London
24 February 1998

J.P. Morgan Securities Ltd.
European Fixed Income Research
Julian Wiseman (44-171) 325-5823
wiseman_j@jpmorgan.com

JPMorgan

European Fixed Income Research

The Magpie Yield Curve Model

- **The Magpie Yield Curve Model provides a concise description of the state of interest-rate term structures**
- **The model can be used to price, hedge and calculate relative value and forwards for swaps, government bonds and strips, and portfolios of credit instruments**
- **The model is extremely frugal with information, so can be used even if there are very few bonds, and is fast enough to be reoptimised in real-time**
- **The model is provably stable, so hedges and extrapolations behave smoothly**

In December 1994 J.P. Morgan introduced The Exponential Yield Curve Model. The daily runs of this model have proved useful to traders and market-makers, by being effective at identifiying profitable relative value opportunities within a single curve. The model now constitutes part of the basic trading equipment for European government bonds.

However, whilst the model has served well at describing relative value based on a single day's prices in a single curve, it has not been satisfactory at describing curve dynamics, and has not given rise to stable hedges. Further, it requires a large number of input bonds and, in markets where there are no prices for short-dated paper, it does not always produce a natural short curve.

To remedy these defects we have introduced **The Magpie Yield Curve Model**, which is a successor to the (old) exponential model. This has a number of favourable qualities:

- As with the old exponential model, it fits an extremely smooth instantaneous forward rate curve to market prices;
- This means that the theoretical zero-coupon and par curves are extremely smooth;
- It is stable, in that a small change in the price of an instrument can only ever give rise to a small change in the value of a parameter;
- The hedges are almost independent of market price, and so are stable;
- Information can be shared between curves. For example, the prices of money-market instruments can be used to ensure that the short-government curve is sensible, while the long-dated governments can help with the extrapolation of the swap curve;
- Because information can be shared, it is possible to fit a smooth curve using very few bonds, the "missing" information being lifted from a different curve in the same currency. This is particularly useful for fitting credit curves.

We start by describing the old model, as the new model is built on the old, and explain the cause of the flaws in the old model. We describe how the new model works in the one portfolio case, and show that the new model has the stability properties absent in the old. We then describe how a number of portfolios can be linked, to improve the fitting in all of them.

The Old Model

At the heart of both the old model† and the new model is a description of the instantaneous forward rate, which is quoted continuously-compounded. This rate, $r(t)$, is parameterised as

$$r(t) = y_0 + y_1 e^{-t/z_1} + \ldots + y_n e^{-t/z_n}$$

† This equations in this description of the old model might not be visually the same as those in the original December 1994 paper, but are mathematically equivalent to them. The parameters have been renamed and re-arranged to coincide with those in the new model — but the changes are purely presentational.

Morgan Guaranty Trust Company
European Fixed Income Research
Julian Wiseman (44-171) 325 5823
wiseman_j@jpmorgan.com

with $z_1 > z_2 > \ldots > z_{n-1} > z_n > 0$. Time, t, is denominated in days (to avoid difficulties with leap years), and the earliest settlement date is deemed to be $t = 0$.

The forward rate curve can be thought of as being described by a linear sum of "factors". The first factor is flat, the others decay with $^1/_e$ lives equal to the z's. If n=2, then y_0 can be thought of as describing long-dated forward rates, y_1 as controlling the medium-to-long curve, and y_2 as controlling the short-to-medium curve. The two z's determine the temporal location of the soft 'fade-outs' between short and medium, and between medium and long.

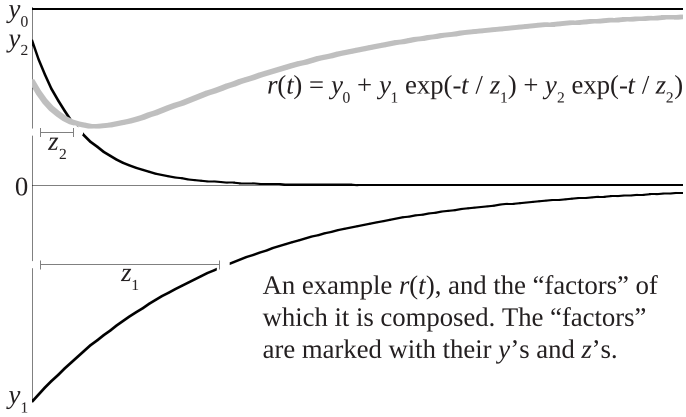

An example $r(t)$, and the "factors" of which it is composed. The "factors" are marked with their y's and z's.

