Free-form interest rate term structure decomposition: a 2nd order optimization problem

Brian Barnard[1]

Introduction

The paper discusses an interest rate term structure decomposition method that breaks from the conventional, in that it does not superimpose any model, form or structure on the decomposition output – hence, the term free-form. The premise is simple: if the model does not presuppose any structure beforehand, and if the structure underlying the input data is indeed that implemented and applied by conventional models[2], and if the model is not internally constrained by method itself, such that its model power can hardly be drawn into question, it should have no problem surfacing such structures. On the other hand, if the structure underlying the input data in fact does not necessarily follow a presupposed structure, the model should achieve greater modelling accuracy, and offer greater explanatory power. This in itself may also offer a better reflection on cases and scenarios where the modelling residual is significant, for instance, in the case of modelling (corporate) vanilla bonds[3].

To realize such a model, a 2nd (higher) order or nested optimization problem – a series of optimizations pertaining to the same problem, and resulting from floating one or more constraints – is defined, and solved. The model is applied to a simulated term structure, as well as a market sample of government vanilla bonds, and the results are discussed.

Methodology

Given the spot rate $\left(r^s\right)$ over the interval from the present to coupon date t, the price of an issue

1 Send correspondence to barnard.b@polka.co.za

2 Fisher et al (1994) as well as Bolder and Gusba (2002) offer excellent reviews of parsimonious models, function-based models, and spline-based models.

3 In this context, the study of Elton et al (2001) focused on explaining the rate spread on corporate bonds.

$\left(P^i\right)$ is given by the sum of its M discounted coupons. Then, given N bond issues, the coupons of a portfolio $\left(C^p\right)$ for the portfolio coupon date t, is given by the sum of all issue coupons $\left(C^i\right)$ on said coupon date t – by introducing zero coupons, an issue can be seen to have a coupon on each portfolio coupon date. Similarly, the portfolio price $\left(P^p\right)$ is given by the sum of the issues' prices $\left(P^i\right)$. A portfolio discount factor $\left(DF^p\right)$ is naturally derived from the portfolio spot rate $\left(r^{sp}\right)$ corresponding to the portfolio coupon dates. The issue price is recast in terms of the K portfolio discount factors, and portfolio coupon dates. The number of portfolio coupon dates is generally greater than or equal to the maximum of the N issues' number of issue coupon dates M.

$$P_n^i = \sum_{m=1}^{M_n} C_{nm}^i \mathrm{e}^{-r_{t_m}^s t_m} \;\; ; \;\; n \in [1,\ldots,N] \qquad [1]$$

$$C_t^p = \sum_{n=1}^{N} C_{nt}^i \;\; ; \;\; P^p = \sum_{n=1}^{N} P_n^i \qquad [2]$$

$$DF_t^p = \mathrm{e}^{-r_t^{sp} t^p} \qquad [3]$$

$$P^i = \sum_{n=1}^{M} C_n^i \mathrm{e}^{-r_{t_n}^s t_n} = \sum_{k=1}^{K} DF_k^p C_k^i \qquad [4]$$

$$K \geq max\left(M_i\right) \;\; ; \;\; i \in [1; N] \qquad [5]$$

A free-form structure essentially becomes an ordinary optimization problem[4]. The sum of the residual between issue market price $\left(P^{im}\right)$ and modelled price $\left(P^i\right)$ across all issues is minimized. An equality constraint is added: the portfolio coupons – the amalgamation of issue coupons – discounted against the resultant portfolio term structure should equal the portfolio value, given as the sum of the portfolio issues' prices. In this case, inequality constraints are also added to prevent negative forward rates across the K portfolio coupon dates.

$$minimize \;\; \sum_{n=1}^{N} \left(P_n^{im} - P_n^i\right)^2 \qquad [6.a]$$

subject to:

$$P^p = \sum_{k=1}^{K} DF_k^p C_k^p \qquad [6.b]$$

$$DF_0^p \leq 1 \qquad [6.c]$$

4 Boyd and Vandenberghe (2004) provides a sound introduction to optimization.

$$DF_k^p \geq DF_{(k+1)}^p \;\; ; \;\; k \in [1,\ldots,K-1] \qquad [6.d]$$

$$DF_K^p \geq 0 \qquad [6.e]$$

When this particular optimization problem is solved, the result may already point to a dependency between the scope and dynamics of the input data, and the resultant model power. In particular, the input data may be insufficient to adequately constrain the output, such that, in this form, high levels of variance may be measured in the output, especially with typical developing market portfolios as input. To constrain unnecessary variance in the output, another constraint is added: the output variance (V^p) may not exceed a certain ceiling (V^{pC}), defined in advance. Output variance is based on the portfolio forward rates (F^p), and is measured as the sum of inter-interval forward rate growth rates.

$$V^p = \sum_{k=1}^{K-1} \frac{\left(F_{t_k^{k+1}}^p - F_{t_{k-1}^k}^p\right)}{t_k^{k+1}} \leq V^{pC} \qquad [6.f]$$

Thus, in addition to modelling residual as primary objective, variance as secondary objective is introduced. However, variance is not directly included in the objective function – typically done by assigning a weight to it – but by modelling it as additional, albeit dynamic, constraint in the form of a predefined ceiling. Given that the minimum variance at which modelling residual is also a minimum, is not known a priori, the primary optimization problem is solved for various variance ceilings, essentially moving the ceiling about. Under-constraining the variance by an overly high ceiling typically results in a comparable – and thus insignificant change in – modelling residual when the ceiling is subsequently lowered. Conversely, over-constraining the variance by means of a low ceiling, typically results in a significant increase in modelling residual accompanying the reduction – or simply change – in the variance ceiling. The result approximates a hyperbolic plot of modelling residual against variance. Also, the distance or spread between the variance sum and variance ceiling may also be seen as an indication of whether the variance ceiling is under-constraining or over-constraining – it can be significant or insignificant, depending on the variance ceiling's imposed effect or force.

A minimum variance-sum solution was sourced to serve as initial solution for all the optimization problems to be solved as part of the overall problem. This solution ignores issue price modelling residuals altogether, and instead minimizes the variance sum alone. The optimization problem itself was also adapted to base some of the modelling residual calculations on absolute sum, not squared sum – it is found that this exposes greater accuracy to some degree. An algorithm was written to search a knee point where the residual sum starts to rapidly degrade and, the variance-sum-ceiling distance starts to disappear. In this case a fixed knee point was used – the point at which the variance sum-ceiling distance as proportion of the variance ceiling dropped below 1E-03. The model was implemented on a graphics processing unit (gpu), given that it resembles a series of optimization problems to solve, thereby posing some compute intensity. Also because it allows substantial parallelization at both the level of individual optimization problems, and the higher level of successive optimization problems – both solving the optimization problems, and optimization problem succession can be parallelized.

