

A Heuristic for Generating Scenario Trees for Multistage Decision Problems

Kjetil Høyland * Michal Kaut † Stein W. Wallace ‡

June 20, 2000; revised May 11, 2001, November 20, 2001

Abstract

In stochastic programming models we always face the problem of how to represent the random variables. This is particularly difficult with multidimensional distributions. We present an algorithm for efficient generation of scenario trees for single or multi-stage problems.

The presented algorithm produces a discrete distribution specified by the first four marginal moments and correlations. The scenario tree is constructed by decomposing the multivariate problem into univariate ones, and using an iterative procedure that combines simulation, Cholesky decomposition and various transformations to achieve the correct correlations without changing the marginal moments.

Our testing shows that the algorithm is faster than a benchmark algorithm with no decomposition. The speed-up increases with the size of the tree and is more than 100 in a case of 20 random variables and 1000 scenarios. This allows us to increase the number of random variables. With the benchmark algorithm, the maximal number of random variables that could be generated in a ‘reasonable’ time was about 12, while the new algorithm allows for a far larger number.

Keywords: stochastic programming, scenario tree generation, asset allocation, Cholesky decomposition, heuristics

1 Introduction

Gjensidige Nor Asset Management (GNAM) has NOK 65 billion (7 billion US\$) under management. During the last few years they have used stochastic-programming-based asset allocation models in their asset allocation processes. The most important step in that process is to establish the market expectations, i.e. to establish what are their beliefs for the major asset categories (bonds, stocks, cash, commodities and currencies) in different major regions of the world. The decision-makers prefer to express their expectations in terms of marginal distributions of the return / interest rates for the different asset classes in addition to correlations. The challenge in converting these expectations into an asset allocation mix is twofold: First we need to convert the expectations into a format which a stochastic programming model can handle, second we need an optimization model which gives us the optimal asset mix, given this input.

Practical experience has told us that the first part, the scenario generation, can in fact be the most challenging and (computer-) time consuming one. For the purpose of generating the input data, GNAM has been using the method described in *Høyland and Wallace, 2001*, a method developed in 1996. For larger problems, with many asset classes, the scenario generation became the bottleneck in the portfolio optimization process. This paper presents an algorithm that reduces the computing time for the scenario generation substantially.

The most well-known applications of asset allocation models are the Russell Yasuda Kasai model in *Cariño and Ziemba, 1998* and the models implemented by Towers Perrin in *Mulvey, 1996*. Other applications can be found in *Consigli and Dempster, 1998* and *Dert, 1995*.

*Gjensidige Nor Asset Management, POBox 276, N-1326 Lysaker, Norway

†Dept. of Industrial Economics and Technology Management, NTNU, N-7491 Trondheim, Norway

‡Molde University College, Servicebox 8, N-6405 Molde, Norway

These models are all focused on long term strategic asset liability planning. Stochastic processes are widely used for scenario generation in such models. The big challenge with such processes is to calibrate the parameters so that the generated scenarios are consistent with the decision-maker's beliefs of the future development of the asset classes. In many applications the parameters are calibrated so that the future scenarios are consistent with the past. This might be acceptable for long term strategic planning. For tactical short term planning, i.e. for the question of how to construct an asset allocation mix relative to an asset allocation benchmark, such approaches are, however, inappropriate. The user wishes to express views on the future which deviate from the past. It is then important that the decision-maker can express the market expectations in a way that he or she finds convenient and that these expectations are converted to model input in a consistent manner.

This is the basic philosophy of the scenario generation method proposed in *Høyland and Wallace, 2001*. The user specifies his or her market expectations in terms of marginal distributions for each asset class in addition to correlations among the different asset classes and possibly other statistical properties. The stochastic, possibly multi-period, asset allocation model requires discrete outcomes for the uncertain variables. To generate these outcomes a least squares model is applied. The idea is to minimize the distance between the statistical properties of the generated outcomes and the specified statistical properties (specified either directly or derived from the marginal distributions, which may themselves be calculated from data or specified explicitly), see *Høyland and Wallace, 2001* for more details.

In the general form of the algorithm presented in *Høyland and Wallace, 2001*, all the scenarios (returns for all the asset classes) are generated simultaneously. Such an approach becomes slow when the number of random variables increases. In this paper we present an algorithm that decomposes the scenario generation problem. We generate one marginal distribution at a time and create the multivariate distribution by putting the univariate outcomes together in the following way: All marginal distributions are generated with the same number of realizations, and the probability of the i th realization is the same for each marginal distribution. The i th scenario, that is, the i th realization of the multivariate distribution, is then created by using the i th realization from each univariate distribution, and given the corresponding probability.

the following paragraph is really deleted

We create a scenario, that is, a realization of the multivariate distribution, by randomly (without replacement) picking one realization from each univariate distribution. All scenarios will end up with the same probabilities, and there will be as many scenarios as there are realizations in each marginal distribution.

We then apply various transformations in an iterative loop to reach the target moments and correlations.

The presented algorithm is inspired by the work of *Fleishman, 1978*, *Vale, 1983* and *Lurie and Goldberg, 1998*. Fleishman presents a cubic transformation that transforms a sample from a univariate normal distribution to a distribution satisfying some specified values for the first four moments. Vale and Maurelli address the multivariate case and analyse how the correlations are changed when the cubic transformation is applied. The algorithm assumes that we start out with multivariate normals. The initial correlations are adjusted so that the correlations after the cubic transformation are the desired ones. The algorithm is only approximate with no guarantee about the level of the error.

Lurie and Goldberg outlined an algorithm that is of similar type as ours. They also generate marginals independently and transform them in an iterative procedure. There are, however, two major differences between the two algorithms. One is in the way they handle the (inevitable) change of distribution during the transition to the multivariate distribution – while they modify the correlation matrix in order to end up with the right distribution, we modify the starting moments. The other major difference is that they start out with parametric univariate distributions whereas we start out with the marginal moments. We believe that specifying marginal moments is a more flexible approach and we certainly could also derive the marginal moments (up to the desired number) from the parametric distribution and apply our approach.

In Section 2 we present the algorithm. To simplify the presentation the algorithm is illustrated for the single period case. Høyland and Wallace argue that the quality¹ of the scenario tree improves if several small trees are aggregated to a larger tree rather than generating the larger

¹See *Høyland and Wallace, 2001* for how to measure ‘quality’ of the scenario tree.

tree directly. We present the algorithm for a single period with no aggregation. In Section 3 we discuss how the algorithm can be adjusted to incorporate tree aggregation and multiple periods. Numerical results are presented in Section 4, while possible future research areas are discussed in Section 5.

2 The algorithm

In Høyland and Wallace’s method a scenario tree can in principle be constructed to fit any statistical properties. In this section we will assume that the specified properties are the first four marginal moments and the correlations. This is consistent with many other studies, as well as our own empirical analysis in *Høyland and Wallace, 2001* of what were the important statistical properties in the given case study. The presented methodology is more general than this, we could in principle specify even higher moments, but it is more restrictive than our original approach, where a wide range of statistical properties could be specified.

The general idea of the algorithm is as follows: Generate univariate outcomes for n random variables, each satisfying a specification for the first four moments. Transform these outcomes so that the new outcomes are consistent with a given correlation matrix. The transformation will distort the higher than the second moments of the original univariate outcomes. Hence we need to start out with a different set of higher moments so that the higher moments that we end up with are the right ones.

The procedure would lead to the exact desired values for the correlations and the marginal moments if the generated univariate outcomes were independent. This is, however, true only when the number of scenarios goes to infinity. With a limited number of scenarios, the marginal moments and the correlations will therefore not fully match the specifications. To be able to secure that the error is within a pre-specified range we have developed an iterative algorithm, which is an extension of the core algorithm.

