European Journal of Operational Research xxx (xxxx) xxx

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Interfaces with Other Disciplines

Improved scalability and risk factor proxying with a two-step principal component analysis for multi-curve modelling

Philip J. Atkins[a], Mark Cummins[b,*]

[a] BP Oil International Ltd (retired), United Kingdom
[b] Dublin City University, Dublin, Ireland

ARTICLE INFO

Article history:
Received 15 January 2021
Accepted 29 April 2022
Available online xxx

Keywords:
Risk management
Multiple price curve modelling
Two-step PCA
Risk factor proxying
Energy markets

ABSTRACT

We consider the practice-relevant problem of modelling multiple price curves to support activities such as price curve simulation and risk management. In this multi-curve setting, the challenge is to jointly capture the risk-factor relationships within each curve and the risk-factor relationships between the curves. Contributing to the existing literature, we develop a novel two-step Principal Component Analysis (PCA) method, which we label PCA^2, that addresses this challenge. The concept of PCA^2 first derives components describing the dynamics of each curve, and then, second, combines these to describe the dynamics across all the curves. The benefits of PCA^2 over PCA are: (i) improved scalability allowing for greater computational efficiency and smaller data structures rendering multi-threading more feasible; (ii) components that remain identifiable at the curve level; and (iii) leveraging the last property, PCA^2, unlike PCA, offers the capability of proxying new curves for which limited historical data exists, using the first-step components from a related curve and estimating second-level correlations empirically. PCA^2 is a novel multi-curve modelling approach that will appeal, for these reasons, to many practitioners, especially those working in risk management.

1. Introduction

The modelling of complex curve dynamics in financial markets is recognised as a significant challenge, whether, for example, yield curves in interest rate markets or forward price curves in commodity markets. Factor analysis has provided a useful multifactor modelling framework to capture these complex curve dynamics. Schmidt (2011) provides a recent review of the interest rate modelling literature, while Carmona & Culon (2014) provide a useful survey of the literature on structural forward price curve models in commodity markets.

In the interest rate setting, some of earliest studies applying factor analysis include Stambaugh (1988), Steeley (1990), and Litterman & Sheinkman (1991), with the latter providing the well-established 'shift', 'tilt' and 'bend' interpretations of the first three principal components that provide a high degree of explainability of within curve variation. Principal Component Analysis (PCA) provides a suitable data-driven statistical method for isolating orthogonal factors that explain curve variation and that allows one to then reduce the dimensionality of the problem space. Alexander (2008, 2009) provides an excellent overview of PCA application for multifactor modelling and risk management. Early applications of PCA include Barber & Cpper (1996) who use PCA to factorise interest rate curve movements for the purpose of immunising interest rate sensitive portfolios from multi-directional curve shifts. Singh (1997) considers the use of PCA for dimension reduction in the calculation of Value-at-Risk. PCA has also seen application in the multifactor modelling framework of Heath, Jrrow, & Mrton (1992) (HJM). PCA provides a parsimonious way of reducing the dimensionality of the HJM model, while allowing for efficient estimation of the within curve covariances. An example of PCA used in this way is the model specification analysis conducted by Buhler, Urig-Homburg, Wlter, & Wber (1999), where the performance of a range of multifactor HJM type models are compared when applied for interest rate option valuation. Beyond the Gaussian framework, PCA has been applied to non-Gaussian problems. For example, Ballotta, Fsai, Lregian, & Prez (2019) use PCA to reduce the dimensionality of the problem in a multivariate Lvy model setting, which then allows for consistent estimation of model parameters in a manner that is notably more computationally efficient.

With much of the literature focused on identifying the appropriate number of factors for given modelling situations, some studies have sought to critically appraise the performance of factor

* Corresponding author.
E-mail address: mark.cummins@dcu.ie (M. Cummins).

https://doi.org/10.1016/j.ejor.2022.04.044

analysis, and PCA in particular. For example, Lekkos (2000) examines whether, in the interest rate setting, apparent explanatory power factor analysis is an artifact induced by the arbitrage constraints. In a multi-country study, the author finds that three factors explain a high proportion of spot rate variance, but that performance in respect of forward rates is less impressive. Falkenstein & Hnweck (1997) examine principal component based approaches to interest rate risk management, through a comparison of principal-components hedges to covariance-consistent hedges. The authors conclude that the principal-components hedges can become unstable when dealing with three or more hedging instruments and hedging performance can therefore deteriorate over time. Soto (2004) highlights similar issues in the application of PCA-based multifactor modelling, providing evidence of a notable deterioration in model performance over time. Lord & Plsser (2007) revisit the shift-tilt-bend interpretations to rigorously investigate whether the interpretations are genuinely driven by the underlying curve movements or a mere artefact of the PCA implementation, with a mixed conclusion that both may be true under certain conditions.

Borrowing from this interest rate literature, we have seen wide application of factor analysis in other markets. In the credit space, for instance, factor analysis has been used for corporate failure prediction and credit modelling. West (1985), in a seminal study, uses factor analysis to generate the factor inputs to a logit estimation of bank difficulty. Canbas, Cbuk, & Klic (2005) in a similar manner use PCA to determine a set of bank characteristic factors for input into logit/probit based models for the prediction of commercial bank failure. Shie, Chen, & Liu (2012) use PCA as a dimension reduction step in their application of a hybrid Support Vector Machine based approach to corporate failure prediction. In the commodity space, Cortazar & Shwartz (1994) provide a seminal study on the application of PCA for the modelling of multiple factors in commodity forward curves, with application to derivatives valuation. Clewlow & Srickland (1999) similarly show how PCA can be used in a multifactor HJM model with application to energy derivatives valuation, while Blanco, Sronow, & Seszyn (2002) provide a practitioner perspective on the application of PCA for energy risk management. Tolmasky & Hindanov (2002) and Chantziara & Siadopoulos (2008) apply PCA to model the factors driving petroleum futures, while electricity futures curve modelling is the context for the application of PCA by Koekebakker & Olmar (2005).

The literature cited thus far, and the wider associated literature, has primarily focused on the case of single curve modelling, yet in practice, risk management activities, trading strategy design, and exotic derivatives valuation, require understanding and modelling the dynamics across curves as well as within curves. Literature with a focus on multi-curve modelling is more sparse and so our study contributes to the literature. Amongst this literature, there are some notable studies that consider the multi-curve setting and that motivate our work. Morino & Rnggaldier (2014) and Grasselli & Mglietta (2016), for example, explore multi-curve modelling within the short rate class of models, while the LIBOR market model is the context for the study of Mercurio (2010). Moreni & Pllavicini (2014), Crepey, Gbac, & Nuyen (2011), Cuchiero, Fntana, & Goatto (2016) and Sabelli, Poppi, Stzia, & Brmetti (2018) all present multi-curve HJM applications in an interest rate setting. Crepey, Gbac, Nor, & Sovmand (2015) and Fanelli (2016) develop defaultable HJM models for the simultaneous evolution of multiple interest rate curves, with the latter extending the HJM model to a class of credit models.

In line with our own focus on PCA, Gerhart & Ltkebohmert (2020) seek to capture cross-tenor dependencies in forecasting multiple yield curves simultaneously. Applying PCA to the covariance matrix of multiple curves, the authors are able to identify common factors driving a portfolio of yield curves. These factors are shown to follow the usual shift-tilt-bend interpretation of individual interest rate curves. Laurini & Oashi (2015) highlight, however, that in the interest rate setting, there is a propagation of measurement error into the PCA from the yield curve bootstrapping step. This is not an issue in the commodity market setting we consider as forward price curves are directly observable.[1] In a commodity market context, Tolmasky & Hindanov (2002) and Chantziara & Siadopoulos (2008) consider the dynamics between forward price curves. As in the case of Gerhart & Ltkebohmert (2020), these authors apply a single PCA step to the multiple forward curves jointly.

In a somewhat related strand of literature, it is considered whether a number of populations share factors. Flury (1984) calls this the 'hypothesis of common "PC's" and introduces the term 'common principal components (CPC's)' for which is described a derivation using maximum likelihood estimation. Trendalov (2010) subsequently proposes a stepwise method for obtaining CPCs. However, the objective of factor analysis is to gain a statistical understanding of population discriminants, whereas for multi-curve modelling the PCs are a means to an end, not an end in themselves. We require the PCs to capture price curve dynamics for subsequent simulation and risk management work, and so have no reason to expect the CPC hypothesis of Flury (1984) to hold in this domain of application. Indeed, we account for a setting where PCs may be meaningfully different across multiple curves.

In considering the practice-relevant problem of multiple price curve modelling, we contribute to the literature through the development of a novel two-step PCA approach, which we label PCA^2. This two-step PCA offers considerable benefits over a one-step PCA. Our approach is most closely aligned to the independent work of Ellefsen & Slavounos (2009), who also propose a two-step PCA approach. However, Ellefsen & Slavounos (2009) implement a hierarchical approach within a hybrid framework, whereas with PCA^2 we combine the two stages into a set of "shock curves" that can be used akin to a one-step PCA. To clarify this distinction, we apply the second stage PCA to the covariance matrix derived from the component returns histories, whereas Ellefsen & Slavounos (2009) apply their second stage PCA to the covariance matrix derived from the raw returns histories. We address the merits of our PCA^2 approach next.

The concept of PCA^2 is to break the derivation of representative components into two steps, in which the first derives components describing *within*-curve relationships (i.e. volatilities and correlations), and the second combines these in a way that captures *between*-curve relationships (i.e. correlations[2]). PCA^2 has been developed to address a number of issues. There is a sense in which the usual single step PCA applied to a portfolio of curves finds the optimal components with which to derive the overall dynamics, but looking at all risk factors for all price curves in one go in practical settings with many price curve exposures leads to a very large covariance matrix, which scales quadratically as more price curves or risk factors are added. This means the computational cost of performing PCA becomes prohibitive when a large number of risk factors and curves are involved. Such large data structures also create challenges to effective parallelisation of computation, where the processing costs of distributing large blocks of data to the processors become prohibitive also. Furthermore, the 'all-or-nothing'

[1] Laurini & Oashi (2015) note that the covariance matrix estimation method can affect the results of PCA for forward curves, in addition to the problem raised in respect of measurement errors. Even with observed forward curves the presence of temporal dependence can influence the PCA decomposition. This is not a problem for the contribution of our study, which does not directly discuss the estimation of the covariance matrix. We thank an anonymous referee for this insight.

