliness*' (§2.7). In effect this is the number of parameters absorbed by each fitted instrument, and hence the number of parameters absorbed by any set of instruments.

x: The fitted y_i and x_i have a consistent meaning (§2.16). Hence day-to-day or week-to-week changes are meaningful, and regression or eigenvector analysis of these changes in y_i or x_i can further reduce the dimensionality of risk.

xi: In currencies big and small, developed and emerging, existing and pre-euro, whether the maximum maturity be long or less so (final ¶ of §2.2), Magpie has been used and found sensible, stable, and worthwhile.

Acknowledgements

Over the years the model has had help and discussion from many people, not all remembered. Particularly worthy of mention are David Theobald, my boss at J. P. Morgan who thought it good and in late 1997 allowed its implementation; and Brian Barney who coded it in C.

Further help came in the form of reading drafts of this paper, and suggesting improvements to it, most recently from Nick Dawson, Rodrigo de Azevedo, and Tharun Puvaneswaran.

References

B. S. Bernanke. FOMC statement, December 2008. URL `http://www.federalreserve.gov/newsevents/pressreleases/monetary20081216b.htm`.

T. S. Bitsberger. August 2005 Quarterly Refunding Statement, 2005. URL `http://www.treasury.gov/press-center/press-releases/Pages/20058311511122645.aspx`.

A. J. G. Cairns. Stability of descriptive models for the term structure of interest rates. *British Actuarial Journal*, May 1997. URL `http://www.semanticscholar.org/paper/8e0afaf6acbaf36b940b0bbaa1e95b5316d292b4`.

A. J. G. Cairns and D. J. Pritchard. Stability of Descriptive Models for the Term Structure of Interest Rates with Application to German Market Data. *British Actuarial Journal*, 2001. URL http://pdfs.semanticscholar.org/8e0a/faf6acbaf36b940b0bbaa1e95b5316d292b4.pdf.

K. E. Carfine. Treasury Calls 11–3/4 Percent Bonds of 2009-14, 2009. URL http://web.archive.org/web/20090902140527/https://treasurydirect.gov/news/pressroom/pressroom_comcall071509.htm.

M. Carney. Minutes of the Monetary Policy Committee Meeting held on 31 July and 1 August 2013, August 2013. URL http://www.bankofengland.co.uk/minutes/2013/monetary-policy-committee-august-2013.

G. M. Dobbie and A. D. Wilkie. The F.T.-Actuaries Fixed Interest Indices. *Journal of the Institute of Actuaries*, 105(1):15–26, 1978. ISSN 00202681.

M. Draghi. Introductory statement to the press conference, July 2013. URL http://www.ecb.europa.eu/press/pressconf/2013/html/is130704.en.html.

P. R. Fisher. Remarks at the November 2001 Quarterly Refunding, 2001. URL http://www.treasury.gov/press-center/press-releases/Pages/po749.aspx.

K. D. Garbade. Managing the Treasury Yield Curve in the 1940s. *Federal Reserve Bank of New York*, (913), 2020. URL http://www.newyorkfed.org/research/staff_reports/sr913.

H. Kuroda. New Framework for Strengthening Monetary Easing, September 2016. URL http://www.boj.or.jp/en/announcements/release_2016/k160921a.pdf.

P. Lowe. Monetary Policy Decision, March 2020. URL http://www.rba.gov.au/media-releases/2020/mr-20-08.html.

C. Nelson and A. F. Siegel. Parsimonious Modeling of Yield Curves. *The Journal of Business*, 60(4): 473–89, 1987. URL http://EconPapers.repec.org/RePEc:ucp:jnlbus:v:60:y:1987:i:4:p:473-89.

Ø. Olsen. Rate decision May 2020, May 2020. URL http://www.norges-bank.no/en/topics/Monetary-policy/Monetary-policy-meetings/2020/may-2020/.

I. Plenderleith. 7% Treasury Convertible Stock 1997, 1995. URL http://web.archive.org/web/19990224210411/http://www.coi.gov.uk/coi/depts/GBE/coi8632a.ok.

J. H. Powell. Transcript of Chair Powell's Press Conference, April 2020. URL http://www.federalreserve.gov/mediacenter/files/FOMCpresconf20200429.pdf.

L. E. O. Svensson. Estimating and Interpreting Forward Interest Rates: Sweden 1992 - 1994. NBER Working Papers 4871, National Bureau of Economic Research, Inc, September 1994. URL http://ideas.repec.org/p/nbr/nberwo/4871.html.

J. D. A. Wiseman. The Exponential Yield Curve Model. Technical report, J. P. Morgan, 1994.

J. D. A. Wiseman. The Magpie Yield Curve Model. Technical report, J. P. Morgan, 1998. URL http://www.jdawiseman.com/papers/finmkts/19980224_Magpie_JPMorgan.pdf.

J. D. A. Wiseman. Bonds: too many; too small, 2006. URL http://www.jdawiseman.com/papers/finmkts/reopenings.html.

J. D. A. Wiseman. On The Plotting of Yields, 2008. URL http://www.jdawiseman.com/plotting_yields.html.

J. D. A. Wiseman. The Magpie Yield Curve Model at SG. Technical report, Société Géneral, 2012. URL http://www.jdawiseman.com/papers/finmkts/20120301_Magpie_SocGen.pdf.

J. D. A. Wiseman. "Treasury market structure": Fewer and Larger, 2026. URL http://www.jdawiseman.com/papers/finmkts/Fewer_Securities.html.

LaTeX compiled to PDF on 2026-11-23 at 15:30.