Section 2.1 discusses the assumptions we have on the correlation matrix, Section 2.2 introduces necessary notation, Section 2.3 explains the key transformations used in the algorithm, Section 2.4 describes the core module of the algorithm, while Section 2.5 explains the iterative procedure.

2.1 Assumption on the correlation matrix

There are two assumption on the specified correlation matrix R . The first one is a general assumption that R is a possible correlation matrix, i.e. that it is a symmetric positive semi-definite matrix with 1’s on the main diagonal. While implementing the algorithm there is no need to write a special code that checks positive semi-definiteness, because we do a Cholesky decomposition of the matrix R at the very start. If R is not positive semi-definite, the Cholesky decomposition will fail.

Note that having an R that is not positive semi-definite means having some internal inconsistency in the data, so we should re-run our analysis. As an alternative, there exist several algorithms that find a possible correlation matrix that is, in some sense, closest to the specified matrix R . One such an algorithm can be found in *Higham, 2000*. Another approach is used in *Lurie and Goldberg, 1998*, where a QP model is formulated to solve the problem. The latter approach has an advantage of being very flexible and allowing, for example, for specifying weights that express how much we believe in every single correlation. We can also use bounds to restrict the possible values. On the other hand, it has the obvious disadvantage of needing a QP solver.

The other assumption is that all the random variables are linearly independent, so that R is a non-singular – hence positive definite – matrix. For checking this property we can again use the Cholesky decomposition, because the resulting lower-triangular matrix L will have zero(s) on its diagonal in a case of linear dependency.

This is not a serious restriction, since having a linearly dependent variable mean that it can be computed from the others after the generation, and we can thus decrease the size of the problem. If we are not sure about our correlations, we can use the method from *Lurie and Goldberg, 1998* mentioned in the previous paragraph, to enforce the positive definiteness. Since they have the elements of the Cholesky-decomposition-matrix L_{ij} as variables in the QP model, we can simply add a constraint $\{L_{ii} \geq \varepsilon > 0\}$ for every random variable i .

We also suggest a simple heuristic to handle the singularity of the correlation matrix inside our algorithm, see Section 2.5.1. It is, however, a very crude procedure, and we recommend to avoid singular correlation matrices if possible by going back to the analysis of the data.

2.2 Notation

To formulate the model we introduce the following notation. Note that vectors are columns by default.

n	number of random variables
s	number of scenarios
\tilde{X}	n -dimensional random variable $\rightarrow \tilde{X} = (\tilde{X}_1, \tilde{X}_2, \dots, \tilde{X}_n)$ \rightarrow every moment of \tilde{X} is a vector of size n \rightarrow the correlation matrix of \tilde{X} is a matrix of size $n \times n$
X	one realization of the random vector \tilde{X} – vector of size n
\mathbb{X}	matrix of s scenarios of the random vector \tilde{X} – \mathbb{X} has dimension $n \times s$
\mathbb{X}_i	row vector of outcomes for random variable \tilde{X}_i – \mathbb{X}_i has size s
$\tilde{X}(mom; corr)$	\tilde{X} has moments $mom = mom_1 \dots mom_4$ and a correlation matrix $corr$ \rightarrow every mom_i is a vector of size n , and $corr$ is a matrix of size $n \times n$

Our goal is to generate scenarios \mathbb{Z} for an n -dimensional random variable \tilde{Z} with properties $\tilde{Z}(TARMOM; R)$, i.e. $TARMOM$ is the matrix (of size $4 \times n$) of target moments and R is the target correlation matrix (of size $n \times n$).

2.3 Key transformations

The core module, which will be presented in the next section, has two key transformations. One is a cubic transformation used for generating univariate distributions with specified moments. The other is a matrix transformation used to transform the multivariate distribution to obtain a given correlation matrix.

2.3.1 Cubic transformation

This transformation comes from *Fleishman, 1978*, where a method to generate a univariate non-normal random variable \tilde{Y}_i with given first four moments is introduced.² It takes a $\mathcal{N}(0, 1)$ variable \tilde{X}_i and uses a cubic transformation

$$\tilde{Y}_i = a + b\tilde{X}_i + c\tilde{X}_i^2 + d\tilde{X}_i^3$$

to obtain \tilde{Y}_i with the target moments. Parameters a , b , c and d are found by solving a system of non-linear equations. Those equations utilize normality of the variable \tilde{X}_i .

The problem of this approach is that \tilde{X}_i must have the first 12 moments equal to those of $\mathcal{N}(0, 1)$ in order to get exactly the target moments of \tilde{Y}_i . Since this is hard to achieve, either by sampling or discretization, the results with those formulas are only approximate. We have thus dropped the assumption of normality and derived formulas that work with arbitrary random variable \tilde{X}_i – see Appendix C. Parameters a , b , c and d are now given by a system of four implicit equations.

We have used a simple non-linear mathematical-programming (least-squares) model to solve the system. In this model we have a , b , c and d as independent variables, moments of \tilde{Y}_i as dependent variables, and we minimize a distance of the moments from their target values. Our method for generating a sample \mathbb{Y}_i from \tilde{Y}_i is then as follows:

- take some sample \mathbb{X}_i of the same size as \mathbb{Y}_i
- calculate the first 12 moments of \mathbb{X}_i (as a discrete distribution)
- compute the parameters a , b , c , d
- construct \mathbb{Y}_i as $\mathbb{Y}_i = a + b\mathbb{X}_i + c\mathbb{X}_i^2 + d\mathbb{X}_i^3$

²The index i is obsolete in this section, but we use it for consistence with the rest of the paper.

2.3.2 Matrix transformation

Our other main tool in the algorithm is a matrix transformation

$$\tilde{Y} = L \times \tilde{X}$$

where L is the lower-triangular matrix. The matrix L always comes from a Cholesky decomposition of the correlation matrix R , so we have $L \times L^T = R$.³

From statistical theory we know that if \tilde{X} is an n -dimensional $\mathcal{N}(0, 1)$ random variable with correlation matrix I (and therefore with \tilde{X}_i mutually independent), then the $\tilde{Y} = L \times \tilde{X}$ is an n -dimensional $\mathcal{N}(0, 1)$ random variable with correlation matrix $R = L \times L^T$.

Since we do not have normal variables, we need a more general result. To make the formulas as simple as possible, we restrict ourselves to the case of zero means and variances equal to 1. In the beginning of Section 2.4 we show how to deal with this restriction. Note that $E[\tilde{X}] = 0$ leads to $mom_i = E[\tilde{X}^i]$. We will thus use the two interchangeably for the rest of the section.

In Appendix B we show that if each \tilde{X}_i has zero mean and variance equal to 1, that the \tilde{X}_i are mutually independent, and that $R = L \times L^T$ is a correlation matrix, then $\tilde{Y} = L \times \tilde{X}$ is an n -dimensional random variable with zero means, variances equal to 1, and correlation matrix (which is also the variance-covariance matrix) R . The higher moments can be found as follows:

$$\begin{aligned} \mathbb{E}[\mathbb{Y}_i^3] &= \sum_{j=1}^i L_{ij}^3 \mathbb{E}[\mathbb{X}_j^3] \\ \mathbb{E}[\mathbb{Y}_i^4] - 3 &= \sum_{j=1}^i L_{ij}^4 (\mathbb{E}[\mathbb{X}_j^4] - 3) \end{aligned}$$

We will need also the opposite direction of the transformation:

$$\tilde{X} = L^{-1} \times \tilde{Y}.$$

Since the above formulas are triangular, it is easy to invert them:

$$\begin{aligned} \mathbb{E}[\mathbb{X}_i^3] &= \frac{1}{L_{ii}^3} \left(\mathbb{E}[\mathbb{Y}_i^3] - \sum_{j=1}^{i-1} L_{ij}^3 \mathbb{E}[\mathbb{X}_j^3] \right) \\ \mathbb{E}[\mathbb{X}_i^4] - 3 &= \frac{1}{L_{ii}^4} \left[\mathbb{E}[\mathbb{Y}_i^4] - 3 - \sum_{j=1}^{i-1} L_{ij}^4 (\mathbb{E}[\mathbb{X}_j^4] - 3) \right] \end{aligned}$$

We divide only by the diagonal elements L_{ii} in these formulas. We can do it since L_{ii} are positive due to regularity (positive definiteness) of R . Note however that if L_{ii} are close to zero, we can experience numerical problems. See remark at the end of Appendix B for how to check for this problem in advance, or Section 2.5.1 for how to deal with it inside the algorithm.