[2] The risk-factor volatilities are fully captured at the first step, so at the second step only correlations are required, as explained in Sections 2.3 and 2.4 below.

nature of standard PCA requires full datasets for all price curves and risk factors to be modelled. This is often not the case, for example, in the energy trading sector when new oil grades or trading hubs are introduced. An alternative approach is therefore needed, which PCA^2 offers.

The benefits of PCA^2 over PCA are therefore: (i) improved scalability allowing for greater computational efficiency and smaller data structures rendering multi-threading more feasible; (ii) components that remain identifiable at the curve level; and (iii) leveraging the last property, PCA^2, unlike PCA, offers the capability of proxying new curves for which limited historical data exists, using the first-step components from a related curve and estimating second-level correlations empirically. PCA^2 is a novel multi-curve modelling approach that will appeal, for these reasons, to many practitioners, especially those working in risk management.

The rest of the paper is organised as follows: Section 2 provides a concise recap of PCA then builds on this to develop our novel PCA^2 methodology; Section 3 provides a numerical example of the empirical application of PCA^2 using data from oil forward markets; Section 4 outlines simulation equations for PCA and PCA^2 variants; Section 5 gives an extended example running through calibration, simulation, and valuation to the derivation of valuation percentiles for risk management purposes; Section 6 lays out the benefits of PCA^2; and finally, Section 7 provides concluding remarks.

2. PCA^2 methodology

2.1. PCA for single and multiple forward curves

To begin, we provide a brief overview of PCA, applying it first in the single curve setting, and then in the multiple curve setting. We consider some underlying asset of interest and assume the returns history for the associated forward curve is captured by the $N \times n$ matrix $\mathbf{H}$. Each row represents a term structure of forward price returns observed over n contract maturities. The contract maturities increase across the term structure, such that the first point on the term structure is the prompt-date forward contract. Each column represents a time series of forward price returns observed over N trade dates, corresponding to a given forward contract maturity. We then compute the $n \times n$ covariance matrix $\mathbf{C}$, which we assume is positive semi-definite. From this we compute the $n \times n$ matrix of eigenvectors $\mathbf{E}$, consisting of the eigenvectors as rows, with the most significant eigenvector at the top, and the eigenvalue matrix $\mathbf{\Lambda}$, with the respective eigenvalues along its diagonal.

From standard PCA theory we know the following:

$$\mathbf{E}\,\mathbf{C} = \mathbf{\Lambda}\,\boldsymbol{E} \tag{1}$$

$$\mathbf{C} = \mathbf{E}^\top\,\mathbf{\Lambda}\,\boldsymbol{E} \tag{2}$$

For our purposes it is useful to introduce the notion of 'scaled eigenvectors' with the definition:

$$\mathbf{S} := \mathbf{\Lambda}^{\frac{1}{2}}\,\mathbf{E} \tag{3}$$

We know that the rows of $\mathbf{E}$ are mutually orthogonal, (hence) linearly independent, and $rank(\mathbf{E}) = n$. By construction these properties are inherited by $\mathbf{S}$.

From Eq. (3) we deduce:

$$\mathbf{S}^\top\,\mathbf{S} = \mathbf{E}^\top\,\mathbf{\Lambda}^{\frac{1}{2}}\,\mathbf{\Lambda}^{\frac{1}{2}}\,\mathbf{E} = \mathbf{C} \tag{4}$$

This shows that the covariance matrix can be recovered from the scaled eigenvectors. We can also deduce:

$$\mathbf{S}\,\mathbf{S}^\top = \mathbf{\Lambda}^{\frac{1}{2}}\mathbf{E}\,\mathbf{E}^\top\,\mathbf{\Lambda}^{\frac{1}{2}} = \mathbf{\Lambda} \tag{5}$$

This confirms that the rows of $\mathbf{S}$ are orthogonal.

With this notation established, we next demonstrate PCA for multiple forward curves. We consider two assets for ease of exposition, but generalising to any number of assets is straightforward. We consider underlying assets a and b, and their associated forward curves. The returns history for each forward curve is assumed be given by an $N \times n$ matrix, which we denote $\mathbf{H}_a$ and $\mathbf{H}_b$ respectively. The rows and columns have the same interpretation as before: rows capturing the term structure of forward prices on a cross sectional basis, and columns capturing the time series of forward prices for the defined contract maturities.

We can now combine these into an $N \times 2n$ block matrix of returns histories, as long as we ensure the observation dates of the rows match:

$$\mathbf{H}_{ab} := [\mathbf{H}_a\,\mathbf{H}_b] \tag{6}$$

The next step is to compute the $2n \times 2n$ covariance matrix $\mathbf{C}_{ab}$. Eigenvector-decomposition of this covariance matrix then gives the $2n \times 2n$ matrix of eigenvectors, $\mathbf{E}_{ab}$, and the $2n \times 2n$ matrix of eigenvalues, $\mathbf{\Lambda}_{ab}$. Leveraging this eigenvector-decomposition we can then define the $2n \times 2n$ matrix of scaled eigenvectors $\mathbf{S}_{ab}$:

$$\mathbf{S}_{ab} := \mathbf{\Lambda}_{ab}^{\frac{1}{2}}\,\mathbf{E}_{ab} \tag{7}$$

We can think of $\mathbf{S}_{ab}$ as a block matrix:

$$\left[\mathbf{S}_{a|ab}\;\;\mathbf{S}_{b|ab}\right] := \mathbf{S}_{ab} \tag{8}$$

The $2n \times n$ left block determines the price behaviour of forward contracts on asset a, and the $2n \times n$ right block determines the price behaviour of forward contracts on asset b.

2.2. Impact of dimension reduction

Utilising scaled eigenvectors allows us to use uncorrelated random shocks within forward curve simulations, which is very convenient. But it also gives us the opportunity to reduce the number of factors used to model the curves.

The eigenvectors have a canonical order based on the amount of variance captured - which is reflected in the declining value of the corresponding eigenvalues. Eq. (4) shows that if we retain all of the computed eigenvectors, then we recover the original covariance matrix exactly. But suppose we reduce the number of principal components used to describe the dynamics of the curve, by selecting the first m components. We can write

$$\mathbf{S} = \begin{bmatrix}\mathbf{U}\\ \mathbf{V}\end{bmatrix} \tag{9}$$

where $\mathbf{U}$ is the $m \times n$ matrix consisting of the selected components, and $\mathbf{V}$ is the $(n-m) \times n$ matrix consisting of the rejected components. As the rows of $\mathbf{S}$ are orthogonal, so are the rows of $\mathbf{U}$; one consequence being that $rank(\mathbf{U}) = m$

We can see that:

$$\mathbf{C} = \mathbf{S}^\top\mathbf{S} = \begin{bmatrix}\mathbf{U}^\top & \mathbf{V}^\top\end{bmatrix}\begin{bmatrix}\mathbf{U}\\ \mathbf{V}\end{bmatrix} = \mathbf{U}^\top\mathbf{U} + \mathbf{V}^\top\mathbf{V} \tag{10}$$

i.e. the original covariance matrix is the sum of the two effective covariance matrices corresponding to $\mathbf{U}$ and $\mathbf{V}$, and the latter constitutes the "error" introduced by reducing $\mathbf{S}$ to $\mathbf{U}$. From Eq. (5) we deduce:

$$\mathbf{\Lambda} = \mathbf{S}\mathbf{S}^\top = \begin{bmatrix}\mathbf{U}\\ \mathbf{V}\end{bmatrix}\begin{bmatrix}\mathbf{U}^\top & \mathbf{V}^\top\end{bmatrix} = \begin{bmatrix}\mathbf{U}\mathbf{U}^\top & \mathbf{U}\mathbf{V}^\top\\ \mathbf{V}\mathbf{U}^\top & \mathbf{V}\mathbf{V}^\top\end{bmatrix} \tag{11}$$

Hence, $\mathbf{U}\mathbf{U}^\top$ and $\mathbf{V}\mathbf{V}^\top$ are diagonal, and $\mathbf{U}\mathbf{V}^\top = \mathbf{V}\mathbf{U}^\top = \mathbf{0}$. Consequently,

$$\mathbf{S}\mathbf{S}^\top = \mathbf{\Lambda}_\mathbf{U} \oplus \mathbf{\Lambda}_\mathbf{V} \tag{12}$$

where $\oplus$ represents the 'direct sum', i.e. the block diagonal matrix with $\mathbf{\Lambda}_\mathbf{U}$ and $\mathbf{\Lambda}_\mathbf{V}$ on the leading diagonal and zeros elsewhere.

2.3. PCA^2 without dimension reduction

Building on the above foundations, we move next to develop the PCA^2 methodology. We demonstrate the two steps of the approach for the case of two underlying assets and no dimension reduction. Extending to any number of assets is straightforward.

First step

We again consider underlying assets a and b and take $N \times n$ forward price returns histories $\mathbf{H}_a$ and $\mathbf{H}_b$. We then compute the $n \times n$ covariance matrices $\mathbf{C}_a$ and $\mathbf{C}_b$ respectively. From these we compute the $n \times n$ matrices of eigenvectors $\mathbf{E}_a$ and $\mathbf{E}_b$, and eigenvalues $\mathbf{\Lambda}_a$ and $\mathbf{\Lambda}_b$, and hence the $n \times n$ matrices of scaled eigenvectors $\mathbf{S}_a$ and $\mathbf{S}_b$. Now in order to combine these to form components across all curves, we first compute the 'component returns histories' for each set, and then compute the covariances between these. Since $\mathbf{S}_a$ and $\mathbf{S}_b$ are invertible, we can translate the returns histories for assets a and b relative to the original risk factors into two $N \times n$ returns histories relative to the scaled eigenvectors $\mathbf{S}_a$ and $\mathbf{S}_b$, as follows:

$$\mathbf{H}_\alpha = \mathbf{H}_a \mathbf{S}_a^{-1}, \quad \mathbf{H}_\beta = \mathbf{H}_b \mathbf{S}_b{}^{-1} \tag{13}$$

The component returns histories can now be combined to form an $N \times 2n$ matrix

$$\mathbf{H}_{\alpha\beta} = \begin{bmatrix} \mathbf{H}_\alpha & \mathbf{H}_\beta \end{bmatrix} \tag{14}$$

We can now compute the $2n \times 2n$ covariance matrix $\mathbf{C}_{\alpha\beta}$. Since the individual sets of principal components are orthogonal, we know that

$$\mathbf{C}_{\alpha\beta} = \begin{bmatrix} \mathbf{I}_n & \mathbf{C}_{\alpha\times\beta} \\ \mathbf{C}_{\alpha\times\beta}^\top & \mathbf{I}_n \end{bmatrix} \tag{15}$$

where $\mathbf{C}_{\alpha\times\beta}$ is composed of the covariances between the two curves' component returns.