The equation for $r(t)$ implies that the zero-coupon rate from settlement s to maturity t is

$$zero_s(t) = y_0 + y_1 \frac{z_1(e^{-s/z_1} - e^{-t/z_1})}{t-s} + \ldots + y_n \frac{z_n(e^{-s/z_n} - e^{-t/z_n})}{t-s}$$

which is also linear in the y's. This in turn gives the discount from some date t back to settlement s

$$disc_s(t) = \exp\Big(y_0(s-t) + y_1 z_1(e^{-t/z_1} - e^{-s/z_1}) + \ldots + y_n z_n(e^{-t/z_n} - e^{-s/z_n})\Big)$$

Naturally, the theoretical price of a bond is the sum of its discounted cashflows.

In the old model the y's and z's were chosen to minimise a fitting error, calculated from the market prices and theoretical prices of each instrument. (Each instrument was deemed to be a bond, paying fixed cashflows, and costing a dirty price payable on a settlement day.) We define $E = E(y_0, \ldots y_n, z_1, \ldots, z_n, \textit{market prices})$ to be

$$E = \sum_{k \in \text{bonds}} \left(\frac{\ln\left(\frac{theory_k}{market_k}\right)}{\frac{\partial theory_k}{\partial y_0} \frac{1}{theory_k}} \right)^2$$

This formula is almost equal to the sum of the square of the yield errors:

- The numerator (the log of the ratio of the prices) is the price error, as a proportion;
- The denominator ($\partial \text{theory}_k / \partial y_0 \times 1/\text{theory}_k$) is the proportionate sensitivity to parallel moves in the forward rate curve, and is hence a form of duration;
- So the ratio of these two is a yield change and hence a yield error, and E is just the sum of the square of these yield differences.

This particular form for the error function is computationally slick, and as we shall see later, also mathematically convenient. The new model uses the same error function.

Recall that the y's and z's were chosen to as to minimise the error. This minimisation has proved to be extremely difficult, as there is a large region of the parameter space (many different values of the y's and z's) in which the error is close to its global minimum, and that this large region contains many local minima. Of course, one of these local minima is the global minimum for the error: but which? This is chaotic: a small change in the prices (and hence a small change in the shape of the error as a function of the y's and z's) can change which of the local minima is the global minimum (see chart).

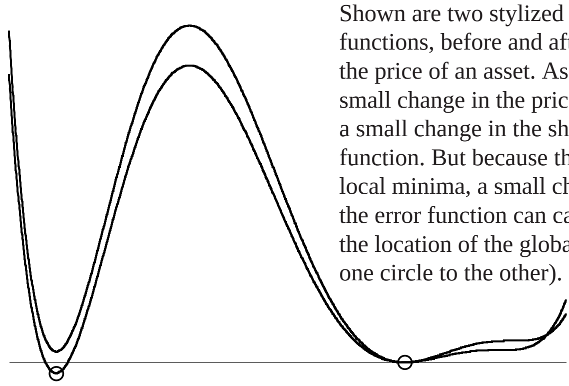

Shown are two stylized one-parameter error functions, before and after a small change in the price of an asset. As one would expect, a small change in the price of an asset causes a small change in the shape of the error function. But because the error has multiple local minima, a small change in the shape of the error function can cause a large jump in the location of the global minimum (from one circle to the other).

Such jumping from one local minimum to another necessitates only a small change in the shape of the yield curve and only a small change in the level of the error, but it can cause a large change in the values of the parameters†. In effect, there are substantially different parameter values that give similar values for the theoretical prices of the bonds, and hence give a similar value for the error.

The jumping parameter values still give reasonably stable values for all the fitted numbers — these being the prices of the bonds. The stability of the fitted numbers implies that relative value measures are also stable to jumps in the parameter values. However, the jumping parameter values can cause some non-fitted numbers to jump as well. Non-fitted numbers typically include the level of yields at maturities shorter than that of the shortest instrument, the level of yields at maturities longer than that of the longest instrument, and also the hedging matrices. Thus the instability in the old model has not hindered its role in relative value finding, but has prevented the model being used to describe hedges, to extrapolate the curve and to describe curve dynamics.

The new model eliminates the parameter jumps, and so can be used for hedging and curve extrapolation as well as relative value.