2.4 The core algorithm

This section presents the core algorithm. It runs as follows: Find the target marginal moments from stochastic processes, from statistics or by specifying them directly. Generate n univariate discrete distributions that have those moments. Create a multivariate distribution by combining the univariate distributions as explained in the Introduction. Transform these outcomes so that they have the desired correlations and marginal moments. If the discrete distributions \mathbb{X}_i were independent, we would end up with a sample having exactly the desired properties.

To facilitate the reading, we have divided the algorithm in two parts. In the *input phase* we read the target properties specified by the user and transform them to a form needed by the algorithm. In the *output phase* we generate the outcomes and transform them to the original properties.

³Note that L always exists, since we assume R to be positive semi-definite.

2.4.1 The input phase

In this phase we work only with the target moments and correlations, we do not have any outcomes yet. This means that all the operations are fast and independent of the number of scenarios s .

Our goal is to generate an n -dimensional random variable \tilde{Z} with moments $TARMOM$ and a correlation matrix R . Since the matrix transformation needs zero means and variances equal to 1, we have to change the targets to match this demand. Instead of \tilde{Z} we will thus generate random variables \tilde{Y} with properties $\tilde{Y}(MOM; R)$, such that $MOM_1 = 0$, and $MOM_2 = 1$. \tilde{Z} is then computed at the very end of the algorithm as

$$\tilde{Z} = \alpha\tilde{Y} + \beta$$

In Appendix A we show that the values leading to the correct \tilde{Z} are:

$$\begin{aligned} \alpha &= TARMOM_2^{1/2} & MOM_3 &= \frac{TARMOM_3}{\alpha^3} \\ \beta &= TARMOM_1 & MOM_4 &= \frac{TARMOM_4}{\alpha^4} \end{aligned}$$

The final step in the input phase is to derive moments of independent univariate random variables \tilde{X}_i such that $\tilde{Y} = L \times \tilde{X}$ will have the target moments and correlations. To do this we need to find the Cholesky-decomposition matrix L , i.e. a lower-triangular matrix L so that $R = L \times L^T$.

The input phase then contains the following steps:

1. Specify the target moments $TARMOM$ and target correlation matrix R of \tilde{Z}
2. Find the normalized moments MOM of \tilde{Y}
3. Compute L and find the transformed moments $TRSFMOM$ of \tilde{X} — see Section 2.3.2




Figure 1: Input Phase

2.4.2 The output phase

In this phase we generate the outcomes for the independent random variable, transform them to get the intermediate-target moments and the target correlations and finally transform them to the moments specified by the user. Since the last transformation is a linear one, it will not change the correlations. All the transformations in this phase are with the outcomes, so the computing time needed for this phase is longer and increases with the number of scenarios s .

We start by generating discretizations of the univariate distributions \tilde{X}_i independently, one by one. This is a well-known problem and there are several possible ways to do it. We have used a method from *Fleishman, 1978*, in a way described in Section 2.3.1. The method starts by sampling from $\mathcal{N}(0, 1)$ and afterwards uses the cubic transformation to get the desired moments. For the starting $\mathcal{N}(0, 1)$ sample we use a random-number generator. An alternative would be to use a discretization of the distribution, or some other method for discretizing $\mathcal{N}(0, 1)$, for example the method described in *Høyland and Wallace, 2001*. They formulate a least squares model with outcomes of \tilde{X}_i as independent variables and the first four moments as dependent variables. The distance (sum of squares) of the actual moments from the target moments is then minimized.

Once we have generated the outcomes \mathbb{X}_i for the random variables \tilde{X}_i , we can proceed with the transformations. First $\mathbb{Y} = L \times \mathbb{X}$ to get the target correlations and then $\mathbb{Z} = \alpha\mathbb{Y} + \beta$ to get the user-specified moments⁴.

The output phase of the core algorithm consists of the following steps:

4. Generate outcomes \mathbb{X}_i of 1-dimensional variables \tilde{X}_i (independently for $i = 1 \dots n$)
5. Transform \mathbb{X} to target correlations: $\mathbb{Y} = L \times \mathbb{X} \rightarrow \mathbb{Y}$ is a set of scenarios for $\tilde{Y}(MOM; R)$
6. Transform to original moments: $\mathbb{Z} = \alpha\mathbb{Y} + \beta \rightarrow \mathbb{Z}$ is a set of scenarios for $\tilde{Z}(TARMOM; R)$



Figure 2: Output Phase

2.5 The modified algorithm

We know that the core algorithm gives us the exact results only if the random variables having support \mathbb{X}_i are independent. Since we generate each set of outcomes \mathbb{X}_i separately and have a limited number of outcomes (s), the sampled co-moments will most certainly differ from zero and the results of the core algorithm will be only approximate.

The modified algorithm is iterative, so all results are approximate, but with a pre-specified maximal error. Hence, we can control the quality of the results. We will again use the matrix transformation $\tilde{Y} = L \times \tilde{X}$, this time both forward and backward, as mentioned in 2.3.2. Recall that this transform allows us to obtain a desired correlation matrix, but it changes the moments while doing so.

In Section 2.3.1 we showed that the cubic transformation allows us to transform the general univariate random variable \tilde{X}_i to a variable with some target moments. This transformation changes the correlations, but if the starting \tilde{X}_i are close to the target distributions, the changes in the correlations are expected to be small.

Our idea is thus to apply the cubic and the matrix transformation, alternately, hoping that the algorithm converges and outcomes with both right moments and correlations are achieved. Hence, on one hand we cannot prove that the procedure always converges, on the other hand, the results are easily checked for correctness. This introduces iterative loops to the core algorithm, namely to Steps 4 and 5, in the following way:

In Step 4 we would like to generate the independent outcomes \mathbb{X}_i . Since independence is very hard to achieve, we change our target to generating uncorrelated outcomes \mathbb{X}_i , i.e. we seek to get \mathbb{X} with properties $\tilde{X}(TRSFMOM; I)$. We use an iterative approach to achieve those properties.

Since we do not control the higher co-moments, they will most likely not be zero, and the \mathbb{Y} obtained in Step 5 will not match the target properties. We thus need another loop there to end up with the desired \mathbb{Y} .

The iterative versions of Steps 4 and 5 are⁵:

Step 4

- 4.i. Generate n marginal distributions with target moments $TRSFMOM$ (independently) \rightarrow we get \mathbb{X} with correlation matrix R_1 close to I due to independent generation
- 4.ii. **let** $p = 1$ and $\mathbb{X}_1 = \mathbb{X}$
- 4.iii. **while** $dist(R_p; I) > \varepsilon_x$ **do**
- 4.iv. **do** Cholesky transformation: $R_p = L_p \times L_p^T$

⁴Note that we have stopped to speak about distributions and started to speak about outcomes, so we have to change the notation from \tilde{X} to \mathbb{X}

⁵As we have not found any better notation, in the rest of this section lower index p denotes an iteration counter, not a matrix column.