Note that, since $\mathbf{S}_a$ and $\mathbf{S}_b$ capture the risk-factor volatilities, the component returns are standardised - so the covariance matrices described above are also correlation matrices.

Second step

PCA is now performed by computing the eigenvectors and eigenvalues of $\mathbf{C}_{\alpha\beta}$, namely the $2n \times 2n$ matrices $\mathbf{E}_{\alpha\beta}$ and $\mathbf{\Lambda}_{\alpha\beta}$ respectively. From these, the corresponding scaled eigenvectors can be computed:

$$\mathbf{S}_{\alpha\beta} = \mathbf{\Lambda}_{\alpha\beta}^{\frac{1}{2}} \mathbf{E}_{\alpha\beta} \tag{16}$$

Component construction

We now have all the ingredients required to form components spanning all the original risk factors:

$$\mathcal{S}_{ab} = \mathbf{S}_{\alpha\beta}(\mathbf{S}_a \oplus \mathbf{S}_b) \tag{17}$$

Since we have arranged the eigenvectors as rows, the left half of $\mathbf{S}_{\alpha\beta}$ relates to asset a and the right half to asset b, so we can think of $\mathbf{S}_{\alpha\beta}$ as a block matrix:

$$\begin{bmatrix} \mathbf{S}_{\alpha|\alpha\beta} & \mathbf{S}_{\beta|\alpha\beta} \end{bmatrix} := \mathbf{S}_{\alpha\beta} \tag{18}$$

Therefore, Eq. (17) can be re-factored:

$$\mathcal{S}_{ab} = \begin{bmatrix} \mathbf{S}_{\alpha|\alpha\beta}\mathbf{S}_a & \mathbf{S}_{\beta|\alpha\beta}\mathbf{S}_b \end{bmatrix} \tag{19}$$

In Eq. (19) above, the $2n \times n$ left block determines the behaviour of asset a, and the $2n \times n$ right block determines the behaviour of asset b.

In Appendix A we prove that without dimension reduction PCA^2 returns the original covariance matrix.

2.4. PCA^2 with dimension reduction at the first step

First step

With dimension reduction, we perform the components partition we saw in Eq. (9) for the underlying assets a and b, specifically:

$$\mathbf{S}_a = \begin{bmatrix} \mathbf{U}_a \\ \mathbf{V}_a \end{bmatrix}, \quad \mathbf{S}_b = \begin{bmatrix} \mathbf{U}_b \\ \mathbf{V}_b \end{bmatrix} \tag{20}$$

where $\mathbf{U}_a$ and $\mathbf{U}_b$ are the $m \times n$ matrices consisting of the selected components, and $\mathbf{V}_a$ and $\mathbf{V}_b$ are the $(n-m) \times n$ matrices consisting of the rejected components. Since $\mathbf{U}_a$ and $\mathbf{U}_b$ are not square, to obtain the corresponding component returns histories, denoted $\mathbf{H}_\gamma$ and $\mathbf{H}_\delta$, we first compute the Moore-Penrose pseudo-inverses. By construction, we know that $rank(\mathbf{U}_a) = rank(\mathbf{U}_b) = m$, hence the pseudo-inverses can be computed directly:

$$\mathbf{U}_a^\dagger = \mathbf{U}_a^\top \left(\mathbf{U}_a \mathbf{U}_a^\top\right)^{-1}, \quad \mathbf{U}_b^\dagger = \mathbf{U}_b^\top \left(\mathbf{U}_b \mathbf{U}_b^\top\right)^{-1} \tag{21}$$

This allows us to translate the returns histories for assets a and b relative to the original risk factors into two $N \times m$ returns histories relative to the components $\mathbf{U}_a$ and $\mathbf{U}_b$ as follows:

$$\mathbf{H}_\gamma = \mathbf{H}_a \mathbf{U}_a^\dagger, \quad \mathbf{H}_\delta = \mathbf{H}_b \mathbf{U}_b^\dagger \tag{22}$$

The component returns histories for the two curves can be combined to give the $N \times 2m$ matrix

$$\mathbf{H}_{\gamma\delta} = \begin{bmatrix} \mathbf{H}_\gamma & \mathbf{H}_\delta \end{bmatrix} \tag{23}$$

We know that the covariance matrix

$$\mathbf{C}_{\gamma\delta} = \begin{bmatrix} \mathbf{C}_\gamma & \mathbf{C}_{\gamma\times\delta} \\ \mathbf{C}_{\gamma\times\delta}^\top & \mathbf{C}_\delta \end{bmatrix} \tag{24}$$

and

$$\mathbf{C}_{\gamma\delta} = \mathbf{H}_{\gamma\delta}^\top \mathbf{H}_{\gamma\delta}/N = \begin{bmatrix} \mathbf{H}_\gamma^\top \\ \mathbf{H}_\delta^\top \end{bmatrix} \begin{bmatrix} \mathbf{H}_\gamma & \mathbf{H}_\delta \end{bmatrix} /N = \begin{bmatrix} \mathbf{H}_\gamma^\top \mathbf{H}_\gamma & \mathbf{H}_\gamma^\top \mathbf{H}_\delta \\ \mathbf{H}_\delta^\top \mathbf{H}_\gamma & \mathbf{H}_\delta^\top \mathbf{H}_\delta \end{bmatrix} /N \tag{25}$$

Re-expressing $\mathbf{C}_\gamma$ in terms of $\mathbf{U}_a$ and $\mathbf{C}_a$ gives

$$\mathbf{C}_\gamma = \mathbf{H}_\gamma^\top \mathbf{H}_\gamma /N = \mathbf{U}_a^{\dagger\top} \mathbf{H}_a^\top \mathbf{H}_a \mathbf{U}_a^\dagger /N = \mathbf{U}_a^{\dagger\top} \mathbf{C}_a \mathbf{U}_a^\dagger \tag{26}$$

Similarly, we can show that for $\mathbf{C}_\delta$:

$$\mathbf{C}_\delta = \mathbf{H}_\delta^\top \mathbf{H}_\delta /N = \mathbf{U}_b^{\dagger\top} \mathbf{H}_b^\top \mathbf{H}_b \mathbf{U}_b^\dagger /N = \mathbf{U}_b^{\dagger\top} \mathbf{C}_b \mathbf{U}_b^\dagger \tag{27}$$

Substituting these expressions (and their analogues) into Eq. (24) then yields:

$$\mathbf{C}_{\gamma\delta} = \begin{bmatrix} \mathbf{U}_a^{\dagger\top} \mathbf{C}_a \mathbf{U}_a^\dagger & \mathbf{U}_a^{\dagger\top} \mathbf{C}_{a\times b} \mathbf{U}_b^\dagger \\ \mathbf{U}_b^{\dagger\top} \mathbf{C}_{a\times b}^\top \mathbf{U}_a^\dagger & \mathbf{U}_b^{\dagger\top} \mathbf{C}_b \mathbf{U}_b^\dagger \end{bmatrix} = \left(\mathbf{U}_a^{\dagger\top} \oplus \mathbf{U}_b^{\dagger\top}\right) \mathbf{C}_{ab} \left(\mathbf{U}_a^\dagger \oplus \mathbf{U}_b^\dagger\right) \tag{28}$$

This completes the first step.

Note that, as an analogue of what we saw in Section 2.3 above, since $\mathbf{U}_a$ and $\mathbf{U}_b$ capture the risk-factor volatilities, the component returns are standardised - so the covariance matrices described above are also correlation matrices.

Second step

The second step of PCA can now be performed by computing the eigenvectors and eigenvalues of $\mathbf{C}_{\gamma\delta}$, namely the $2m \times 2m$ matrices $\mathbf{E}_{\gamma\delta}$ and $\mathbf{\Lambda}_{\gamma\delta}$ respectively. From these, the corresponding scaled eigenvectors can be computed:

$$\mathbf{S}_{\gamma\delta} = \mathbf{\Lambda}_{\gamma\delta}^{\frac{1}{2}} \mathbf{E}_{\gamma\delta} \tag{29}$$

We know that the rows of $\mathbf{E}_{\gamma\delta}$ are mutually orthogonal, (hence) linearly independent, and $rank(\mathbf{E}_{\gamma\delta}) = 2m$. By construction these

Table 1
Covariance of daily log-returns × 1000.