To test the model, a market term structure was simulated by selecting a number of intervals, randomly selecting the forward rates of the interval points between a set minimum (0) and maximum (0.3), and interpolating the forward rates of the remaining points by means of straight lines across the intervals. All government zero coupon and vanilla bonds from the South African market for a given date (2014-07-09) were taken, and discounted by the simulated term structure, to obtain simulated values for these issues, using their original coupon data. The issues' simulated prices and original coupon data were used to decompose a term structure, and the result was compared to the simulated structure. The model was also run on the original price data of the said government bond sample. The given sample constitutes 27 issues, with 211 portfolio coupon dates, resulting in 211 portfolio discount factors.

Analysis

Figure 1 to 5, and table 1 pertain to the simulated case; figure 6 to 10, and table 2 pertain to the market case. In all of the below, modelling residual is stated as absolute sum, not squared sum. All modelling residual values are thus essentially understated in terms of squared sum.

Figure 1 and 2 show the simulated and modelled spot rate and forward rate term structures. Table 1 shows the various variance ceilings selected by the algorithm, with their resultant modelling residuals and variance sums. A lot can perhaps be said about the simulated term structure. The first question may be whether it is indeed practical – a subsequent paragraph returns to this point. Whether it is practical or not, it may point out some of the intricacies embedded in pre-form models. For instance, even though the models were not run, a spline-based model would arguably have achieved greater accuracy over the Nelson Siegel model in this case, and the Nelson Siegel model may have overly averaged the case. With the spline-based model, it is clear that the number of intervals selected, may play a significant role in the outcome; the number of relevant intervals may also vary across decompositions.

The correlation or degree of fit between the simulated and modelled term structure was not quantified. However, modelling residual measures so low, such that it can be viewed as insignificant or essentially zero, and modelling residual fails to be a differentiating factor any longer. From the perspective of the model, the solution is seen as optimal, and any discrepancy between simulated and modelled term structure becomes an attribute of input data, and model power. More input data could have sustained a differentiating modelling residual for longer, or more uniquely, giving the model more to respond to. Also, the rather arbitrary and absolute selection of a knee point to terminate on, may be seen as an artifact of model power. Figure 3 shows the issues contributing coupons to the portfolio coupons and thus term structure coefficients. It essentially reflects on data density per portfolio term structure coefficient. Evidently, some areas have greater density than others. The data density plot can of course be mapped back to the decomposed term structure plots.

Figure 4 and 5 show the approximate hyperbolic plot of modelling residual against variance sum. It also shows how the distance between the variance sum and variance ceiling shrinks as the variance ceiling is lowered and consequently becomes a greater constraining factor. At some point, the variance ceiling becomes a saturating constraint, and thereafter, the modelling residual starts to rapidly degrade as a consequence.

With the market case, modelling residual can not be discarded as insignificant, as in the previous case, even though it is still relatively low, with very high correlation between modelled and market prices: as a percentage of total portfolio value, modelling residual is less than 1%. As figure 8 and 9 show, variance sum against modelling residual, and the distance between the variance sum and variance ceiling follow the previously described pattern.

In placing the modelling residual and modelled term structure in context, if it is argued that the model attained relatively high accuracy in the previous case, the question is whether this particular modelled term structure is an accurate representation. Has the model somehow lost its power this time, or not?

It is pointed out that the constraints of the model – preventing negative forward rates, and preserving the value of the portfolio – may preliminarily be put forward as potential contributor – removing these constraints may provide greater flexibility, aiding the model. But, this brings into question whether the constraints are indeed sensible and practical. Except for price, the underlying issue data is essentially the same. Hence, any potential model power degradation may then perhaps rather point to a discrepancy between embedded variance in the simulated case, and the market case – with that of the market case being greater – such that a relative decline in information provided, occurred. In addition, it may simply be that the model is attempting to adjust to individual issues, and their impact on the portfolio. In the simulated case, the impact of individual issues was essentially limited – actually completely removed – by discounting all issues with the same portfolio term structure. In a market context, issues may not necessarily share a common term structure, such that issue-based modelling residual measured by means of a portfolio term structure is bound to be non-zero.

To further reflect on this, the modelled market term structure was again used as a simulated case – the issues' coupons were discounted against the modelled market term structure, and their prices recalculated accordingly. A decomposition was again run on the data, to evaluate to what extent the model could rediscover the solution. Figure 10 and 11, and table 3 show the results.