- 4.v. do backward transform $\mathbb{X}_p^* = L^{-1} \times \mathbb{X}_p \rightarrow$ zero correlations, wrong moments
- 4.vi. do cubic transform of \mathbb{X}_p^* with *TRSFMOM* as the target moments;
store results as $X_{p+1} \rightarrow$ right moments, wrong corr.
- 4.vii. compute correlation matrix R_{p+1}
- 4.viii. let $p = p + 1$
- 4.ix. let $\mathbb{X} = \mathbb{X}_p \rightarrow \mathbb{X}$ is from $\tilde{X}(TRSFMOM; I)$ with correlation error $dist(R_p; I) \leq \varepsilon_x$

The sum of squares of differences of elements is used as a distance of two matrices in 4.iii. The same distance is used also at points 5.iii and 5.ix below. Since \mathbb{X} is not a final output, the maximum error ε_x in Step 4 is typically higher than the corresponding ε_y in Step 5.

There are two possible outcomes from Step 4: X_p corresponding to random variables with right moments and slightly off correlations, and X_{p-1}^* corresponding to random variables with slightly off moments and zero correlations. We start with the latter in the Step 5, and denote it X^* .

Step 5

- 5.i. $\mathbb{Y}_1 = L \times \mathbb{X}^* \rightarrow$ both moments and correlations are incorrect
(due to higher co-moments different from zero)
- 5.ii. let $p = 1$ and let R_1 be the correlation matrix of \tilde{Y}
- 5.iii. while $dist(R_p; R) > \varepsilon_y$ do
- 5.iv. do Cholesky transformation: $R_p = L_p \times L_p^T$
- 5.v. do backward transform $\mathbb{Y}_p^* = L_p^{-1} \times \mathbb{Y}_p$
 \rightarrow zero correlations, incorrect moments
- 5.vi. do forward transform $\mathbb{Y}_p^{**} = L \times \mathbb{Y}_p^*$
 \rightarrow correct correlations (R), incorrect moments
- 5.vii. do cubic transform of \mathbb{Y}_p^{**} with *MOM* as the target moments;
store results as $\mathbb{Y}_{p+1} \rightarrow \mathbb{Y}_{p+1}$ has property $\tilde{Y}_{p+1}(MOM; R_{p+1})$
- 5.viii. let $p = p + 1$
- 5.ix. let $\mathbb{Y} = \mathbb{Y}_p \rightarrow \tilde{Y}$ has property $\tilde{Y}(MOM; R_p)$ with correlation error $dist(R_p; R) \leq \varepsilon_y$

Note that we can again choose two different outcomes from Step 5 (and thus from the whole algorithm). After Step 5.ix, \tilde{Y} has the right moments and (slightly) off correlation. If we prefer exact correlations, we can either use the last \mathbb{Y}_p^{**} , or repeat Steps 5.iv – 5.vi just before we go to Step 5.ix.

Note also that Steps 5.v and 5.vi are written as individual steps for the sake of clarity. Since we have always $s > n$ (usually $s \gg n$), it is much more efficient to join the two steps and calculate $\mathbb{Y}_p^{**} = (L \times L_p^{-1}) \times \mathbb{Y}_p$ directly.

2.5.1 Handling singular and ‘almost-singular’ correlation matrices

Since a singular R leads to zero(s) on a diagonal of the Cholesky-decomposition matrix L , we cannot compute the inverse transformation in step 3 of the algorithm – see Section 2.3.2. In Appendix B we argue that in the case of a badly conditioned (almost singular) R , we can evaluate the formulas, but the results will possibly be very inaccurate. We also suggest precautionary steps to avoid this problem. This section describes what to do when we nevertheless decide to proceed with the scenario generation.

The formulas found in Appendix B give us the best starting distribution \tilde{X} for our iterative procedure to find \tilde{Y} . More precisely, a starting point, such that if we find independent \tilde{X} with the given moments, we will not need any iterations at all. Yet, it is not the only possible starting point. According to our numerical tests, the convergence of the iterative procedure is very strong and the algorithm converges (almost) regardless the starting distribution \tilde{X} .

Therefore, if we are not able to find the starting \tilde{X}_i (due to division by zero), or its calculated properties are extreme, or even impossible (for example a negative kurtosis), we can just replace it by some other starting distribution. The question of which starting distribution to choose depends on individual taste, and we offer two possibilities.

One is to substitute every ‘incorrect’ number with the corresponding value from $\mathcal{N}(0, 1)$, the other is to find the closest ‘reasonable’ value. For both of them we need to set up what we

believe to be ‘reasonable’ values. Intervals like $E[\tilde{X}^3] \in (-3, 3)$ and $E[\tilde{X}^4] \in (0.5, 10)$ should be non-restrictive, yet sufficient to avoid numerical difficulties.

The procedure for computing the target moments *TRSFMOM* is then:

- Define the ‘safe’ intervals for $E[\tilde{X}^3]$ and $E[\tilde{X}^4]$
- for $i = 1$ to n do
- compute $E[\tilde{X}^3]$ and $E[\tilde{X}^4]$
- if $E[\tilde{X}^k]$ is out of range, replace it by
 - a) the corresponding value from $\mathcal{N}(0, 1)$
 - b) the border value

3 Extensions

3.1 Tree aggregation

Høyland and Wallace, 2001 argue that the quality of the scenario tree improves if we generate several small trees and aggregate them to a large one, instead of generating the large tree directly. By aggregation they understand creating the large tree as the union of the scenarios of all the small trees. If p_{ij} is the probability of scenario i in the j ’th tree, then that scenario obtains the probability $\frac{1}{N}p_{ij}$ in the union of N trees.

It is straightforward to implement the tree aggregation in our algorithm. If we choose to generate u subtrees of size t (where $ut = s$), we simply call the generation procedure u times with t as the number of scenarios, collect the scenario sets from all runs and report the complete set of scenarios at the end. For guidance on the size of u and t see *Høyland and Wallace, 2001*.

3.2 Multi-period case

In the multi-period case there is much freedom with respect to the generation of a scenario tree. *Høyland and Wallace, 2001* recommend a sequential approach. They start with the generation of the first period outcomes. The specifications for the second period will depend on the first period outcomes to capture time dependent effects like mean reversion and volatility clumping. Hence the second period outcomes can be generated one at a time with specifications dependent on the first period outcomes. Third period outcomes are generated in a similar way, dependent on the path for the first and the second period outcomes. This process then continues until enough periods are generated.

The important observation here is that the multi-period tree is generated by constructing single period trees, one by one. In the same way, by generalizing the single period case described in Section 2, the methods of this paper can also address the multi-period case.

4 Numerical results

This section presents the times needed to generate scenario trees with different numbers of random variables and scenarios. As a benchmark we use the algorithm from *Høyland and Wallace, 2001* where no decomposition is used. We have fixed probabilities in the benchmark model to make the algorithms comparable.

Both algorithms are implemented in AMPL with MINOS 5.5 as a solver. The test was done on a Pentium III 500 MHz machine with 256 MB memory, running Windows NT 4.0. As a distance used for evaluating the quality of the distributions in Steps 4 and 5 of the algorithm, we have used an absolute-mean-error defined as

$$MAE = \sqrt{\frac{1}{N_{el}} \sum_k (value_k - TARGET_k)^2}$$

where N_{el} is the number of elements in the sum. The distance was evaluated separately for moments ($N_{el} = 4n$) and for correlations ($N_{el} = \frac{n(n-1)}{2}$).

The stopping-values were $\varepsilon_x = 0.1$ ($\varepsilon_x = 0.2$ in the case of 20 r.v.) and $\varepsilon_y = 0.01$. Note that the distances are evaluated inside the algorithm, where all the data are scaled to variance equal

to 1, so they are scale-independent. The formulas in Appendix A show that the maximum error of the i -th moment in the final output is $TARMOM_i^i \cdot \varepsilon_y$.