Covariance	CO1	CO6	CO12	CO18	CO24	QS1	QS6	QS12	QS18	QS24
CO1	0.526	0.403	0.335	0.290	0.254	0.304	0.266	0.226	0.197	0.176
CO6	0.403	0.350	0.300	0.263	0.233	0.247	0.223	0.194	0.171	0.153
CO12	0.335	0.300	0.261	0.232	0.208	0.207	0.189	0.166	0.148	0.133
CO18	0.290	0.263	0.232	0.208	0.188	0.180	0.164	0.145	0.131	0.118
CO24	0.254	0.233	0.208	0.188	0.172	0.157	0.144	0.129	0.116	0.106
QS1	0.304	0.247	0.207	0.180	0.157	0.371	0.284	0.238	0.207	0.183
QS6	0.266	0.223	0.189	0.164	0.144	0.284	0.256	0.219	0.192	0.170
QS12	0.226	0.194	0.166	0.145	0.129	0.238	0.219	0.192	0.170	0.151
QS18	0.197	0.171	0.148	0.131	0.116	0.207	0.192	0.170	0.153	0.137
QS24	0.176	0.153	0.133	0.118	0.106	0.183	0.170	0.151	0.137	0.128

This table reports the covariance matrix of daily log-returns derived from Brent Crude (tickers CO1, CO6, CO12, CO18 and CO24) and Low Sulphur Gasoil (tickers QS1, QS6, QS12, QS18 and QS24) futures contracts, spanning the first-, sixth-, twelfth-, eighteenth- and twenty-fourth-month maturities. The sample period is 26 September 2011 to 24 September 2021.

properties are inherited by $\mathbf{S}_{\gamma\delta}$. The left half of $\mathbf{S}_{\gamma\delta}$ relates to asset a and the right half to asset b, so we can think of $\mathbf{S}_{\gamma\delta}$ as a block matrix:

$$\begin{bmatrix}\mathbf{S}_{\gamma|\gamma\delta} & \mathbf{S}_{\delta|\gamma\delta}\end{bmatrix} := \mathbf{S}_{\gamma\delta} \tag{30}$$

Component construction

We now have all the ingredients required to form new components spanning all the original risk factors, but based on an abbreviated set of eigenvectors from the first step of PCA:

$$\mathbb{S}_{ab} = \begin{bmatrix}\mathbf{S}_{\gamma|\gamma\delta}\mathbf{U}_a & \mathbf{S}_{\delta|\gamma\delta}\mathbf{U}_b\end{bmatrix} = \mathbf{S}_{\gamma\delta}(\mathbf{U}_a \oplus \mathbf{U}_b) \tag{31}$$

By construction we know that $rank(\mathbb{S}_{ab}) = 2m$ and its rows are linearly independent. The left $2m \times n$ block of $\mathbb{S}_{ab}$ determines the behaviour of asset a, and the right $2m \times n$ block determines the behaviour of asset b.

The effective covariance matrix

The effective covariance matrix corresponding to $\mathbb{S}_{ab}$ is:

$$\begin{aligned}\mathbb{C}_{ab} &= \mathbb{S}_{ab}^\top \mathbb{S}_{ab} = \left(\mathbf{U}_a^\top \oplus \mathbf{U}_b^\top\right)\mathbf{S}_{\gamma\delta}^\top \mathbf{S}_{\gamma\delta}\left(\mathbf{U}_a \oplus \mathbf{U}_b\right)\\ &= \left(\mathbf{U}_a^\top \oplus \mathbf{U}_b^\top\right)C_{\gamma\delta}\left(\mathbf{U}_a \oplus \mathbf{U}_b\right)\\ &= \left(\mathbf{U}_a^\top \oplus \mathbf{U}_b^\top\right)\left(\mathbf{U}_a^{\dagger\top} \oplus \mathbf{U}_b^{\dagger\top}\right)\mathbf{C}_{ab}\left(\mathbf{U}_a^\dagger \oplus \mathbf{U}_b^\dagger\right)\left(\mathbf{U}_a \oplus \mathbf{U}_b\right)\end{aligned} \tag{32}$$

Equivalently,

$$\begin{cases}\mathbb{C}_{ab} = \mathbf{J}_{ab}^\top \mathbf{C}_{ab}\mathbf{J}_{ab}\\ \text{where } \mathbf{J}_{ab} = \left(\mathbf{U}_a^\dagger\mathbf{U}_a \oplus \mathbf{U}_b^\dagger\mathbf{U}_b\right)\end{cases} \tag{33}$$

Appendix B outlines how Eq. (33) can be simplified to

$$\begin{cases}\mathbb{C}_{ab} = \mathbf{C}_{ab}\mathbf{J}_{ab}\\ \text{where } \mathbf{J}_{ab} = \left(\mathbf{U}_a^\dagger\mathbf{U}_a \oplus \mathbf{U}_b^\dagger\mathbf{U}_b\right)\end{cases} \tag{34}$$

The analogues of Eq. (12) also apply:

$$\mathbf{S}_a\mathbf{S}_a^\top = \mathbf{\Lambda}_a = \mathbf{\Lambda}_{U_a} \oplus \mathbf{\Lambda}_{V_a}, \quad \mathbf{S}_b\mathbf{S}_b^\top = \mathbf{\Lambda}_b = \mathbf{\Lambda}_{U_b} \oplus \mathbf{\Lambda}_{V_b} \tag{35}$$

Consequently,

$$\begin{aligned}\mathbb{S}_{ab}\mathbb{S}_{ab}^\top &= \mathbf{S}_{\gamma\delta}\left(\mathbf{U}_a \oplus \mathbf{U}_b\right)\left(\mathbf{U}_a^\top \oplus \mathbf{U}_b^\top\right)\mathbf{S}_{\gamma\delta}^\top\\ &= \mathbf{S}_{\gamma\delta}\left(\mathbf{U}_a\mathbf{U}_a^\top \oplus \mathbf{U}_b\mathbf{U}_b^\top\right)\mathbf{S}_{\gamma\delta}^\top\\ &= \mathbf{S}_{\gamma\delta}\left(\mathbf{\Lambda}_{U_b} \oplus \mathbf{\Lambda}_{V_b}\right)\mathbf{S}_{\gamma\delta}^\top\end{aligned} \tag{36}$$

This shows that the rows of $\mathbb{S}_{ab}$ (i.e. our PCA^2 components) are not generally orthogonal, although we showed above that they are nonetheless linearly independent.

In principle, dimension reduction can be applied at the second step of PCA, but in our practical applications of the PCA^2 method we have avoided this, therefore further details are relegated to Appendix C.

Remark on the generality of n and m

To close our discussion, it should be noted that the choices of n and m can be made on a curve-by-curve basis because all of the reasoning (and equations) continues to hold in cases where curves for assets a and b have different n's and m's. For example, let us assume we have n_a risk factors for curve a and n_b risk factors for curve b, and N days of returns data for both. Then $\mathbf{H}_a$ and $\mathbf{H}_b$ are $N \times n_a$ and $N \times n_b$ matrices respectively. Hence, $\mathbf{S}_a$ and $\mathbf{S}_b$ are $n_a \times n_a$ and $n_b \times n_b$ matrices respectively. If we select m_a first-step components for curve a and m_b first-step components for curve b, then $\mathbf{U}_a$ and $\mathbf{U}_b$ are $m_a \times n_a$ and $m_b \times n_b$ matrices respectively. Consequently, $\mathbf{U}_a^\dagger$ and $\mathbf{U}_b^\dagger$ will be $n_a \times m_a$ and $n_b \times m_b$ matrices, and $\mathbf{H}_\gamma$ and $\mathbf{H}_\delta$ will be $N \times m_a$ and $N \times m_b$ matrices respectively. The second step of PCA will take the $N \times (m_a + m_b)$ matrix $\mathbf{H}_{\gamma\delta}$ as input and produce the $\mathbf{C}_{\gamma\delta}$, $\mathbf{E}_{\gamma\delta}$, $\mathbf{\Lambda}_{\gamma\delta}$ and $\mathbf{S}_{\gamma\delta}$ matrices, which are all $(m_a + m_b) \times (m_a + m_b)$. From Eq. (31) we can now see that since $\mathbb{S}_{ab} = \mathbf{S}_{\gamma\delta}(\mathbf{U}_a \oplus \mathbf{U}_b)$, and $\mathbf{U}_a \oplus \mathbf{U}_b$ is an $(m_a + m_b) \times (n_a + n_b)$ matrix, $\mathbb{S}_{ab}$ will be the $(m_a + m_b) \times (n_a + n_b)$ matrix in which the left $(m_a + m_b) \times n_a$ block, $\mathbf{S}_{\gamma|\gamma\delta}\mathbf{U}_a$, relates to curve a and the right $(m_a + m_b) \times n_b$ block, $\mathbf{S}_{\delta|\gamma\delta}\mathbf{U}_b$, relates to curve b.

3. Numerical demonstration of PCA^2

We now proceed to demonstrate the PCA^2 methodology with a numerical example based on oil futures market data. We seek to be as explicit as possible in demonstrating the main stages as outlined in Section 2, ensuring readers can fully replicate the results and readily implement the methodology for their own purposes. We make reference to the relevant equation numbering as required to aid the cross-referencing. To this end, we extract Bloomberg price data for Brent Crude (tickers CO1, CO6, CO12, CO18 and CO24) and Low Sulphur Gasoil (tickers QS1, QS6, QS12, QS18 and QS24) futures contracts, spanning the first-, sixth-, twelfth-, eighteenth- and twenty-fourth-month maturities. We compute daily log-returns over the ten-year period from 26 September 2011 to 24 September 2021.

In describing the first step, Table 1 shows the covariances of log-returns for the combined curve setting, and Table 2 shows the derived risk factor volatilities. The scaled eigenvectors (see

Table 2
Volatilities of daily log-returns.

Risk Factor	$CO1$	$CO6$	$CO12$	$CO18$	$CO24$	$QS1$	$QS6$	$QS12$	$QS18$	$QS24$
Volatility	0.0229	0.0187	0.0162	0.0144	0.0131	0.0192	0.016	0.0139	0.0124	0.0113

This table reports the volatilities of daily log-returns derived from Brent Crude (tickers CO1, CO6, CO12, CO18 and CO24) and Low Sulphur Gasoil (tickers QS1, QS6, QS12, QS18 and QS24) futures contracts, spanning the first-, sixth-, twelfth-, eighteenth- and twenty-fourth-month maturities. The sample period is 26 September 2011 to 24 September 2021.

Table 3
Scaled eigenvectors for Brent Crude $\times 1000$.

$\mathbf{S}_a$	$CO1$	$CO6$	$CO12$	$CO18$	$CO24$
$PCa1$	22.005	18.606	15.965	14.074	12.536
$PCa2$	6.446	−0.770	−2.425	−3.135	−3.565
$PCa3$	0.685	−1.772	−0.561	0.577	1.495
$PCa4$	0.034	−0.333	0.428	0.186	−0.321
$PCa5$	0.001	−0.029	0.137	−0.204	0.095

This table reports the scaled eigenvectors derived according to Eq. (20) from Brent Crude (tickers CO1, CO6, CO12, CO18 and CO24) futures contracts, spanning the first-, sixth-, twelfth-, eighteenth- and twenty-fourth-month maturities. The five principal components are labelled accordingly $PCa1 - PCa5$, where the inclusion of the letter 'a' in the labelling emphasises our designation of Brent Crude as the generic asset 'a' defined in Section 2.