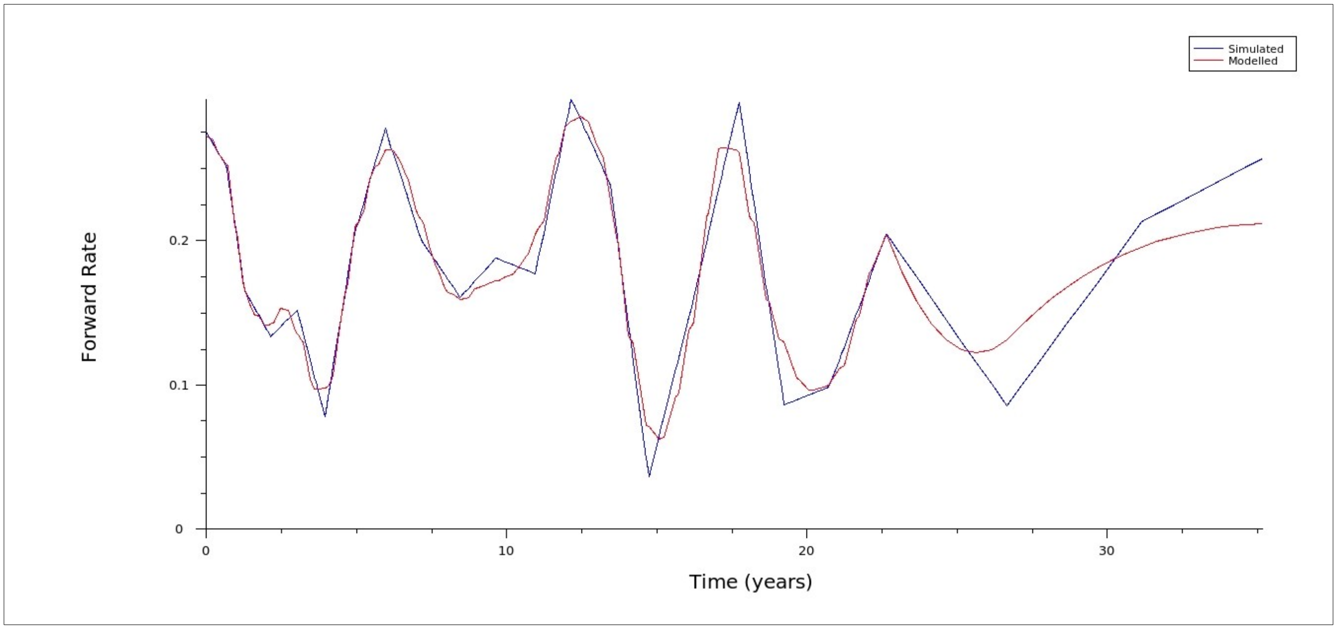

Figure 1: Simulated and modelled forward rate plot

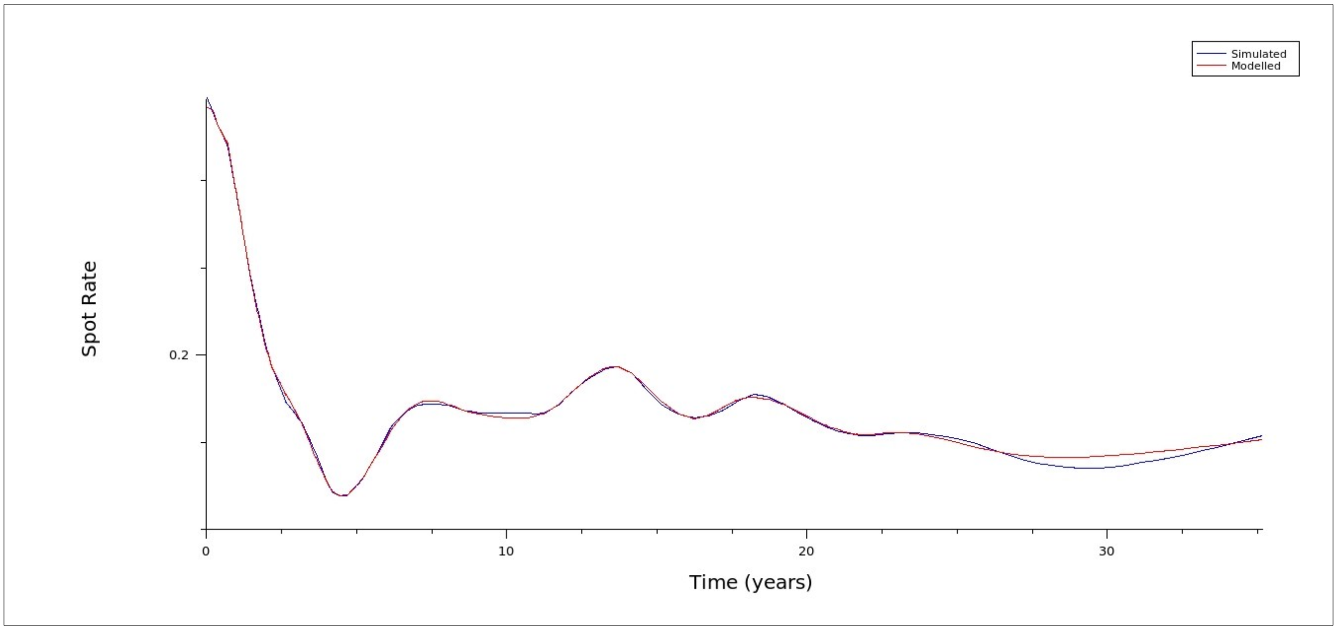

Figure 2: Simulated and modelled spot rate plot

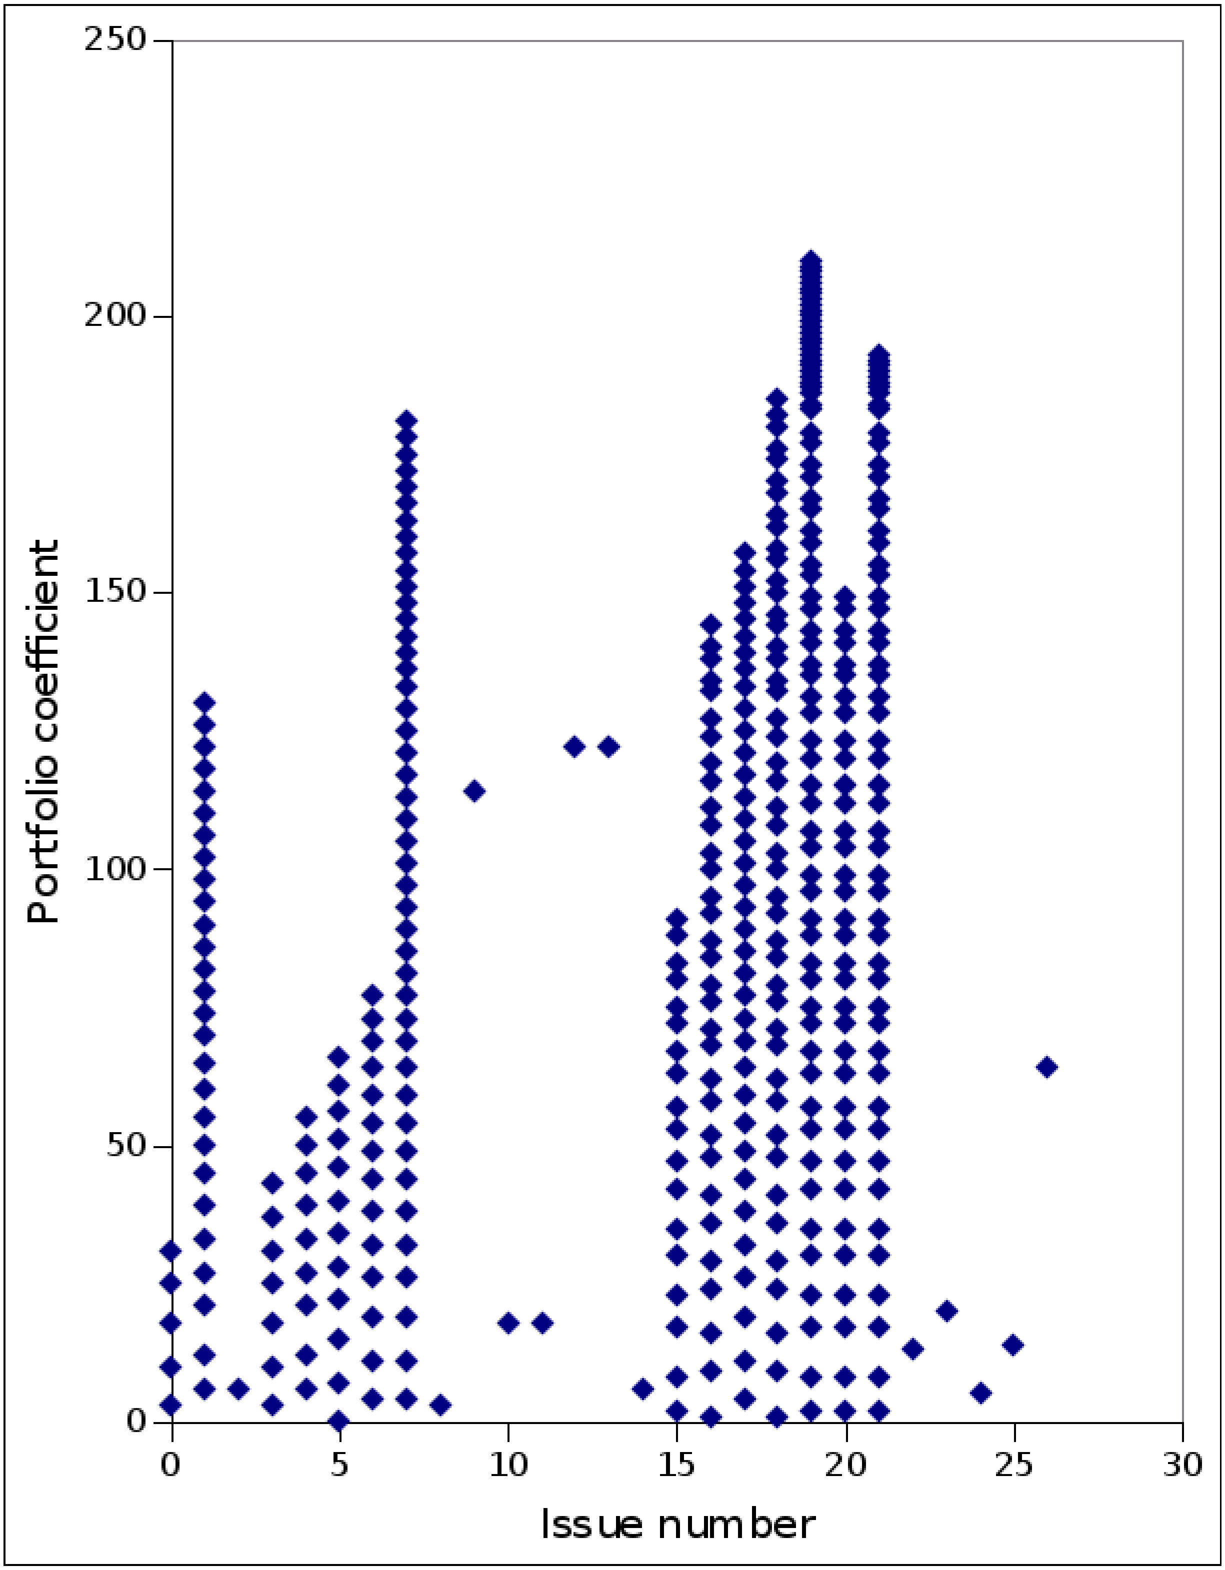

Figure 3: Issues per portfolio coupon date

Table 1: Residual and variance sums for variance ceilings selected – simulated case

	Modelling Residual	Variance Sum	Variance Ceiling
1	7.2677507164e-05	1.2718862337e+01	1.5266840874e+01
2	6.0175864292e-11	6.2324410009e+00	7.6334204369e+00
3	8.0589046547e-10	3.0506717597e+00	3.8167102184e+00
4	5.3500315289e-11	1.5260966662e+00	1.9083551092e+00
5	1.4944490090e-11	8.2876689407e-01	9.5417755461e-01
6	5.4321786510e-01	4.7708662669e-01	4.7708877730e-01
7	1.3708145730e-10	7.4582186419e-01	8.2876689407e-01
8	3.7075054138e-10	6.9275065159e-01	7.4582186419e-01
9	2.1037749320e-10	6.5971128322e-01	6.9275065159e-01
10	1.1990408666e-11	6.3976015629e-01	6.5971128322e-01
11	1.9557688802e-10	6.2828686451e-01	6.3976015629e-01
12	1.6583889817e-10	6.2227415332e-01	6.2828686451e-01
13	5.5868376592e-10	6.1967740580e-01	6.2227415332e-01
14	1.0557954511e-10	6.1895811593e-01	6.1967740580e-01
15	3.1402791478e-10	6.1888819114e-01	6.1895811593e-01