The following tables summarise the results. The two blank cells are left because the expected running times were very high.

r.v.	40	100	200	1000
4	00:01	00:17	01:33	1:00:18
8	00:21	03:05	17:25	7:04:28
12	00:43	07:14	44:55	
20	03:47	37:22	4:00:00	

Results of the benchmark algorithm
([h]:mm:ss)

r.v.	40	100	200	1000
4	00:01	00:01	00:01	00:05
8	00:04	00:04	00:03	00:10
12	00:09	00:07	00:06	00:16
20	01:05	00:44	00:31	00:48

Results of our algorithm (mm:ss)

r.v.	40	100	200	1000
4	1.1×	15×	67×	770×
8	5.8×	50×	354×	2605×
12	4.6×	60×	465×	
20	3.5×	51×	464×	

Speed-Up

We see that the new algorithm is always faster, with speed-ups ranging from almost nothing for the smallest tree to thousands for the biggest ones. The speed-up increases with both the number of scenarios and the number of random variables, even if there are exceptions to the latter case. We can thus expect the speed-up to be significant for the real-life applications.

The biggest difference is in the way running time changes with the number of scenarios. While the time for the benchmark model increases approximately quadratically⁶, the time for our model increases less than linearly, in fact it decreases in most cases. This strange behaviour is caused by a better convergence for the large trees, i.e. the fact that the number of iterations decreases with the size of a tree: with 40 scenarios we need typically 1–2 iterations in Step 4 of the algorithm (generation of \tilde{X}) and 2–3 iterations in Step 5 (generation of \tilde{Y}), whereas with 1000 scenarios we usually need only 1 iteration in both parts. Since we cannot have less than 1 iteration, there will be no more improvement in the convergence for trees with more than 1000 scenarios. We can thus expect approximately linear time-dependency for very large trees.

We should however realise that the above results are ‘not fair’ to the benchmark algorithm, because we do not use it the recommended way. Instead of trying to create a tree of 1000 scenarios directly, we should generate 10 trees of 100 (or 25 trees of 40) scenarios and aggregate them, as described in Section 3, or in more detail in *Høyland and Wallace, 2001*. We should thus look at the speed-per-scenario. To make the results easier to read, the next table shows the fastest way to create scenario trees of different sizes. Note that it means to use the smallest tree (40 scenarios) for the benchmark model and the biggest tree (1000 scenarios) for our model.

r.v.	benchmark	our alg.	speed-up
4	00:35	00:05	7.5×
8	08:39	00:10	53×
12	17:54	00:21	52×
20	1:34:46	00:51	113×

1000 scenarios with the best tree-aggregation
([h]:mm:ss)

5 Future work

Even though we have achieved a substantial increase of speed compared to the benchmark algorithm, the tested implementation is still far from efficient. In the current AMPL implementation,

⁶This is not surprising, since the number of specifications in the model is $4n + \frac{n(n-1)}{2} = 0.5(n^2 + 7n)$

the algorithm spends most of the time with the Cholesky transformation and the communication with the solver before and after the cubic transform.

Both these critical times can be eliminated or decreased by implementing the algorithm in a compiled form. A C++ implementation is currently being developed at the SINTEF⁷ research institute, and according to the first tests is more than 10 times faster than our AMPL code.

The presented algorithm is well suited for parallel implementation, because most of its parts can be processed independently for every random variable. The only step of the algorithm that uses considerable amount of time and cannot be parallelised in this simple way is the Cholesky transformation of the correlation matrix, but parallel codes for the Cholesky transform can also be found.

Another area of future research is the scenario-tree aggregation. We will have to find out whether the recommendation to aggregate trees coming from *Høyland and Wallace, 2001* is valid also for our model. Note that while in their model there was a strong numerical argument for aggregating, our algorithm works better on larger trees. If it is nevertheless the case that we should aggregate, we will have to find the optimal number of trees to aggregate, given the number of random variables and scenarios.

6 Conclusion

We have presented an algorithm to generate scenario trees for multistage decision problems. The purpose of the algorithm is to speed up an existing scenario generation algorithm, which constructs multi-dimensional scenario trees that are consistent with specified moments and correlations.

The original algorithm constructs the multidimensional scenario tree by solving a single, potentially very large least squares problem, based. The main idea of the new algorithm is to decompose the least squares problem so that each marginal distribution is constructed separately. To combine the different marginal distributions so that the multivariate distribution satisfies the specified correlations, we apply the Cholesky decomposition and a cubic transformation in an iterative procedure.

Testing shows that our algorithm is significantly faster, with speed-up increasing with the size of a tree. With 20 random variables and 1000 scenarios our algorithm is more than 100 times faster than the benchmark algorithm with no decomposition.

This improvement is of great practical importance since it allows for generating bigger trees and for more interactive use of the model.

Acknowledgments

We would like to thank Erik Kole from the Maastricht University for pointing an error in the earlier version of the paper, and Matthias Nowak from SINTEF, Trondheim, for suggesting some important changes in notation and structure of the paper. Much of this work was done while Stein W. Wallace visited the Centre for Advanced Study at the Academy of Science and Letters in Oslo.

References

- Cariño, D. R. and Ziemba, W. T. (1998). Formulation of the Russell-Yasuda Kasai financial planning model. *Operations Research*, 46(4):443–449.
- Consigli, G. and Dempster, M. A. H. (1998). Dynamic stochastic programming for asset-liability management. *Annals of Operations Research*, 81:131–162.
- Dert, C. (1995). *Asset Liability Management for Pension Funds, A Multistage Chance Constrained Programming Approach*. PhD thesis, Erasmus University, Rotterdam, The Netherlands.

⁷The Foundation for Scientific and Industrial Research at the Norwegian Institute of Technology

- Fleishman, A. I. (1978). A method for simulating nonnormal distributions. *Psychometrika*, 43:521–532.
- Higham, N. J. (2000). Computing the nearest correlation matrix. Numerical Analysis Report No. 369, Department of Mathematics, University of Manchester.
- Høyland, K. and Wallace, S. W. (2001). Generating scenario trees for multi stage decision problems. *Management Science*, 47(2):295–307.
- Lurie, P. M. and Goldberg, M. S. (1998). An approximate method for sampling correlated random variables from partially-specified distributions. *Management Science*, 44(2):203–218.
- Mulvey, J. M. (1996). Generating scenarios for the towers perrin investment system. *Interfaces*, 26:1–13.
- Vale, C. David & Maurelli, V. A. (1983). Simulating multivariate nonnormal distributions. *Psychometrika*, 48(3):465–471.

Appendix

A Linear transformation $\tilde{Z} = \alpha\tilde{Y} + \beta$

The purpose of the algorithm is to generate an n -dimensional random variable \tilde{Z} with the target moments $TARMOM$. The algorithm needs the first two moments of all individual random variables \tilde{Z}_i equal to 0 and 1, respectively, hence we generate the random variables \tilde{Y}_i with $E[\tilde{Y}_i] = 0$ and $Var[\tilde{Y}_i] = 1$ instead, and compute \tilde{Z} as $\tilde{Z} = \alpha\tilde{Y} + \beta$ afterwards.

For every individual random variable \tilde{Z}_i we need to find parameters α_i , β_i and moments MOM_{i3} and MOM_{i4} of \tilde{Y}_i , so that $\tilde{Z}_i = \alpha_i\tilde{Y}_i + \beta_i$ will have the desired moments $TARMOM_i$. We state the moments of \tilde{Z}_i as functions of α_i , β_i and moments of \tilde{Y}_i , utilizing $E[\tilde{Y}_i] = 0$ and $Var[\tilde{Y}_i] = 1$, and solve the equations so that \tilde{Z}_i has the target moments.