Table 4
Scaled eigenvectors for Low Sulphur Gasoil $\times 1000$.

$\mathbf{S}_b$	$QS1$	$QS6$	$QS12$	$QS18$	$QS24$
$PCb1$	18.372	15.827	13.662	12.050	10.771
$PCb2$	5.715	−1.020	−2.136	−2.515	−2.725
$PCb3$	0.594	−1.873	−0.661	0.568	1.941
$PCb4$	0.086	−0.649	0.373	1.040	−0.829
$PCb5$	0.033	−0.378	0.795	−0.491	0.040

This table reports the scaled eigenvectors derived according to Eq. (20) from Low Sulphur Gasoil (tickers QS1, QS6, QS12, QS18 and QS24) futures contracts, spanning the first-, sixth-, twelfth-, eighteenth- and twenty-fourth-month maturities. The five principal components are labelled accordingly $PCb1 - PCb5$, where the inclusion of the letter 'b' in the labelling emphasises our designation of Low Sulphur Gasoil as the generic asset 'b' defined in Section 2.

Table 5
Moore-Penrose pseudo-inverses.

$\mathbf{U}_a^\dagger$	$PCa1$	$PCa2$	$\mathbf{U}_b^\dagger$	$PCb1$	$PCb2$
$CO1$	15.276	91.354	$QS1$	17.735	109.865
$CO6$	12.916	−10.911	$QS6$	15.279	−19.616
$CO12$	11.083	−34.361	$QS12$	13.189	−41.069
$CO18$	9.770	−44.425	$QS18$	11.633	−48.349
$CO24$	8.702	−50.530	$QS24$	10.398	−52.385

This table reports the Moore-Penrose pseudo-inverses calculated as per Eq. (21) and that are subsequently used to compute the principal component returns histories $\mathbf{H}_\gamma$ and $\mathbf{H}_\delta$ as per Eq. (22). Brent Crude is given by the tickers CO1, CO6, CO12, CO18 and CO24, and Low Sulphur Gasoil is given by the tickers QS1, QS6, QS12, QS18 and QS24, corresponding to the futures contracts spanning the first-, sixth-, twelfth-, eighteenth- and twenty-fourth-month maturities. Only two principal components are retained for both Brent Crude and Low Sulphur Gasoil, where $PCa1 - PCa2$ and $PCb1 - PCb2$ are the principal components as noted in Tables 3 and 4.

Eq. (20)) for Brent Crude ($\mathbf{S}_a$) and Low Sulphur Gasoil ($\mathbf{S}_b$) can then be derived and these are reported in Tables 3 and 4. Table 5 shows the Moore-Penrose pseudo-inverses (see Eq. (21)) used to compute the principal component returns histories $\mathbf{H}_\gamma$ and $\mathbf{H}_\delta$ (see Eq. (22)), where we retain only two principal components for Brent Crude and two principal components for Low Sulphur Gasoil. The covariances of the principal component returns histories (see Eq. (24)) are shown in Table 6. The result of the second step of PCA is shown in Table 7, which gives the block matrix $\mathbf{S}_{\gamma\delta}$ (see Eq. (30)). The final PCA^2 components spanning the original risk factors, $\mathbb{S}_{ab}$ (see Eq. (31)), are shown in Table 8. Table 9 shows the effective co-

Table 6
Principal component covariances.

$\mathbf{C}_{\gamma\delta}$	$PCa1$	$PCa2$	$PCb1$	$PCb2$
$PCa1$	1.000	0.000	0.745	−0.075
$PCa2$	0.000	1.000	0.031	0.285
$PCb1$	0.745	0.031	1.000	0.000
$PCb2$	−0.075	0.285	0.000	1.000

This table reports the covariances of the principal component returns histories as per Eq. (24). $PCa1 - PCa2$ and $PCb1 - PCb2$ are respectively the principal components linked to Brent Crude and Low Sulphur Gasoil as noted in Tables 3 and 4.

Table 7
Second step of PCA^2.

$\mathbf{S}_{\gamma\delta}$	$PCa1$	$PCa2$	$PCb1$	$PCb2$
$PC1$	0.935	0.003	0.930	−0.093
$PC2$	0.003	−0.802	−0.080	−0.798
$PC3$	0.002	−0.597	0.059	0.594
$PC4$	0.354	0.001	−0.353	0.035

This table reports the result of the second step of PCA, which derives the block matrix $\mathbf{S}_{\gamma\delta}$ as per Eq. (30). $PCa1 - PCa2$ and $PCb1 - PCb2$ are respectively the principal components linked to Brent Crude and Low Sulphur Gasoil as noted in Tables 3 and 4, and taken from the first step of PCA^2. $PC1 - PC4$ are the principal components derived from the second step of PCA^2.

variance matrix $\mathbb{C}_{ab}$ as per Eq. (34). Comparison of Table 1 and Table 9 shows that the PCA^2 components capture the original covariance structure extremely well - the differences are shown in Table 10.

Remark on orthogonality

It is a well-known property of principal components that they are orthogonal, and this is often stated to be essential for their use in simulation using uncorrelated random variates. Table 11 shows the scalar product between the PCA^2 components, and the translation of that into the angle between them (in degrees). This table clearly shows that these components are not orthogonal. Orthogonality of components is valued because it guarantees linear independence, but linear independence is actually sufficient without orthogonality. We saw in Section 2.4 that the PCA^2 components are linearly independent, consequently they can be used in exactly the same way as one-step principal components within a Monte Carlo simulation framework - which we show in the simulation section below.

4. PCA^2 application to forward curve simulation

In this section we leverage the methods derived above to develop the equations required to simulate movements of one or more forward curves. We use an HJM-based multifactor model similar to that summarised by Blanco et al. (2002), however, in order to align with earlier sections of this paper we outline a matrix-based representation.

To aid the exposition, Fig. 1 provides a schematic of the forward curve simulation. Each cell represents the price F_{ij} for the jth absolute forward contract at simulation time-step i. We as-

Table 8
Final PCA^2 components $\times 1000$.

$\mathbb{S}_{ab}$	CO1	CO6	CO12	CO18	CO24	QS1	QS6	QS12	QS18	QS24
PC1	20.599	17.395	14.920	13.149	11.710	16.563	14.821	12.910	11.445	10.275
PC2	−5.104	0.673	1.992	2.556	2.897	−6.025	−0.447	0.616	1.047	1.317
PC3	−3.801	0.501	1.484	1.904	2.157	4.487	0.333	−0.459	−0.780	−0.980
PC4	7.809	6.595	5.656	4.985	4.439	−6.279	−5.619	−4.894	−4.339	−3.895

This table reports the final PCA^2 components spanning the original risk factors, $\mathbb{S}_{ab}$ as per Eq. (31). $PC1 - PC4$ are the principal components derived from the second step of PCA^2, as noted in Table 7. Brent Crude is given by the tickers CO1, CO6, CO12, CO18 and CO24, and Low Sulphur Gasoil is given by the tickers QS1, QS6, QS12, QS18 and QS24, corresponding to the futures contracts spanning the first-, sixth-, twelfth-, eighteenth- and twenty-fourth-month maturities.

Table 9
Covariances recovered from PCA^2 components $\times 1000$.

$\mathbf{C}_{ab}$	CO1	CO6	CO12	CO18	CO24	QS1	QS6	QS12	QS18	QS24
CO1	0.526	0.404	0.336	0.289	0.253	0.306	0.262	0.226	0.199	0.178
CO6	0.404	0.347	0.299	0.264	0.236	0.245	0.221	0.192	0.171	0.153
CO12	0.336	0.299	0.261	0.232	0.209	0.206	0.189	0.165	0.147	0.132
CO18	0.289	0.264	0.232	0.208	0.188	0.180	0.166	0.146	0.130	0.117
CO24	0.253	0.236	0.209	0.188	0.170	0.158	0.148	0.130	0.116	0.105
QS1	0.306	0.245	0.206	0.180	0.158	0.370	0.285	0.239	0.207	0.182
QS6	0.262	0.221	0.189	0.166	0.148	0.285	0.252	0.218	0.193	0.173
QS12	0.226	0.192	0.165	0.146	0.130	0.239	0.218	0.191	0.170	0.153
QS18	0.199	0.171	0.147	0.130	0.116	0.207	0.193	0.170	0.152	0.137
QS24	0.178	0.153	0.132	0.117	0.105	0.182	0.173	0.153	0.137	0.123

This table reports the effective covariance matrix $\mathbb{C}_{ab}$ as per Eq. (34). Brent Crude is given by the tickers CO1, CO6, CO12, CO18 and CO24, and Low Sulphur Gasoil is given by the tickers QS1, QS6, QS12, QS18 and QS24, corresponding to the futures contracts spanning the first-, sixth-, twelfth-, eighteenth- and twenty-fourth-month maturities.

Table 10
Covariance differences $\times 1000$.

	CO1	CO6	CO12	CO18	CO24	QS1	QS6	QS12	QS18	QS24
CO1	0.000	0.001	0.000	0.000	−0.001	0.002	−0.003	0.000	0.002	0.003
CO6	0.001	−0.003	−0.001	0.001	0.003	−0.002	−0.002	−0.001	0.000	0.000
CO12	0.000	−0.001	−0.001	0.000	0.001	−0.001	0.000	0.000	−0.001	−0.001
CO18	0.000	0.001	0.000	0.000	−0.001	0.000	0.002	0.001	−0.001	−0.001
CO24	−0.001	0.003	0.001	−0.001	−0.002	0.001	0.004	0.002	0.000	−0.001
QS1	0.002	−0.002	−0.001	0.000	0.001	0.000	0.001	0.000	0.000	−0.001
QS6	−0.003	−0.002	0.000	0.002	0.004	0.001	−0.004	−0.001	0.002	0.003
QS12	0.000	−0.001	0.000	0.001	0.002	0.000	−0.001	−0.001	0.000	0.002
QS18	0.002	0.000	−0.001	−0.001	0.000	0.000	0.002	0.000	−0.002	0.000
QS24	0.003	0.000	−0.001	−0.001	−0.001	−0.001	0.003	0.002	0.000	−0.004

This table reports the differences between the original covariance matrix set out in Table 1 and the covariance matrix recovered from the PCA^2 methodology, as set out in Table 9. Brent Crude is given by the tickers CO1, CO6, CO12, CO18 and CO24, and Low Sulphur Gasoil is given by the tickers QS1, QS6, QS12, QS18 and QS24, corresponding to the futures contracts spanning the first-, sixth-, twelfth-, eighteenth- and twenty-fourth-month maturities.