The new model has a second benefit. Information is embedded within the model in a manner that allows it to be shared. So, for example, money-market instruments can help form the shape of the early part of the government curve, and the government curve can help form the shape of long part of the swap curve. One can build two credit curves for a corporate issuer, one senior and one subordinated, these two curves sharing some information with each other, and copying any "missing" pieces of information from the swap or government curve.

First we describe the model in the special case of there being one portfolio, and then we describe how multiple portfolios can be linked.

Magpie: the one portfolio case

As with the old model, the new model describes the instantaneous forward rate quoted continuously-compounded, $r(t)$ as:

$$r(t) = y_0 + y_1 e^{-t/z_1} + \ldots + y_n e^{-t/z_n}$$

where time, t is denominated in days and the trade date is deemed to be t=0.

The error function is the same as for the old model, so is again effectively the sum of the square of the differences between the market and theoretical yields:

$$E = \sum_{k \in \text{bonds}} \left(\frac{\ln\left(\frac{theory_k}{market_k}\right)}{\frac{\partial theory_k}{\partial y_0} \frac{1}{theory_k}} \right)^2$$

What is different is the method of fitting. Rather than a single stage fitting to minimise the error, the fitting is done in two stages: an inner stage, in which the y's are determined whilst the z's are known and fixed, and an outer stage in which the z's are determined.

The inner stage is very simple. At this stage z's are known and fixed, and only the y's are varied. The y's are chosen to minimise the error. Although the joint y- and z-minisation of the error in the old model is fraught with difficulties, a purely y-minimisation, with z's held constant, is quite different. The y-minimisation of E has a unique local minimum, and that unique local minimum is also the global minimum. The y-location of this global minimum (described by the optimised y's) moves continuously and stably as bond prices are varied continuously. In fact, not only has the fitting these desirable qualities, the error as a function of the y's is very close to being quadratic, so the fitting is also very fast.

So how are the z's determined? The z's cannot be chosen to minimise the error, as this would then be equivalent to the old model. Instead we recall that $r(t)$ was chosen to be a linear sum of factors. We have already said that the amounts of each factor (the y's) are chosen to minimise the error. So asking how to choose the z's is equivalent to asking how to choose the factors.

† The author thanks Andrew J. G. Cairns MA PhD FFA for drawing attention to these "catastrophic" parameter jumps. For the UK gilt curve Dr Cairns suggests that n=4 and that the z's be fixed at 5 years, $2\frac{1}{2}$ years, $1\frac{1}{4}$ years and $\frac{5}{8}$ of a year *(Proceedings of the 6th AFIR International Colloquium, Nuremberg, October 1996).*

Morgan Guaranty Trust Company
European Fixed Income Research
Julian Wiseman (44-171) 325 5823
wiseman_j@jpmorgan.com

We want the factors to be different, because if two of the z's are the same, then the y's associated with these z's become indeterminate, (in that one can add anything to one of them and subtract the same from the other without altering the curve). Likewise, if any z is so large as to be effectively infinite, then that factor becomes the same shape as the y_0 factor (the y_0 factor is flat, as if there was a z_0 equal to $+\infty$). If any z is close to 0, then that factor becomes zero everywhere, and hence irrelevant.

In most problems one asks that the factors be in some sense "orthogonal". But directly measuring the orogonality of these factors requires a definition of the inner-product of two factors. In the absence of a natural definition of orthogonality, any such definition would be arbitrary.

Instead define $\mathbf{H} = \mathbf{H}(y_0, \ldots\ y_n, z_1, \ldots, z_n,$ *market prices*) to be the matrix of second partial derivatives of the error wrt the y's:

$$\mathbf{H} = \begin{pmatrix} \frac{\partial^2 E}{\partial {y_0}^2} & \frac{\partial^2 E}{\partial y_1 \partial y_0} & \cdots & \frac{\partial^2 E}{\partial y_n \partial y_0} \\ \frac{\partial^2 E}{\partial y_0 \partial y_1} & \frac{\partial^2 E}{\partial {y_1}^2} & \cdots & \frac{\partial^2 E}{\partial y_n \partial y_1} \\ \cdots & \cdots & \cdots & \cdots \\ \frac{\partial^2 E}{\partial y_0 \partial y_n} & \frac{\partial^2 E}{\partial y_1 \partial y_n} & \cdots & \frac{\partial^2 E}{\partial {y_n}^2} \end{pmatrix}$$

If the factors were completely independent then this would be a diagonal matrix. If the factors were in any of the "bad" cases described above ($z_i \to 0$, $z_i \to \infty$, $z_i \to z_j$ with $0<i<j$) then this matrix would be singular. So our intuition about the z's suggests that we want $\mathbf{H}$ to be as diagonal as possible, and as far from singular as possible.