The formulas can be stated either for individual random variables \tilde{Z}_i , or as vector equations for all random variables \tilde{Z}_i at once. The only difference is the presence/absence of the index i in all elements. To make the formulas easier to read, we use the vector form:

$$\begin{aligned} \mathbb{E}[\tilde{Z}] &= \mathbb{E}[\alpha\tilde{Y} + \beta] = \alpha\mathbb{E}[\tilde{Y}] + \beta = \beta = TARMOM_1 \quad \longrightarrow \quad \text{Note: } \tilde{Z} - \mathbb{E}[\tilde{Z}] = \alpha\tilde{Y} \\ Var[\tilde{Z}] &= \mathbb{E}[(\tilde{Z} - \mathbb{E}[\tilde{Z}])^2] = \mathbb{E}[(\alpha\tilde{Y})^2] = \alpha^2\mathbb{E}[\tilde{Y}^2] = \alpha^2Var[\tilde{Y}] = \alpha^2 = TARMOM_2 \\ Skew[\tilde{Z}] &= \mathbb{E}[(\tilde{Z} - \mathbb{E}[\tilde{Z}])^3] = \dots = \alpha^3Skew[\tilde{Y}] = \alpha^3MOM_3 = TARMOM_3 \\ Kurt[\tilde{Z}] &= \mathbb{E}[(\tilde{Z} - \mathbb{E}[\tilde{Z}])^4] = \dots = \alpha^4Kurt[\tilde{Y}] = \alpha^4MOM_4 = TARMOM_4 \end{aligned}$$

Solving for the unknown parameters we get:

$$\begin{aligned} \alpha &= \sqrt{TARMOM_2} \\ \beta &= TARMOM_1 \\ MOM_3 &= \frac{TARMOM_3}{\alpha^3} = \frac{TARMOM_3}{TARMOM_2^{3/2}} \\ MOM_4 &= \frac{TARMOM_4}{\alpha^4} = \frac{TARMOM_4}{TARMOM_2^2} \end{aligned}$$

B Matrix transformation $\tilde{Y} = L \times \tilde{X}$

We seek an n -dimensional random variable \tilde{Y} with zero means, variances equal to 1, skewness MOM_3 , kurtosis MOM_4 , and a correlation matrix $R = L \times L^T$, where L is the lower-triangular matrix. To generate the outcomes \mathbb{Y} of the random variable \tilde{Y} we first generate a matrix \mathbb{X} of outcomes of the random variable \tilde{X} with \tilde{X}_i as independent random variables, and afterwards compute \mathbb{Y} as $\mathbb{Y} = L \times \mathbb{X}$. Note that the k 'th moment of \tilde{Y} is equal to $\mathbb{E}[\tilde{Y}^k]$ because of the zero means.

During the algorithm we use the matrix transformation in two different cases. First it is used on distributions, which are abstract objects, so we work only with their (statistical) properties. In the second part of the algorithm, we transform the outcomes, which are matrices of numbers. The formulas are the same in both cases, the only difference is the use of \tilde{X} in the first case and \mathbb{X} in the latter.

Note that we can use both notations without changing the rest of the equations: \tilde{X} is a column vector of univariate random variables \tilde{X}_i , while \mathbb{X} can be seen as a column 'vector' of the outcomes \mathbb{X}_i . The only difference is that the \mathbb{X}_i is a row vector of size s , while the \tilde{X}_i is a 1-dimensional random variable.

We use the distribution-notation in the rest of the appendix. The matrix transformation in a column-wise form is then:

$$\tilde{Y}_i = \left(L \times \tilde{X} \right)_i = \sum_{j=1}^n L_{ij} \tilde{X}_j = \sum_{j=1}^i L_{ij} \tilde{X}_j$$

where the last equality comes from the fact that L is a lower-triangular matrix and therefore $L_{ij} = 0$ for $j > i$.

Theorem – properties of $\tilde{Y} = L \times \tilde{X}$

Assume we have an n -dimensional random variable \tilde{X} with the following properties:

- i.* $\mathbb{E}[\tilde{X}^k]$ exists for $k = 1 \dots 4$
- ii.* $\mathbb{E}[\tilde{X}] = 0$ and $\mathbb{E}[\tilde{X}^2] = 1$
- iii.* univariate r.v. \tilde{X}_i, \tilde{X}_j are independent for $i \neq j$

Assume further that L is a lower-triangular matrix of size $n \times n$ such that $R = L \times L^T$, where R is a correlation matrix, i.e. R is a symmetric positive semi-definite matrix with 1's on the main diagonal.

If we then define a random variable \tilde{Y} as $\tilde{Y} = L \times \tilde{X}$, it has the following properties:

- iv.* $\mathbb{E}[\tilde{Y}^k]$ exists for $k = 1 \dots 4$
- v.* $\mathbb{E}[\tilde{Y}] = 0$ and $\mathbb{E}[\tilde{Y}^2] = 1$
- vi.* \tilde{Y} has a correlation matrix $R = L \times L^T$
- vii.* $\mathbb{E}[\tilde{Y}_i^3] = \sum_{j=1}^i L_{ij}^3 \mathbb{E}[\tilde{X}_j^3]$
- viii.* $\mathbb{E}[\tilde{Y}_i^4] - 3 = \sum_{j=1}^i L_{ij}^4 (\mathbb{E}[\tilde{X}_j^4] - 3)$

Consequence

Under the assumptions of the theorem, with an additional assumption that R is regular (and therefore positive-definite), we can express the moments of \tilde{X} as:

$$\begin{aligned} \text{ix. } \mathbb{E}[\tilde{X}_i^3] &= \frac{1}{L_{ii}^3} \left(\mathbb{E}[\tilde{Y}_i^3] - \sum_{j=1}^{i-1} L_{ij}^3 \mathbb{E}[\tilde{X}_j^3] \right) \\ \text{x. } \mathbb{E}[\tilde{X}_i^4] - 3 &= \frac{1}{L_{ii}^4} \left[\mathbb{E}[\tilde{Y}_i^4] - 3 - \sum_{j=1}^{i-1} L_{ij}^4 (\mathbb{E}[\tilde{X}_j^4] - 3) \right] \end{aligned}$$

Proof

iv.

Will be proved by proving the rest.

v. & vi.

From assumptions *ii.* and *iii.* we get the variance-covariance matrix of \tilde{X} as $\mathbb{E}[\tilde{X} \times \tilde{X}^T] = I$. The mean of \tilde{Y} follows directly:

$$\mathbb{E}[\tilde{Y}] = \mathbb{E}[L \times \tilde{X}] = L \times \mathbb{E}[\tilde{X}] = 0$$

For proving the second part of *v.* and *vi.* we compute the variance-covariance matrix of \tilde{Y} :

$$\mathbb{E}[\tilde{Y} \times \tilde{Y}^T] = \mathbb{E}[L \times \tilde{X} \times \tilde{X}^T \times L^T] = L \times \mathbb{E}[\tilde{X} \times \tilde{X}^T] \times L^T = L \times I \times L^T = L \times L^T = R$$

Since R is a correlation matrix, it has 1's on the main diagonal, which proves *v.* It also means that R is both the variance-covariance matrix and the correlation matrix, which proves *vi.*

vii. & ix.

$$\begin{aligned}\mathbb{E} \left[\tilde{Y}_i^3 \right] &= \mathbb{E} \left[\left(\sum_{j=1}^n L_{ij} \tilde{X}_j \right) \left(\sum_{k=1}^i L_{ik} \tilde{X}_k \right) \left(\sum_{l=1}^i L_{il} \tilde{X}_l \right) \right] \\ &= \mathbb{E} \left[\sum_{j,k,l=1}^i L_{ij} L_{ik} L_{il} \tilde{X}_j \tilde{X}_k \tilde{X}_l \right] = \sum_{j,k,l=1}^i L_{ij} L_{ik} L_{il} \mathbb{E} \left[\tilde{X}_j \tilde{X}_k \tilde{X}_l \right]\end{aligned}$$