Table 11
Non-orthogonality of PCA^2 components.

Scalar Product	PC1	PC2	PC3	PC4	Degrees	PC1	PC2	PC3	PC4
PC1	1.000	−0.161	0.179	0.168	PC1	0	99	80	80
PC2	−0.161	1.000	0.067	0.179	PC2	99	0	86	80
PC3	0.179	0.067	1.000	−0.161	PC3	80	86	0	99
PC4	0.168	0.179	−0.161	1.000	PC4	80	80	99	0

This table reports the scalar product between the PCA^2 components, and the translation of that into the angle between them (to the nearest degree). $PC1 - PC4$ are the principal components derived from the second step of PCA^2, as noted in Table 7.

sume that the time between steps matches the time period between contracts, so the dark grey cells correspond to where absolute contracts have expired. The labelled light grey cells represent valid prices. The arrow that is shown represents the derivation of, as an example, F_{23} from F_{13} via the application of a perturbation factor Q_{11}, which we describe below. All forward prices can be easily derived from such a Q-perturbation, and, indeed, it is this Q-perturbation that allows us to readily transition from a PCA setting to a PCA^2 setting.

The rules for the derivation of F_{ij} are:

$$F_{ij} = \begin{cases} F_{0j}, & \text{if } i = 0 \\ \text{NaN}, & \text{if } i > j \\ F_{i-1,j} \cdot Q_{i-1,j-i}, & \text{otherwise} \end{cases} \tag{37}$$

where F_{0j} is the initial price of the jth absolute forward contract and Q is the appropriate perturbation matrix derived from the principal components. We describe the specification of Q for PCA first and then PCA^2. Note that NaN ("Not-a-Number") is used to de-

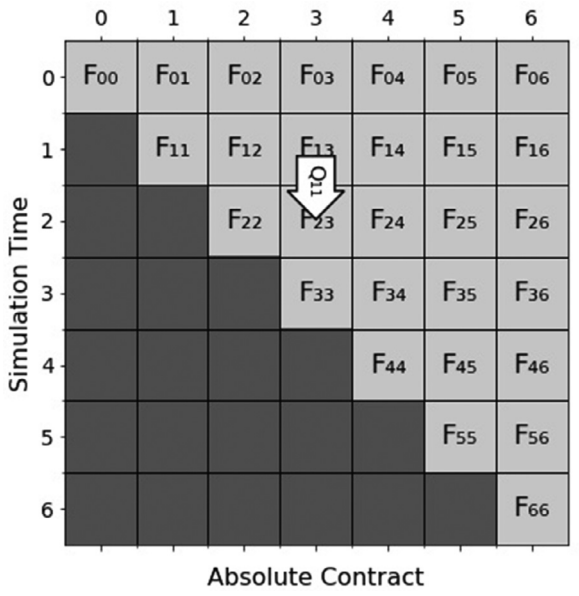

Fig. 1. Price simulation matrix. This figure provides a schematic of the forward curve simulation as described by Eq. (37). Each cell represents the price F_{ij} for the jth absolute forward contract at simulation time-step i. The time between steps matches the time period between contracts, so the dark grey cells correspond to where absolute contracts have expired. The labelled light grey cells represent valid prices. The arrow that is shown represents the derivation of, as an example, F_{23} from F_{13} via the application of a perturbation factor Q_{11}, in line with Eq. (37). The Q-perturbation may be defined appropriate depending on whether a PCA implementation or a PCA^2 implementation is required, as described later in this section.

note no price, corresponding to where the absolute contracts have expired.

PCA for single curve simulation

For the PCA case *without* dimension reduction, it follows from Section 2 that Q is given by:

$$Q = \exp\left[\sqrt{\Delta t}\,\boldsymbol{\varepsilon}\,\mathbf{S} - \frac{1}{2}\Delta t\,\mathbf{1}\,\mathbf{S}^2\right] \tag{38}$$

where $\mathbf{S}$ is the $n \times n$ matrix of scaled eigenvalues from Eq. (3), $\boldsymbol{\varepsilon}$ is an $n \times n$ matrix of uncorrelated standard normal variates, $\mathbf{1}$ is the $n \times n$ matrix consisting entirely of 1's, and $(\cdot)^2$ denotes pointwise squaring. Δt is some defined time step for the simulation. In a similar fashion, for the PCA case *with* dimension reduction, Q is given by:

$$Q = \exp\left[\sqrt{\Delta t}\,\boldsymbol{\varepsilon}\,\mathbf{U} - \frac{1}{2}\Delta t\,\mathbf{1}\,\mathbf{U}^2\right] \tag{39}$$

where $\mathbf{U}$ is the $m \times n$ matrix of scaled eigenvalues from Eq. (9), $\boldsymbol{\varepsilon}$ is an $n \times m$ matrix of uncorrelated standard normal variates, $\mathbf{1}$ is the $n \times m$ matrix consisting entirely of 1's, and $(\cdot)^2$ denotes pointwise squaring.

PCA for multiple curve simulation

Moving to the simulation of multiple forward curves is straightforward. For the case of *two* curves, we have a pair of analogues of Eq. (37):

$$F^a_{ij} = \begin{cases} F^a_{0j}, & \text{if } i = 0 \\ \text{NaN}, & \text{if } i > j \\ F^a_{i-1,j} \cdot Q^a_{i-1,j-i}, & \text{otherwise} \end{cases}$$
$$F^b_{ij} = \begin{cases} F^b_{0j}, & \text{if } i = 0 \\ \text{NaN}, & \text{if } i > j \\ F^b_{i-1,j} \cdot Q^b_{i-1,j-i}, & \text{otherwise} \end{cases} \tag{40}$$

where F^a_{0j} and F^b_{0j} are the initial prices for jth absolute forward contracts for some assets a and b, and $[Q^a\,Q^b] = Q$ is the perturbation matrix derived from the principal components. The "otherwise" cases create the link between the simulations of a and the simulations of b via the multiplication by Q, which applies the same ε's to both curves. It can be shown for the case without dimension reduction that Q is given by:

$$[Q^a\,Q^b] = Q = \exp\left[\sqrt{\Delta t}\,\boldsymbol{\varepsilon}\,\mathbf{S}_{ab} - \frac{1}{2}\Delta t\,\mathbf{1}\,\mathbf{S}^2_{ab}\right] \tag{41}$$

Here $\mathbf{S}_{ab}$ is the $2n \times 2n$ matrix of scaled eigenvalues from Eq. (7), while $\boldsymbol{\varepsilon}$ is a $2n \times 2n$ matrix of uncorrelated standard normal variates, $\mathbf{1}$ is the $2n \times 2n$ matrix consisting entirely of 1's.

PCA^2 for multiple curve simulation

For the PCA^2 case, a similar derivation follows. For ease of exposition, we again consider two curves associated with some assets a and b. We consider the case of dimension reduction here, where we move from n to m components at the first stage. It can be shown that Q is given by:

$$[Q^a\,Q^b] = Q = \exp\left[\sqrt{\Delta t}\,\boldsymbol{\varepsilon}\,\mathbb{S}_{ab} - \frac{1}{2}\Delta t\,\mathbf{1}\,\mathbb{S}^2_{ab}\right] \tag{42}$$

where $\mathbb{S}_{ab}$ is the $2m \times 2n$ matrix of *PCA*^{2} components from Eq. (31), and $\boldsymbol{\varepsilon}$ is a $2n \times 2m$ matrix of uncorrelated standard normal variates, 1 is the $2n \times 2m$ matrix consisting entirely of 1's.

5. PCA^2 application to risk management

In this section, we showcase an extended application of the PCA^2 methodology, benchmarking against a PCA approach. We choose the practice relevant problem of risk management, and work through the following key sequential stages: (1) calibration of PCA^2 (and PCA) parameters from the raw data, (2) simulation of the forward price curves, (3) valuation of an example portfolio under these simulated price curves, and (4) derivation of valuation percentiles for risk management purposes. For this extended application, we enrich our Bloomberg data by including the tickers for intermediate contracts (i.e. CO2-CO5, CO7-CO11, CO13-CO17, and CO19-CO23 for Brent Crude, and QS2-QS5, QS7-QS11, QS13-QS17, and QS19-QS23 for Low Sulphur Gasoil).

5.1. Calibration

Using the full price dataset, we compute the PCA^2 components using only the first three components of Brent Crude and the first three components of Low Sulphur Gasoil at the first step. The left-hand side of Fig. 2 plots the resultant six PCA^2 components. The left half of each component curve corresponds to Brent Crude and the right half to Low Sulphur Gasoil. As a comparison we have also computed the PCA components using only the first six components, shown on the right-hand side of Fig. 2. In this example the relatively high degree of correlation between the two curves has led to S1 and PC1 both reflecting the shared shift of price levels across both curves. However, generally multi-curve components differ and are not open to straightforward interpretation.

Notwithstanding these differences in components, the efficacy of PCA^2 can be established through a comparison of the correlation structure captured by the PCA^2 components with the original correlation structure. These are shown graphically in Fig. 3, where it is evident that the six PCA^2 components are capable of capturing extremely well the correlations between the 48 risk factors.