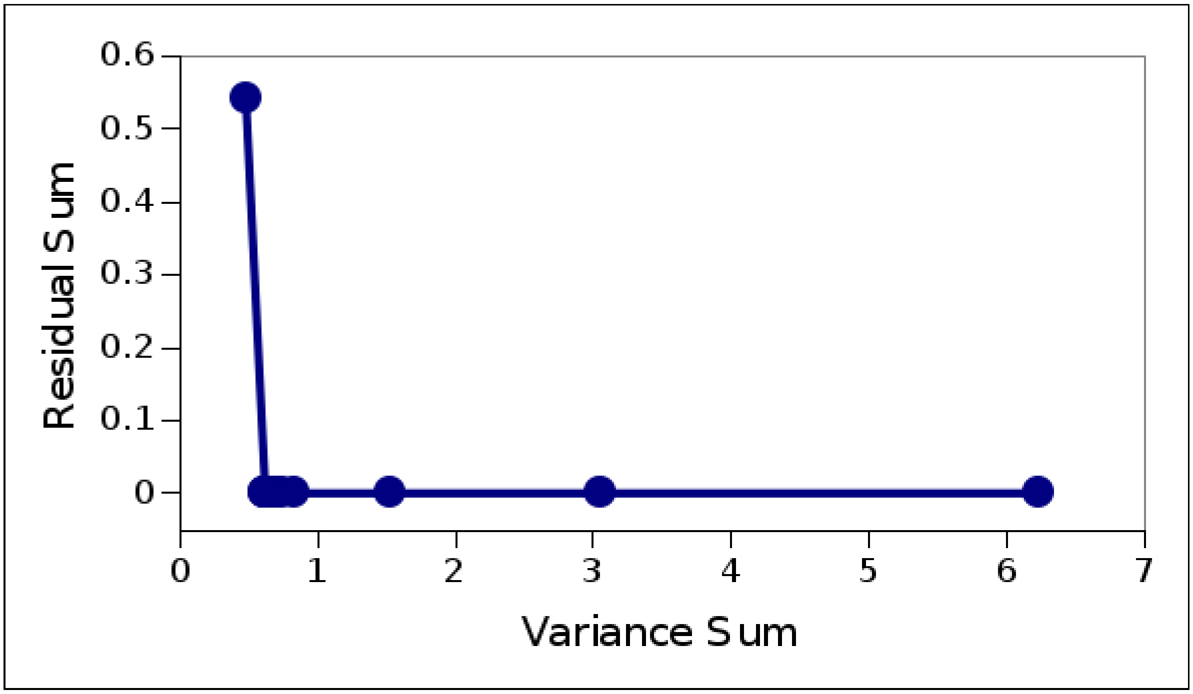

Figure 4: Residual sum, variance sum plot

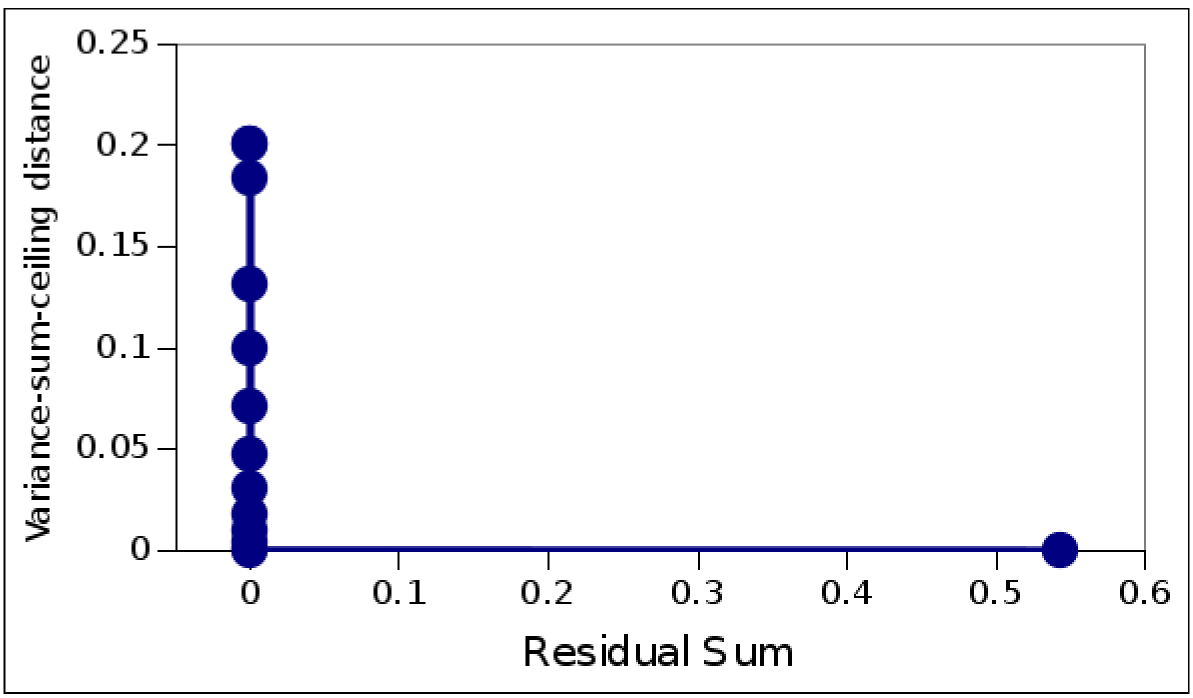

Figure 5: Variance-sum-ceiling distance, residual sum plot

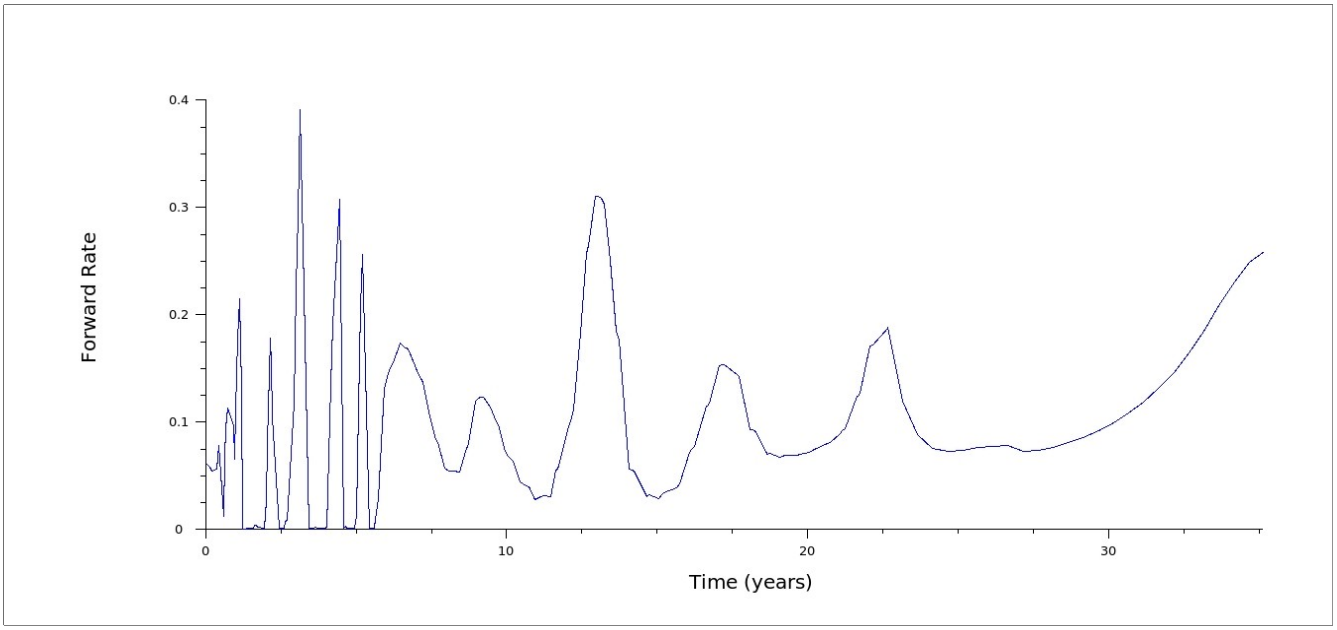

Figure 6: Modelled forward rate plot

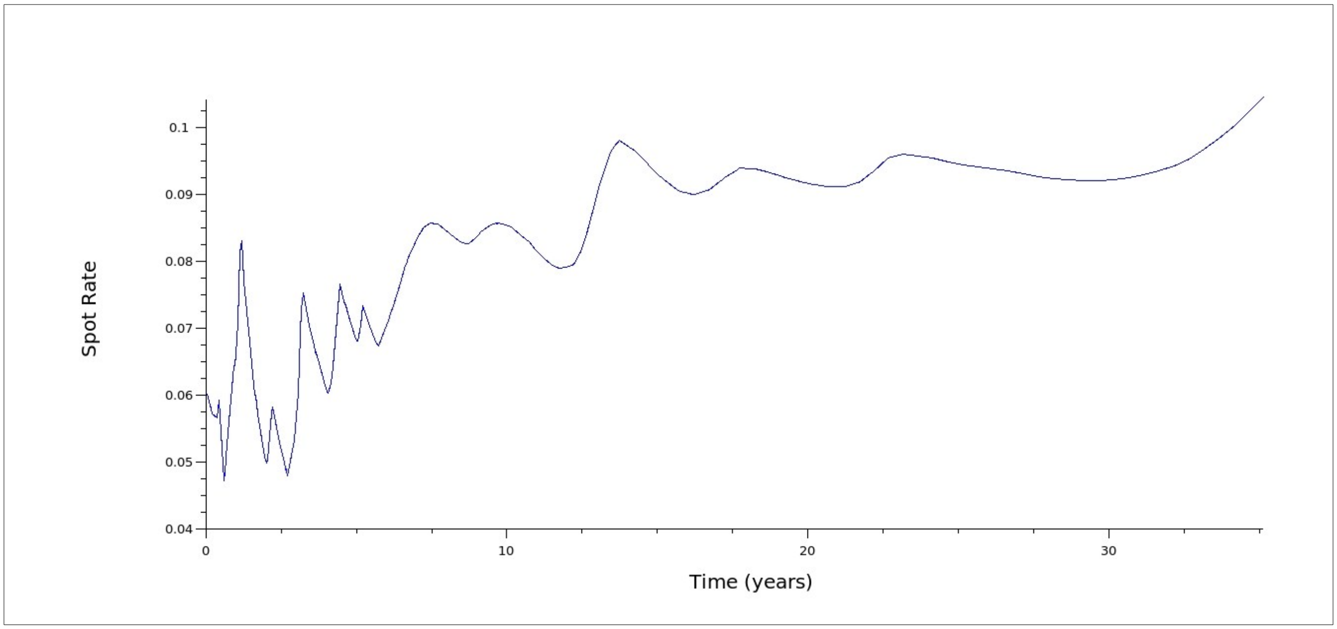

Figure 7: Modelled spot rate plot

Table 2: Residual and variance sums for variance ceilings selected – market case

	Modelling Residual	Variance Sum	Variance Ceiling
1	1.8678669300e+03	6.2146724682e-01	6.5256297673e-01
2	1.2902325735e+01	9.7884446509e-01	9.7884446509e-01
3	1.2316512434e+01	1.4682666976e+00	1.4682666976e+00
4	1.1939423126e+01	2.2024000465e+00	2.2024000465e+00
5	1.1640731330e+01	3.3036000697e+00	3.3036000697e+00
6	1.1296276685e+01	4.9554001045e+00	4.9554001045e+00
7	1.0942112107e+01	7.4331001568e+00	7.4331001568e+00
8	1.0491611042e+01	1.1149650235e+01	1.1149650235e+01