The natural measure of "singularity" is the condition number† of $\mathbf{H}$, written $|||\mathbf{H}|||$. As we want to choose the z's to minimise the singularity of $\mathbf{H}$, we must choose the z's so that $|||\mathbf{H}|||$ is minimised. This minimisation of $|||\mathbf{H}|||$ has the effect of separating the z's from each other, from zero, and from infinity, so that they span the timescale of the market. For example, with n=5 in sterling, $z_1 \ldots z_5$ are approximately 15 years, $2\frac{1}{2}$ years, 6 months, 5 weeks and 4 days‡.

So, in summary, the one portfolio case works as follows. The z's are guessed. With the values of the z's fixed, the y's are chosen so as to minimise E. With the y's at this E-minimum, $\mathbf{H}$ and from this $|||\mathbf{H}|||$ are calculated. The z's are then repeatedly reguessed until this $|||\mathbf{H}|||$ (evaluated at the y-minimum of E) is minimised.

There is an intuitive way to see this. Essentially we do a two stage fitting, and these stages are almost independent.

- Firstly, the z's are fitted to the temporal information structure of the market. So in an emerging market with no instruments longer than a year, the z's will all be measured in weeks and months rather than years, and this is true irrespective of whether yields are 3% or 3000%. In a market with longer-dated instruments, the longer z's will be measured in years, so that the z's "span" the market.
- Second, with the z's chosen to "fit" the information structure, the y's are fitted to market prices. Different levels or shapes of yield curves will result in different y's.

The two fittings can be separated (at least conceptually) because $\mathbf{H}$ is almost independent of the y's (and actually independent if the market consists only of single-cashflow instruments).

The cashflow structure of the market does not vary during the day (excepting new issuance, which we ignore for these purposes), and so the lengthy calculation of the z's can be done overnight, the z's being held constant during the day. However, prices and yields do vary, so the y's must be varied in real-time. Indeed, because the z's are fixed there is no need to hedge any implied z exposure; only y exposure need be hedged. All instruments can therefore be repriced in terms of, and hedged using, only n+1 benchmarks (the n+1 varying pararameters being y_0 to y_n). With n=5 there will need to be 6 benchmarks, and (for example) these might well be overnight money, a 3-month interest rate future, a 2-year, a 5-year, a 10-year, and also the long bond.

† We define the condition number $|||\mathbf{H}|||$ to be $(\sum\lambda_i^2)(\sum\lambda_i^{-2})$, where the λ_i are the eigenvalues of $\mathbf{H}$.

‡ The shorter z's, especially z_4 and z_5, are quite sensitive to the selection of money-market instruments included in the portfolio. We include 1-week, 2-week and 1-month interbank money, and exclude securities with less than 35 days to maturity. The interbank and government portfolios are connected as described in the next section.

There is a simplification that we often apply to this. Because the z's vary little, it can be more convenient just to fix them at their well-separated values. This simplifies historical analysis of the y's.

As previously remarked, the model allows different portfolios of instruments — which will have different fitted curves — to share some information. Let us now discuss how this works.

Magpie: multiple portfolios

Let us consider the sterling interest rate market, which consists primarily of government debt, London deposits (LIBOR) and LIBOR-linked derivatives (including interest rate futures, FRAs, and swaps), and also an overnight index called SONIA with swaps against it.

The UK authorities never issue coupon paying debt shorter than 3 years, and in practice issue sub-4-year paper extremely rarely. T-Bills almost never trade in the secondary market, so sub-1-year there are only old non-par gilts that rarely trade. Beyond 3 years the government market (out to the 6% Dec 2028 and the Jun 2021 principal strip) is active and liquid.

So price discovery in short dated instruments tends not to take place in government debt: short-dated gilt prices are driven by LIBOR products rather than by price discovery in gilts themselves. We would like our yield curve model to reflect this, in that we would the short-dated government curve to be (somehow) guided by the interbank markets.

The old exponential model, and indeed most other curve models, keep portfolios of non-matching instruments completely separated. The new model links portfolios in a manner that allows exactly this sharing of information.

First, within any one currency all portfolios share a common set of z's. This allows the y's (the amount of each factor) in different portfolios to be directly comparable, because these factors have the same shape. Thus information about yields (about the y's) thus become fungible between portfolios.