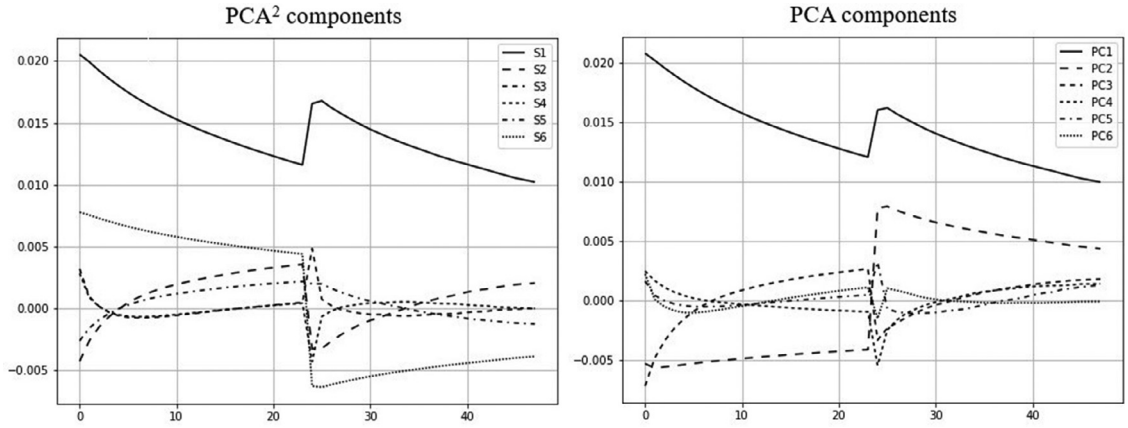

Fig. 2. Comparison of components. This figure provides (i) a plot (left panel) of the six principal components derived under the PCA^2 methodology (where the first three principal components for each of Brent Crude and Low Sulphur Gasoil are used at the first step) and (ii) a plot (right panel) of the first six principal components derived under the standard PCA methodology. The data used to generate these plots is the full 24-month Brent Crude forward curve (CO1-CO24) and the full 24-month Low Sulphur Gasoil forward curve QS1-QS24. The sample period is 26 September 2011 to 24 September 2021. The left half of each component curve corresponds to Brent Crude and the right half to Low Sulphur Gasoil.

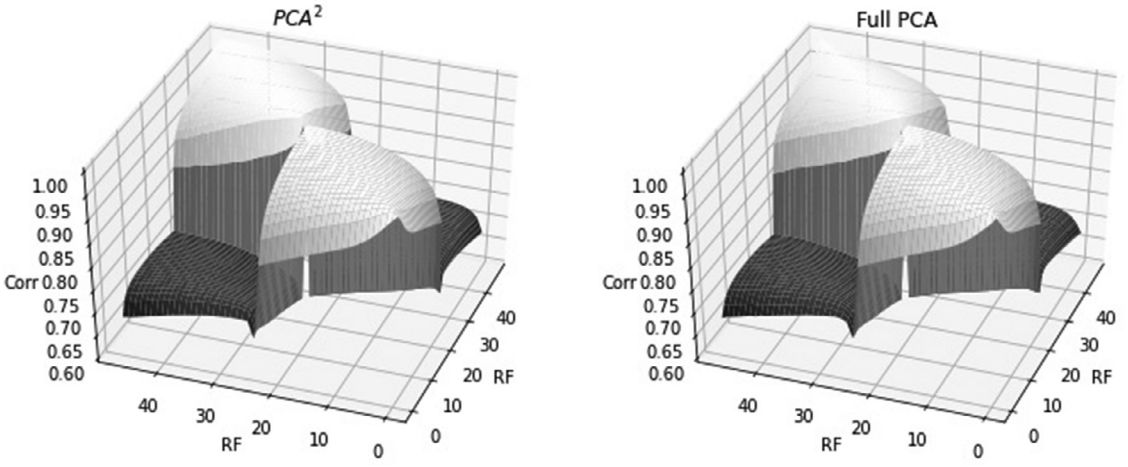

Fig. 3. Correlation structure comparison. This figure provides (i) a plot (left panel) of correlation structure captured by the six principal components based PCA^2 methodology (where the first three principal components for each of Brent Crude and Low Sulphur Gasoil are used at the first step) and (ii) a plot (right panel) of the correlation structure captured by the full 48-risk factor PCA methodology. The data used to generate these plots is the full 24-month Brent Crude forward curve (CO1-CO24) and the full 24-month Low Sulphur Gasoil forward curve QS1-QS24. The sample period is 26 September 2011 to 24 September 2021. The left half of each component curve corresponds to Brent Crude and the right half to Low Sulphur Gasoil.

5.2. Simulation

Using the calibrations above we can now compute price paths for Brent Crude and Low Sulphur Gasoil under the PCA^2 framework, using Eq. (40) and the required Q-perturbation Eq. (42). A matching pair of paths is given as an example in Figs. 4 and 5. Each path begins with the actual price curve at time 0, shown in solid back, followed by a series of compounding perturbations, shown in grey. Each curve is obtained from it's predecessor by dropping the nearest contract (which expires) and perturbing the rest.

5.3. Valuation

The PCA^2 based simulation set out in the previous section can be used in a number of different ways for risk management purposes, but, as an example, we consider a portfolio of forward trades of Brent Crude and Low Sulphur Gasoil, and value this portfolio across 500,000 simulations. Figure 6 summarises the Brent Crude trades, and Fig. 7 summarises the Low Sulphur Gasoil trades. We design a scenario where we have net long positions in Brent Crude (i.e. we have bought more than we have sold), but in Gasoil we have net long positions up to month 14 then net short positions (i.e. we have sold more that we have bought) from month 15 onwards.

The 'Average Trade Price' curve (bold line in Figs. 6 and 7) represents the trade-quantity-weighted average of the trade prices by forward month, and the 'Initial Market Price' curve (dotted line in Figs. 6 and 7) represents the market prices for the forward months as of time zero.

Because we only consider linear instruments in this illustrative example, portfolio valuation at each simulation time-step consists of multiplying each remaining 'Net Trade Quantity' by the difference between the 'Average Trade Price' and the simulated market price and determining the Net Present Value (NPV) of these as of the simulation time. (For simplicity we assume a zero interest rate for the NPV discounting in our example.) Fig. 8 plots the valuations of the Brent Crude and Low Sulphur Gasoil sub-portfolios using simulated price paths of the form shown in Figs. 4 and 5 respectively, together with the overall net valuation. Fig. 9 shows 50 of the 500,000 portfolio valuation paths to illustrate the level of variation within the full dataset.

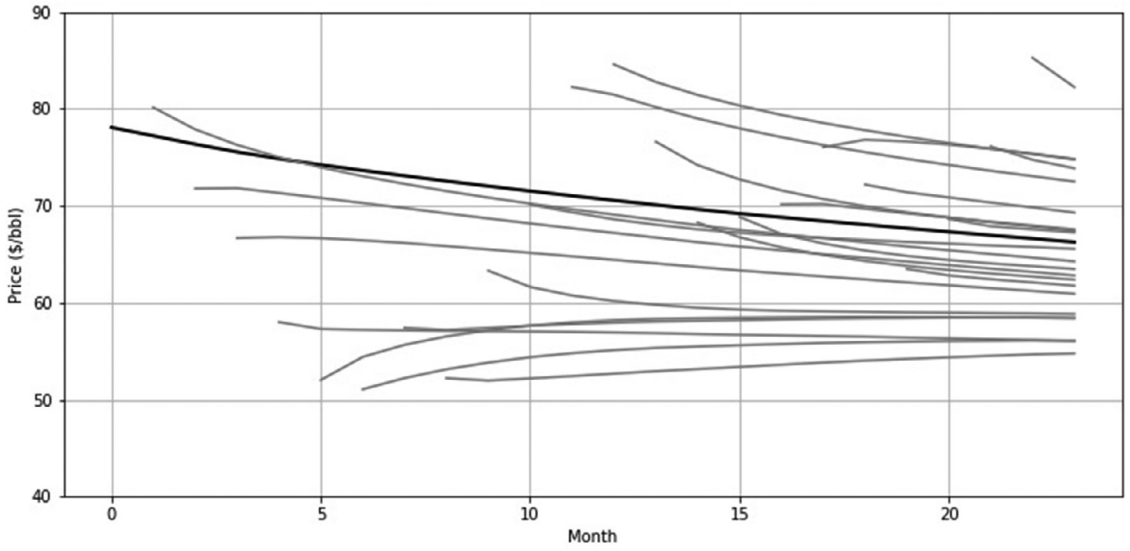

Fig. 4. Brent Crude price path. This figure provides sample price paths for Brent Crude under the PCA^2 framework, using the simulation Eq. (40) and the required Q-perturbation Eq. (42). Each path begins with the actual price curve at time 0, shown in solid back, followed by a series of compounding perturbations, shown in grey. Each curve is obtained from it's predecessor by dropping the nearest contract (which expires) and perturbing the rest.

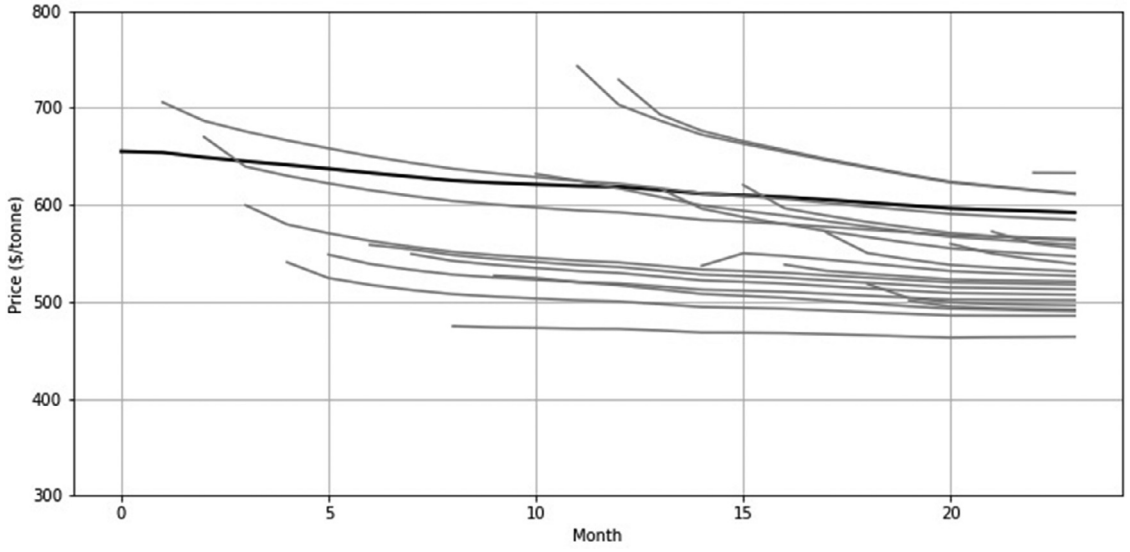

Fig. 5. Low Sulphur Gasoil price path. This figure provides sample price paths for Low Sulphur Gasoil under the PCA^2 framework, using the simulation Eq. (40) and the required Q-perturbation Eq. (42). Each path begins with the actual price curve at time 0, shown in solid back, followed by a series of compounding perturbations, shown in grey. Each curve is obtained from it's predecessor by dropping the nearest contract (which expires) and perturbing the rest.