The assumption of independence of \tilde{X}_i implies that $\mathbb{E}[\tilde{X}_k \tilde{X}_l \tilde{X}_m] = 0$ if at least one index differs from the other two. To show it, let us assume that $j \neq k$ and $j \neq l$. This means that \tilde{X}_j is independent of both \tilde{X}_k and \tilde{X}_l and we can thus write

$$\mathbb{E} \left[\tilde{X}_j \tilde{X}_k \tilde{X}_l \right] = \mathbb{E} \left[\tilde{X}_j \right] \mathbb{E} \left[\tilde{X}_k \tilde{X}_l \right] = 0.$$

Since two equal indices means that one is different, the only case where $\mathbb{E}[\tilde{X}_j \tilde{X}_k \tilde{X}_l] \neq 0$ is for $j = k = l$. This simplifies the triple sum to a single one and proves *vii.*:

$$\mathbb{E} \left[\tilde{Y}_i^3 \right] = \sum_{j=1}^i L_{ij}^3 \mathbb{E} \left[\tilde{X}_j^3 \right].$$

The proof of the consequence *ix.* is straightforward since the transformation is triangular and we can write the inversion directly. We start with the $\mathbb{E}[\tilde{X}_1^3]$, and continue ‘downward’ to $\mathbb{E}[\tilde{X}_n^3]$:

$$\begin{aligned}\mathbb{E} \left[\tilde{X}_1^3 \right] &= \frac{1}{L_{11}^3} \mathbb{E} \left[\tilde{Y}_1^3 \right] \\ \mathbb{E} \left[\tilde{X}_i^3 \right] &= \frac{1}{L_{ii}^3} \left(\mathbb{E} \left[\tilde{Y}_i^3 \right] - \sum_{j=1}^{i-1} L_{ij}^3 \mathbb{E} \left[\tilde{X}_j^3 \right] \right)\end{aligned}$$

Note that we can divide by L_{ii} because we have assumed $R = L \times L^T$ to be positive-definite, so we have $L_{ii} > 0 \forall i$.

viii. & x.

$$\begin{aligned}\mathbb{E} \left[\tilde{Y}_i^4 \right] &= \mathbb{E} \left[\left(\sum_{j=1}^i L_{ij} \tilde{X}_j \right) \left(\sum_{k=1}^i L_{ik} \tilde{X}_k \right) \left(\sum_{l=1}^i L_{il} \tilde{X}_l \right) \left(\sum_{m=1}^i L_{im} \tilde{X}_m \right) \right] \\ &= \mathbb{E} \left[\sum_{j,k,l,m=1}^i L_{ij} L_{ik} L_{il} L_{im} \tilde{X}_j \tilde{X}_k \tilde{X}_l \tilde{X}_m \right] \\ &= \sum_{j,k,l,m=1}^i L_{ij} L_{ik} L_{il} L_{im} \mathbb{E} \left[\tilde{X}_j \tilde{X}_k \tilde{X}_l \tilde{X}_m \right]\end{aligned}$$

With the same argument as for the third moment, the assumption of independent \tilde{X}_i again means that $\mathbb{E}[\tilde{X}_k \tilde{X}_l \tilde{X}_m \tilde{X}_n] = 0$ if at least one of the indices $j, k, l,$ and m is different from all the others. This means that at least two indices must be the same in order to make the mean non-zero. Since three equal indices implies one different, we are left with two possibilities for non-zeroes: either all the indices are equal, or there are two different indices, each one twice.

In the former case we have $\mathbb{E}[\tilde{X}_j \tilde{X}_k \tilde{X}_l \tilde{X}_m] = \mathbb{E}[\tilde{X}_j^4]$. In the latter case, with for example $j = l$ and $k = m$, we get

$$\mathbb{E} \left[\tilde{X}_j \tilde{X}_k \tilde{X}_l \tilde{X}_m \right] = \mathbb{E} \left[\tilde{X}_j^2 \tilde{X}_k^2 \right] = \mathbb{E} \left[\tilde{X}_j^2 \right] \mathbb{E} \left[\tilde{X}_k^2 \right] = 1 \cdot 1 = 1$$

Since we have a sum over all the combinations of indices k, l, m and n , the latter case occurs $\binom{4}{2} = 6$ times for every pair of indices. We thus have to multiply that element by 6 and sum it over $k < l$ to prevent the element from being counted twice.

We have thus simplified the sum to:

$$\mathbb{E} \left[\tilde{Y}_i^4 \right] = \sum_{j=1}^i L_{ij}^4 \mathbb{E} \left[\tilde{X}_j^4 \right] + 6 \sum_{\substack{j,k=1 \\ j < k}}^i L_{ij}^2 L_{ik}^2$$

Now we need to use the fact that L is a matrix from the Cholesky decomposition of the correlations matrix R . Since R has 1's on the main diagonal the following relationship holds:

$$1 = R_{ii} = (L \times L^T)_{ii} = \sum_{j=1}^n L_{ij} (L^T)_{ji} = \sum_{j=1}^n L_{ij}^2 = \sum_{j=1}^i L_{ij}^2 = 1$$

Using this result, we can eliminate the double-sum from the formula for $\mathbb{E} \left[\tilde{Y}_i^4 \right]$. To do this, we express the double-sum in two different ways and then sum them together.

Note that the first equality in the second line comes from a simple switching of letters.

$$\begin{aligned} \sum_{\substack{j,k=1 \\ j < k}}^i L_{ij}^2 L_{ik}^2 &= \begin{cases} \sum_{j=1}^i \sum_{k=j+1}^i L_{ij}^2 L_{ik}^2 = \sum_{j=1}^i L_{ij}^2 \sum_{k=j+1}^i L_{ik}^2 \\ \sum_{k=1}^i \sum_{j=1}^{k-1} L_{ij}^2 L_{ik}^2 = \sum_{j=1}^i \sum_{k=1}^{j-1} L_{ik}^2 L_{ij}^2 = \sum_{j=1}^i L_{ij}^2 \sum_{k=1}^{j-1} L_{ik}^2 \end{cases} \\ 2 \sum_{\substack{j,k=1 \\ j < k}}^i L_{ij}^2 L_{ik}^2 &= \sum_{j=1}^i L_{ij}^2 \left(\sum_{k=1}^{j-1} L_{ik}^2 + \sum_{k=j+1}^i L_{ik}^2 \right) = \sum_{j=1}^i L_{ij}^2 \left(\sum_{k=1}^i L_{ik}^2 - L_{ij}^2 \right) \\ &= \left(\sum_{j=1}^i L_{ij}^2 \right) \left(\sum_{k=1}^i L_{ik}^2 \right) - \sum_{j=1}^i L_{ij}^4 = 1 - \sum_{j=1}^i L_{ij}^4 \end{aligned}$$

Now we substitute the result into the formula for $\mathbb{E} \left[\tilde{Y}_i^4 \right]$ and get:

$$\begin{aligned} \mathbb{E} \left[\tilde{Y}_i^4 \right] &= \sum_{j=1}^i L_{ij}^4 \mathbb{E} \left[\tilde{X}_j^4 \right] + 6 \sum_{\substack{j,k=1 \\ j < k}}^i L_{ij}^2 L_{ik}^2 = \sum_{j=1}^i L_{ij}^4 \mathbb{E} \left[\tilde{X}_j^4 \right] + 3 - 3 \sum_{j=1}^i L_{ij}^4 \\ &= \sum_{j=1}^i L_{ij}^4 \left(\mathbb{E} \left[\tilde{X}_j^4 \right] - 3 \right) + 3 \end{aligned}$$