Second, we use this fungibility. Recalling that, in sterling, $z_5 \approx 4$ days, let us consider the role of y_5. This quick-decaying term will effect the shape of the sub-1-month curve, but will have a negligible effect beyond this. Because there are money-market instruments (such as deposits) with such short maturities, the error function (for the swap portfolio) will increase quickly if the swap y_5 is moved away from its optimal value, and so the swap y_5 will be well determined (steep thin line in chart chart). But there are few government assets with prices significantly affected by this part of the curve, so the gilt error function will be relatively insensitive to movements in the gilt y_5, and the gilt y_5 will be badly fitted (flat thin line in chart).

If we link the two, so that $y_5^{(swap)} = y_5^{(govt)}$, and fit the joint y's so as to minimise the total of the errors from the two portfolios, the precise fitting of the swap y_5 will trickle through into a precise fitting of the joint y_5 (heavy line in chart). In effect, the information content is added rather than diluted, and the y_5 chosen will be the average of the y_5 from each portfolio, this average being information-weighted.

In summary, the multiple portfolio case works as follows. The z's are guessed. Each portfolio is assumed to have the same y's (full linking). With the values of the z's fixed, the n+1 y's are chosen so as to minimise E. With the y's at this E-minimum, $|||\mathbf{H}|||$ is calculated, this being the condition number of the n+1 by n+1 matrix of second derivatives of the error with respect to the y's.

An example of the linking of the y's over a set of portfolios with common z's.

Portfolio \ i, $\sim z_i$	0 Infinity	1 15yrs	2 30mths	3 6mths	4 5wks	5 4days
Swaps FRAs Depo's	0	3	6	9	11	12
Strippables (& strips)	1	4	7	10	11	12
Non-strippable gilts	2	5	8	10	11	12

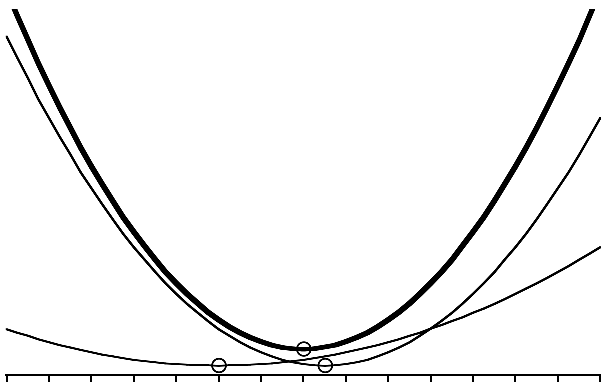

The lines shows stylised errors as functions of a single parameter. The two thin lines shows errors from two portfolios: the flatter error from a portfolio with few appropriate instruments and hence a badly fitted parameter; the steeper thin line shows the error from a portfolio with many and hence a well fitted parameter. Adding errors togther, means that fitting a joint pararameter effectively chooses the information-weighted average parameter.

Morgan Guaranty Trust Company
European Fixed Income Research
Julian Wiseman (44-171) 325 5823
wiseman_j@jpmorgan.com

The *z*'s are then repeatedly reguessed until |||**H**||| is minimised (|||**H**||| being evaulated at the *y*-minimum of *E*). With the *z*'s now known and finally fixed, the full linking of the *y*'s is replaced by partial linking (perhaps as in the example diagram at the top left of the page), and this larger set of *y*'s is optimised one last time, again to minimise the sum of the portfolio errors.

There is an extra tweak on the above that is used when fitting to non-existent instruments. This can be important: consider the case in Italy, where a strips market is due to be introduced soon. If strips were added to the portfolio only after they had prices, there would then be a jump (longer) in the *z*'s, and a commensurate adjustment in their theoretical prices and hedges. We wish to avoid such a jump on the second day of strip trading, so wish to add strips now. This is done as follows. When fitting the *z*'s by minimising |||**H**|||, we assume that strips exist and set their market prices equal to their theoretical prices. Having chosen the *z*'s, the (partially linked) *y*'s now need to be fitted: for this purpose strips are ignored.

Conclusion

The separation of *y*'s and *z*'s gives the Magpie Yield Curve Model both flexibility and a uniquene local minimum. Within an academy the commonality of the *z*'s means that the *y*'s have common meanings — allowing linking and hence the sharing of information. The result is a model that allows its flexibility to be customised to any given market or task: more bonds ⇒ more variables ⇒ fewer links; or fewer bonds ⇒ fewer variables ⇒ more links.