9	9.9789767987e+00	1.6724475351e+01	1.6724475353e+01
10	9.8237139685e+00	2.5086713029e+01	2.5086713029e+01
11	9.7383539495e+00	3.7630069544e+01	3.7630069544e+01
12	9.6914500012e+00	5.6369437344e+01	5.6445104316e+01
13	9.6914500003e+00	8.3363413596e+01	8.4667656473e+01
14	9.7066114751e+00	4.7037586929e+01	4.7037586930e+01
15	9.6961320988e+00	5.1741345619e+01	5.1741345623e+01
16	9.6916087858e+00	5.4093224952e+01	5.4093224969e+01
17	9.6914500024e+00	5.5234110226e+01	5.5269164642e+01
18	9.6914500016e+00	5.5802035936e+01	5.5857134479e+01
19	9.6914500013e+00	5.6085792887e+01	5.6151119397e+01
20	9.6914500014e+00	5.5943929357e+01	5.6004126938e+01
21	9.6914500016e+00	5.5872986512e+01	5.5930630709e+01
22	9.6914500015e+00	5.5837512209e+01	5.5893882594e+01
23	9.6914500016e+00	5.5783091968e+01	5.5837512209e+01
24	9.6914500016e+00	5.5819774321e+01	5.5875508536e+01