To prove *viii*, we need only subtract 3 from both sides of the formula. Note that we get a formula with kurtosis decreased by 3, i.e. with kurtosis relative to $\mathcal{N}(0, 1)$:

$$\mathbb{E} \left[\tilde{Y}_i^4 \right] - 3 = \sum_{j=1}^i L_{ij}^4 \left(\mathbb{E} \left[\tilde{X}_j^4 \right] - 3 \right)$$

Since the formula has the same form as the one for kurtosis, the second part of the consequence is obtained the same way as for kurtosis. This concludes the proof.

$$\begin{aligned} \mathbb{E} \left[\tilde{X}_1^4 \right] - 3 &= \frac{1}{L_{11}^4} \mathbb{E} \left[\tilde{Y}_1^4 \right] - 3 \\ \mathbb{E} \left[\tilde{X}_i^4 \right] - 3 &= \frac{1}{L_{ii}^4} \left[\mathbb{E} \left[\tilde{Y}_i^4 \right] - 3 - \sum_{j=1}^{i-1} L_{ij}^4 \left(\mathbb{E} \left[\tilde{X}_j^4 \right] - 3 \right) \right] \end{aligned}$$

Q.E.D.

Remarks

- Note that the set of moments $(0, 1, 0, 3)$ is an invariant of this transformation. This confirms the known theoretical result that the linear transformations preserve normality. In the context of our algorithm it means that if we generate normal variables, we can skip Step 3 of the algorithm.
- Even if we know that the regularity of R gives us $L_{ii} \neq 0$, we still can have $L_{ii} \approx 0$ in a case of ‘almost-singular’ correlation matrix R . As we divide by L_{ii}^3 and L_{ii}^4 during the inverse transformation, the chance of numerical instability is substantial. The result of those instabilities can be either extreme values of the target skewness or kurtosis of \tilde{X} , or even a negative kurtosis.

The important question is thus when this problem occurs, i.e. when does \tilde{Y} have an almost-singular correlation matrix R ? The basic answer is that this happens if one \tilde{Y}_i is almost a linear combination of the others, in other words, it can be almost fully explained using the other \tilde{Y}_j . The most obvious example is a correlation matrix with very high numbers (abs. value > 0.9), yet we should be aware that with increasing n (dimension of \tilde{Y}) we increase the chance of having a dependency also with ‘moderate’ correlations.

To avoid the possible problems, we should perform some tests before starting the algorithm. The easiest (but not sufficient) is to check if $L_{ii} \geq \varepsilon > 0$ for all asset classes i . Note that if we use the method from *Lurie and Goldberg, 1998* to ensure positive semi-definiteness of the correlation matrix, we can simply add the constraints $\{L_{ii} \geq \varepsilon > 0\}$ to the QP model. Although more difficult, it would be advisable to check the condition number of the correlation matrix. This number is defined as the ratio $\frac{\lambda_{max}(R)}{\lambda_{min}(R)}$ of maximum and minimum eigenvalues of the matrix R and measures how ill-conditioned the matrix is. Higher values indicate that the matrix is close to singular, so the inverse transformation can return incorrect results.

In case the above precautions fail, we suggest a simple heuristic for handling the singularity and ‘almost-singularity’ problem inside our algorithm – see Section 2.5.1. It is a crude procedure, but it should be sufficient in most cases.

C Cubic transformation $\tilde{Y} = a + b\tilde{X} + c\tilde{X}^2 + d\tilde{X}^3$

The purpose of this transformation is to produce a univariate random variable \tilde{Y}_i with specified first four moments $\mathbb{E}[\tilde{Y}_i^k], k = 1, \dots, 4$, given the random variable \tilde{X}_i with known first 12 moments $\mathbb{E}[\tilde{X}_i^k], k = 1, \dots, 12$.

The problem is to find the transform parameters a_i, b_i, c_i and d_i . For this we have to express $\mathbb{E}[\tilde{Y}_i^k]$ as functions of $\mathbb{E}[\tilde{X}_i^k]$. The formulas can be stated either for univariate random variables \tilde{Y}_i , or as a vector equations for the random vector \tilde{Y} . The only difference is the presence/absence of index i in all elements. To make the formulas easier to read, we use the vector form.

$$\begin{aligned}
\mathbb{E}[\tilde{Y}] &= a + b\mathbb{E}[\tilde{X}] + c\mathbb{E}[\tilde{X}^2] + d\mathbb{E}[\tilde{X}^3] \\
\mathbb{E}[\tilde{Y}^2] &= d^2\mathbb{E}[\tilde{X}^6] + 2cd\mathbb{E}[\tilde{X}^5] + (2bd + c^2)\mathbb{E}[\tilde{X}^4] + (2ad + 2bc)\mathbb{E}[\tilde{X}^3] \\
&\quad + (2ac + b^2)\mathbb{E}[\tilde{X}^2] + 2ab\mathbb{E}[\tilde{X}] + a^2 \\
\mathbb{E}[\tilde{Y}^3] &= d^3\mathbb{E}[\tilde{X}^9] + 3cd^2\mathbb{E}[\tilde{X}^8] + (3bd^2 + 3c^2d)\mathbb{E}[\tilde{X}^7] + (3ad^2 + 6bcd + c^3)\mathbb{E}[\tilde{X}^6] \\
&\quad + (6acd + 3b^2d + 3bc^2)\mathbb{E}[\tilde{X}^5] + (a(6bd + 3c^2) + 3b^2c)\mathbb{E}[\tilde{X}^4] \\
&\quad + (3a^2d + 6abc + b^3)\mathbb{E}[\tilde{X}^3] + (3a^2c + 3ab^2)\mathbb{E}[\tilde{X}^2] + 3a^2b\mathbb{E}[\tilde{X}] + a^3 \\
\mathbb{E}[\tilde{Y}^4] &= d^4\mathbb{E}[\tilde{X}^{12}] + 4cd^3\mathbb{E}[\tilde{X}^{11}] + (4bd^3 + 6c^2d^2)\mathbb{E}[\tilde{X}^{10}] \\
&\quad + (4ad^3 + 12bcd^2 + 4c^3d)\mathbb{E}[\tilde{X}^9] + (12acd^2 + 6b^2d^2 + 12bc^2d + c^4)\mathbb{E}[\tilde{X}^8] \\
&\quad + (a(12bd^2 + 12c^2d) + 12b^2cd + 4bc^3)\mathbb{E}[\tilde{X}^7] + (6a^2d^2 + a(24bcd + 4c^3) \\
&\quad + 4b^3d + 6b^2c^2)\mathbb{E}[\tilde{X}^6] + (12a^2cd + a(12b^2d + 12bc^2) + 4b^3c)\mathbb{E}[\tilde{X}^5] \\
&\quad + (a^2(12bd + 6c^2) + 12ab^2c + b^4)\mathbb{E}[\tilde{X}^4] + (4a^3d + 12a^2bc + 4ab^3)\mathbb{E}[\tilde{X}^3] \\
&\quad + (4a^3c + 6a^2b^2)\mathbb{E}[\tilde{X}^2] + 4a^3b\mathbb{E}[\tilde{X}] + a^4
\end{aligned}$$

Note that since the matrix transformation $\tilde{Y}_i = L \times \tilde{X}_i$ does not change the first two moments, we have $E[\tilde{X}_i] = 0$ and $E[\tilde{X}_i^2] = 1$. We could thus simplify the above formulas slightly by assuming that the first two moments are exactly 0 and 1. In our implementation we have, however, used the formulas in the presented form, computing the actual moments.

To solve the equations, we have formulated a non-linear least squares problem with a , b , c and d as independent variables and moments of \tilde{Y}_i as dependent variables. There are of course other possibilities.