Appendix: Partial proof of stability

It is important that both the *y*- and *z*-optimisations should be stable with respect to small changes in the market prices of the various instruments. For simplicity, here we only prove this to be true for single-cashflow assets (typically strips, deposits and FRAs).

Recall that the zero-coupon rate from settlement *s* to maturity *t* is

$$zero_s(t) = \left(\int_s^t r(x)dx \right) / (t-s) = y_0 + y_1 \frac{z_1(e^{-s/z_1} - e^{-t/z_1})}{t-s} + \ldots + y_n \frac{z_n(e^{-s/z_n} - e^{-t/z_n})}{t-s}$$

and that the discount from maturity *t* back to settlement date *s* is

$$disc_s(t) = e^{-(t-s)zero_s(t)} = \exp\Big(y_0(s-t) + y_1 z_1 (e^{-t/z_1} - e^{-s/z_1}) + \ldots + y_n z_n (e^{-t/z_n} - e^{-s/z_n}) \Big)$$

The error is

$$E = \sum_{k \in \text{bonds}} \left(\frac{\ln\left(\frac{theory_k}{market_k}\right)}{\frac{\partial theory_k}{\partial y_0} \frac{1}{theory_k}} \right)^2$$

which is always non-negative. In the special case in which every asset has a single cashflow of unity,

$$E = \sum_{k \in \text{strips}} \left(\frac{-(t_k - s_k)\ \ zero_{s_k}(t_k)\ -\ \ln(market_k)}{-(t_k - s_k)\ \ disc_{s_k}(t_k)\ \ \frac{1}{disc_{s_k}(t_k)}} \right)^2 = \sum_k \left(zero_{s_k}(t_k)\ -\ \frac{\ln(market_k)}{(s_k - t_k)} \right)^2$$

We will need the second derivative of this *E* with respect to the *y*'s. Note that $zero_s(t)$ is linear in the *y*'s, and hence its second derivative with respect to any two *y*'s is 0.

$$\frac{\partial^2 E}{\partial y_i \partial y_j} = \frac{\partial^2}{\partial y_i \partial y_j} \sum_{k \in \text{strips}} \left(zero_{s_k}(t_k) - \frac{\ln(market_k)}{(s_k - t_k)} \right)^2$$

$$= \frac{\partial}{\partial y_i} \sum_k 2 \left(\left(zero_{s_k}(t_k) - \frac{\ln(market_k)}{(s_k - t_k)} \right) \frac{\partial zero_{s_k}(t_k)}{\partial y_j} \right)$$

$$= 2 \sum_k \left((\cdots) \frac{\partial^2 zero_{s_k}(t_k)}{\partial y_i \partial y_j} + \frac{\partial zero_{s_k}(t_k)}{\partial y_i} \frac{\partial zero_{s_k}(t_k)}{\partial y_j} \right)$$

$$= 2 \sum_k \frac{\partial zero_{s_k}(t_k)}{\partial y_i} \frac{\partial zero_{s_k}(t_k)}{\partial y_j}$$

$$= 2 \sum_k \frac{z_i (e^{-s_k / z_i} - e^{-t_k / z_i})}{t_k - s_k} \frac{z_j (e^{-s_k / z_j} - e^{-t_k / z_j})}{t_k - s_k}$$

So, for single-cashflow assets, the second derivative of the error with respect to any two *y*'s is independent of both the *y*'s and the market prices: the second derivative of the error depends solely on the *z*'s and the cashflow structure of the instruments. Thus the matrix of second derivatives is constant over *y*-space, and the error as a function of the *y*'s is quadratic.

We need to optimise the *y*'s so as to minimise the error. Equivalently, we can solve so that the first derivative of the error is zero. Since **H** is constant and non-singular, *E* must be a quadratic function of the *y*'s, and since *E* is bounded below by 0, the unique point at which the first derivative vanishes must be a minimum.

Note that **H** is far from singular, because the *z*'s were chosen to minimise |||**H**|||, and hence small changes in the prices of instruments have an effect on the *y*'s that is not only stable, but is also (in some sense) minimal. This would not be the case if we were fitting splines or polynominals: these classes of component shapes give rise to **H**'s that are near singular, and so a small price change can have a large effect on the implied curve.

So, if each asset has a single cashflow, then H is independent of the market prices, and since the *z*'s are chosen to minimise |||H|||, the *z*'s must be independent of market price. The *z*'s having been fitted to the cashflow structure of the market, the *y*-fitting then has a unique local minima, the location of which is stable to pertubations in market prices.