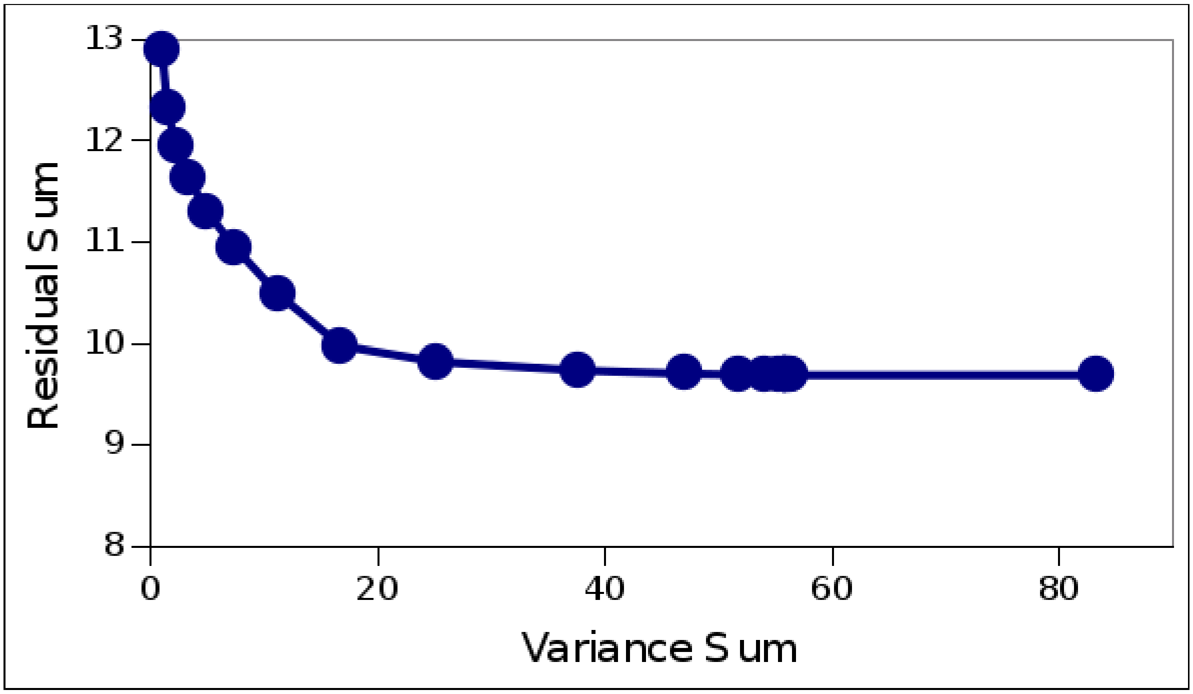

Figure 8: Residual sum, variance sum plot

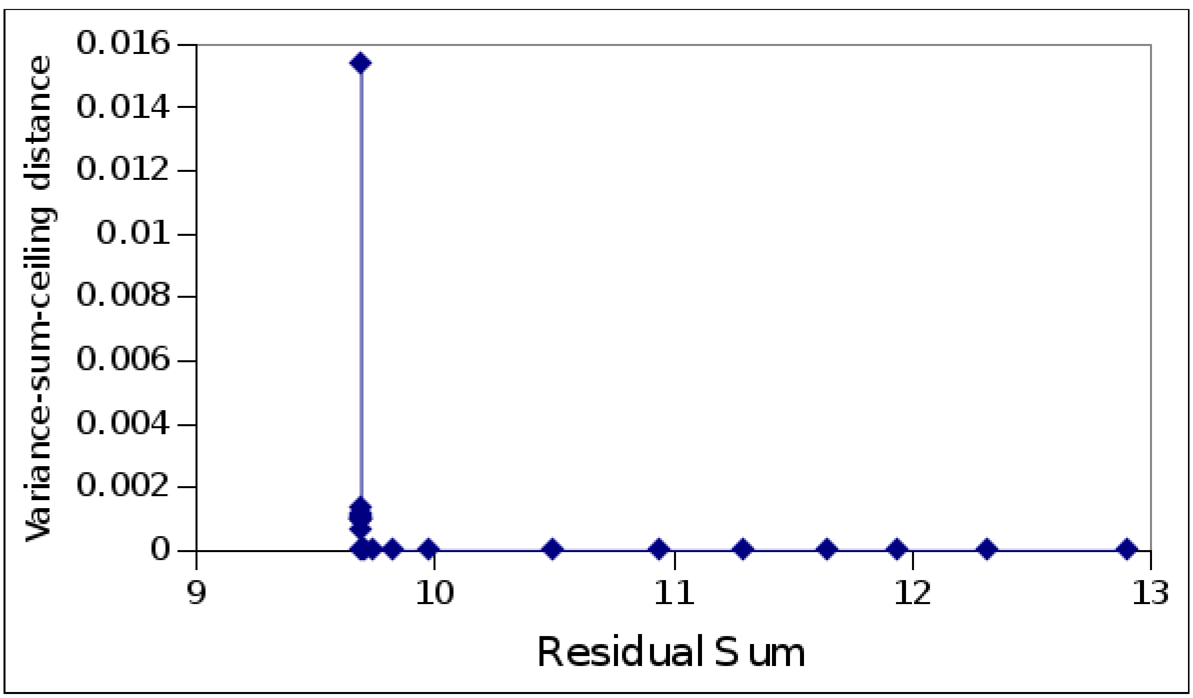

Figure 9: Variance-sum-ceiling distance, residual sum plot

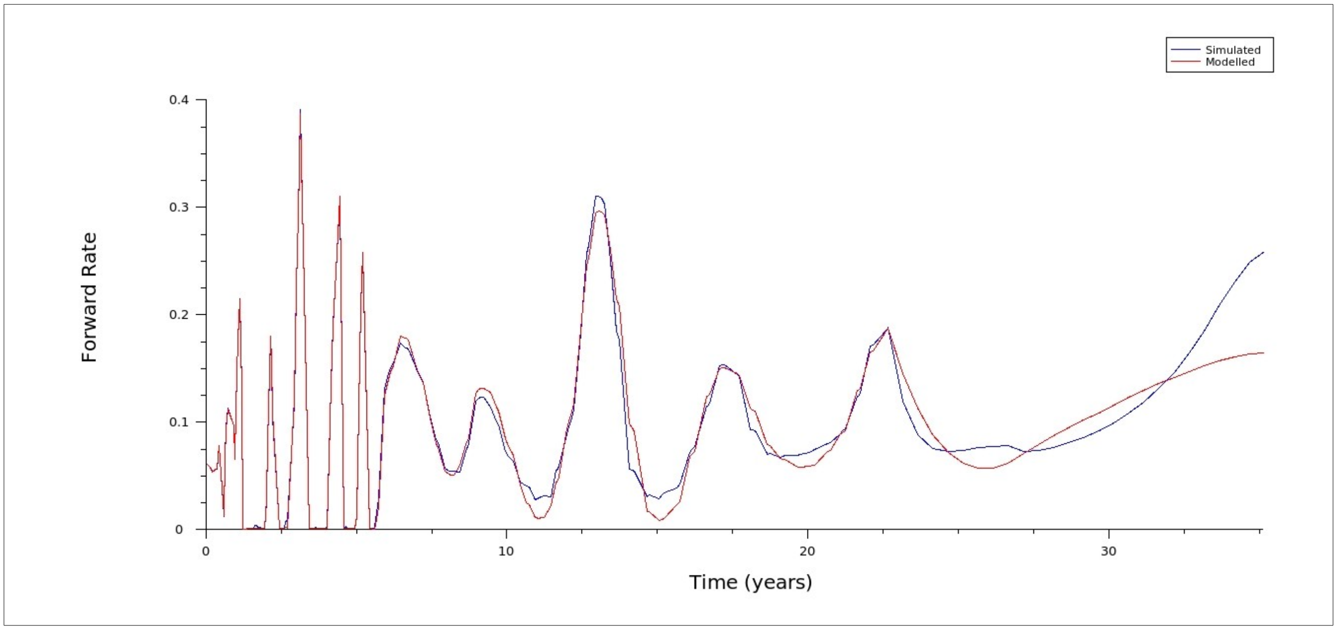

Figure 10: Simulated and modelled forward rate plot

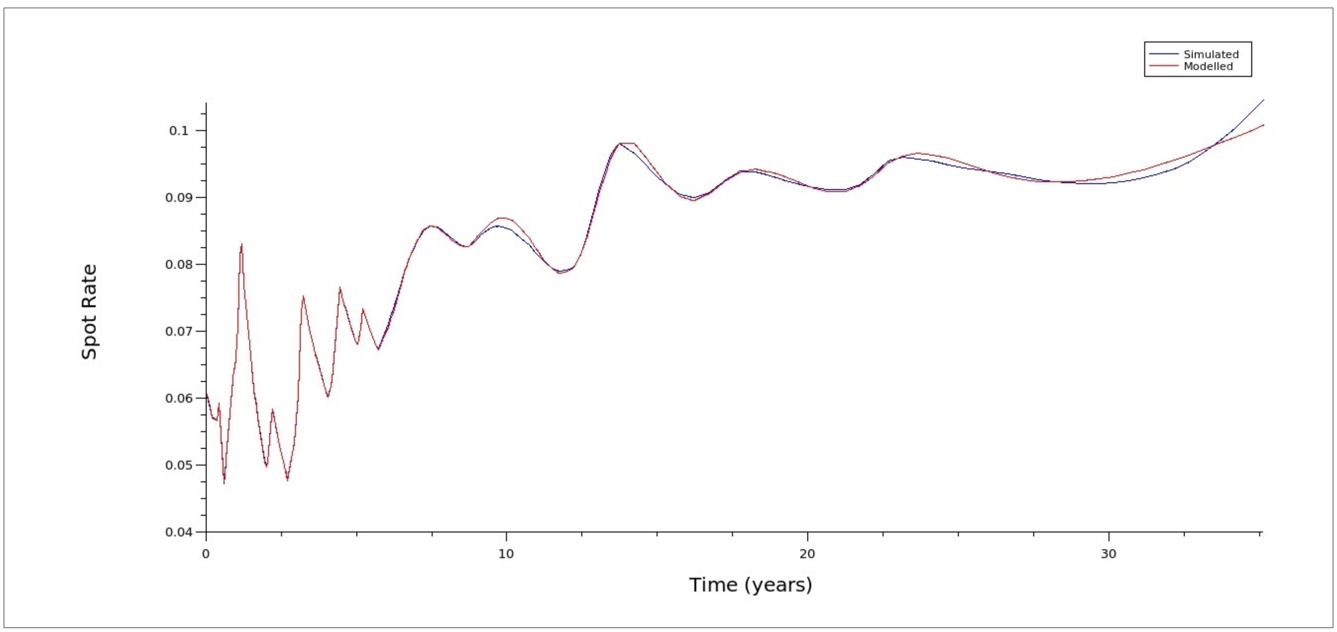

Figure 11: Simulated and modelled spot rate plot

Table 3: Residual and variance sums for variance ceilings selected – recast case

	Modelling Residual	Variance Sum	Variance Ceiling
1	2.5042901714e+00	2.1167351031e+00	2.1167351031e+00
2	2.0241388746e+00	3.1751026546e+00	3.1751026546e+00
3	1.5706415744e+00	4.7626539819e+00	4.7626539819e+00
4	1.1030393782e+00	7.1439809728e+00	7.1439809728e+00
5	6.4648228226e-01	1.0715971459e+01	1.0715971459e+01
6	3.4576928538e-01	1.6073957189e+01	1.6073957189e+01
7	1.6281978839e-01	2.4110935783e+01	2.4110935783e+01
8	5.5889751190e-02	3.6166403675e+01	3.6166403675e+01
9	6.1772453819e-10	5.4247366500e+01	5.4249605512e+01
10	1.8344650243e-09	8.0233333569e+01	8.1374408268e+01

Conclusion

Higher order optimization as method, particularly as applied to the context of term structure decomposition seems promising, with potential. Free-form term structure decomposition offers a new perspective on decomposition. One may argue that it moves away from and against over-emphasizing form in term structure decomposition, to rather examine more closely the forces of individual issues on portfolios. It certainly raises a number of questions, and points out certain premises embedded in pre-form term structure decomposition, and term structure decomposition in general. The relative validity of the constraints implemented by the model could be further examined. Also, it is questioned whether an over-insistence on form itself may not lead to model-induced modelling residual, contrasted to issue-induced modelling residual.

The possible impact of sample size on model accuracy was alluded to. The study offers some basis for a framework to investigate the said dynamic in greater detail. Very basically, an artificial sample can be expanded and retracted to further examine the phenomenon. Equally, simulating term structures in itself offers some method to reflect on model power. The rudimentary simulations of the study can be expanded, not only in terms of form, but also in terms of hierarchy – contemplating simulations at the issue level.

If persistent issue-identical term structures – issues sharing the same term structure with other issues and the portfolio – as concept is questionable, perhaps the contrary should be emphasized and investigated. This may entail revising the view that the portfolio term structure is simply a regression and thus broad averaging of individual issues, to more specifically define the relationship between the portfolio term structure and issue term structures, and it may imply decomposing at the issue level. With higher order optimization as method, this may indeed be tenable.

References

Boyd, S. and Vandenberghe, L. (2004) Convex Optimization, Cambridge University Press

Bolder, D. J. and Gusba, S. (2002) Exponentials, Polynomials, and Fourier Series: More Yield Curve Modelling at the Bank of Canada, Bank of Canada Working Paper

Elton, E. J., Gruber, M. J., Agrawal, D. and Mann, C. (2001) Explaining the Rate Spread on Corporate Bonds, The Journal of Finance, 56(1)

Fisher, M., Nychka, D. and Zervos, D. (1994) Fitting the Term Structure of Interest Rates with Smoothing Splines, U.S. Federal Reserve Board Working Paper