Although the simulated price curves will tend to increasingly converge from their starting values as simulation time increases, because of the relatively high correlation between crude and gasoil the net short positions for Low Sulphur Gasoil from month-15 onwards will tend to hedge the net long positions for Brent Crude. This hedging affects the valuations at all times in the simulation but especially from month-10 onwards. This is clearly seen in Fig. 9, and also in Figs. 11 and 12 in Section 5.4 below.

5.4. Risk management

Various metrics used for the purpose of risk management (e.g. Value-at-Risk, Expected Shortfall, and Potential Future Exposure) are based on the tail percentiles of valuation distributions. The simulation framework outlined in the previous section allows for such tail percentile calculations. For illustrative purposes, Fig. 10 shows the probability density of portfolio valuations for month-10, and Fig. 11 shows the probability density of portfolio valuations from month-1 onwards. The latter allows a range of percentiles for the portfolio valuations to be readily determined across all simulated periods, as shown in Fig. 12. The latter plot clearly shows the hedging effect discussed in the previous section.

As proof of the efficacy of PCA^2, in addition to those given in Section 3, we compute the portfolio valuation percentiles using PCA^2, PCA with six PCs, denoted PCA(6), and PCA with the full 48 PCs. Table 12 presents the percentage error in valuation percentiles between PCA^2 and the full PCA, and PCA(6) and the full PCA. The results show that PCA^2 and PCA(6) both give results very close to full PCA. Therefore, in terms of these valuation errors, there is not much to recommend one over the other on the basis of accuracy. However, the PCA^2 methodology offers a number of key methodological benefits over PCA and we address these extensively in the next section.

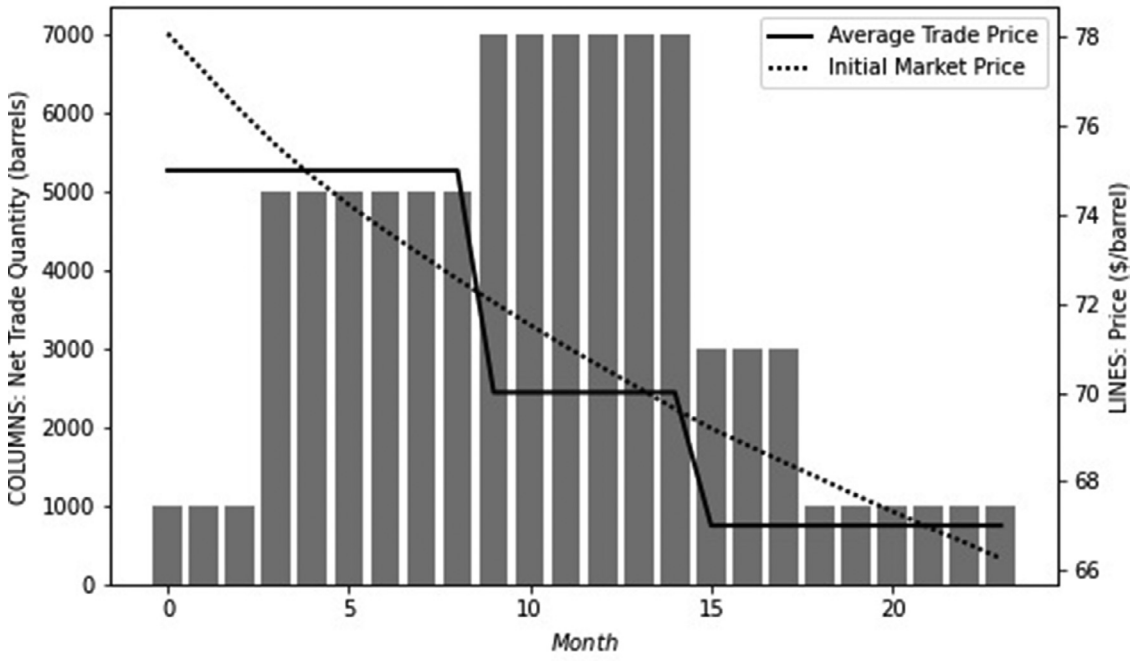

Fig. 6. Brent Crude trades. This figure summarises the Brent Crude trades designed for the valuation exercise under the PCA^2 framework. The scenario design involves net long positions over all months.

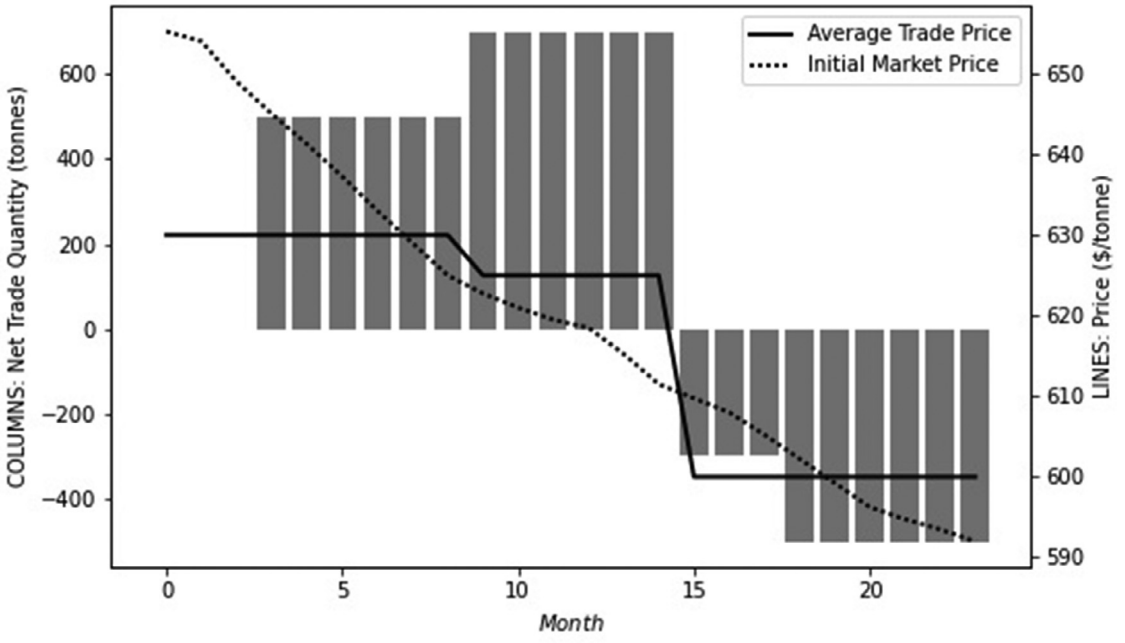

Fig. 7. Low Sulphur Gasoil trades. This figure summarises the Low Sulphur Gasoil trades designed for the valuation exercise under the PCA^2 framework. The scenario design involves net long positions up to month 14 and net short positions from month 15 onwards.

6. Methodological benefits of PCA^2

The previous sections demonstrate the utility of the PCA^2 methodology in terms of how easily it can be implemented for simulation purposes and its practicality in terms of risk management application, while demonstrating its efficacy relative to full- and reduced-dimension PCA. In this section we present the advantages that the PCA^2 methodology offers over standard PCA. Three key methodological benefits relative to PCA are identified and addressed in turn: (1) the computational efficiencies offered by improved scalability; (2) the manner in which components at the curve level can be practically identified; and (3) the capability to proxy curves of interest where insufficient historical data is available.

6.1. Improved scalability

In this section we highlight the scaling benefit of PCA^2 relative to PCA. The scalability of the data structures involved is a very important consideration for the practical application of the methods described here. Very large matrices can create implementation is-

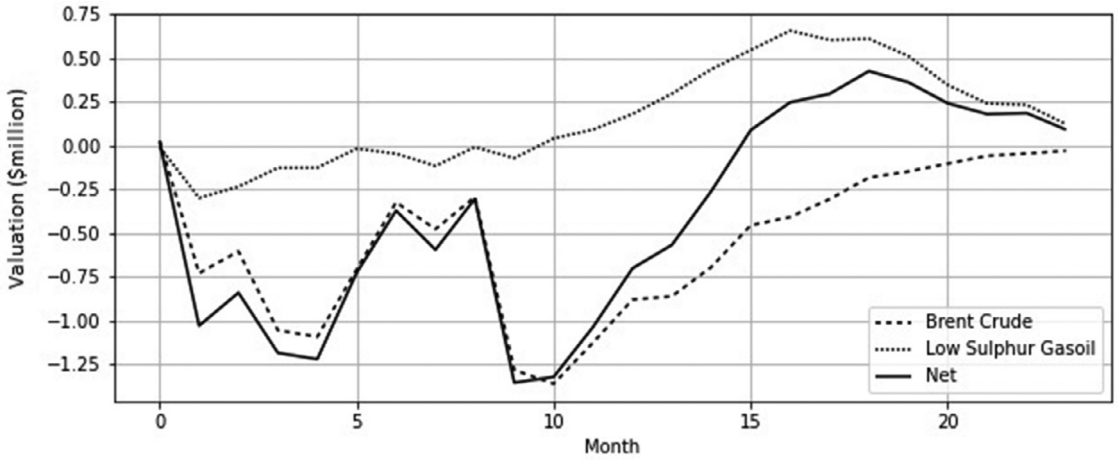

Fig. 8. Portfolio valuation for simulated price paths. This figure plots the valuations of the Brent Crude and Low Sulphur Gasoil sub-portfolios using simulated price paths based on the simulation framework Eq. (40) and the required Q-perturbation Eq. (42). The figure also plots the overall net valuation. Valuations are in millions of dollars.

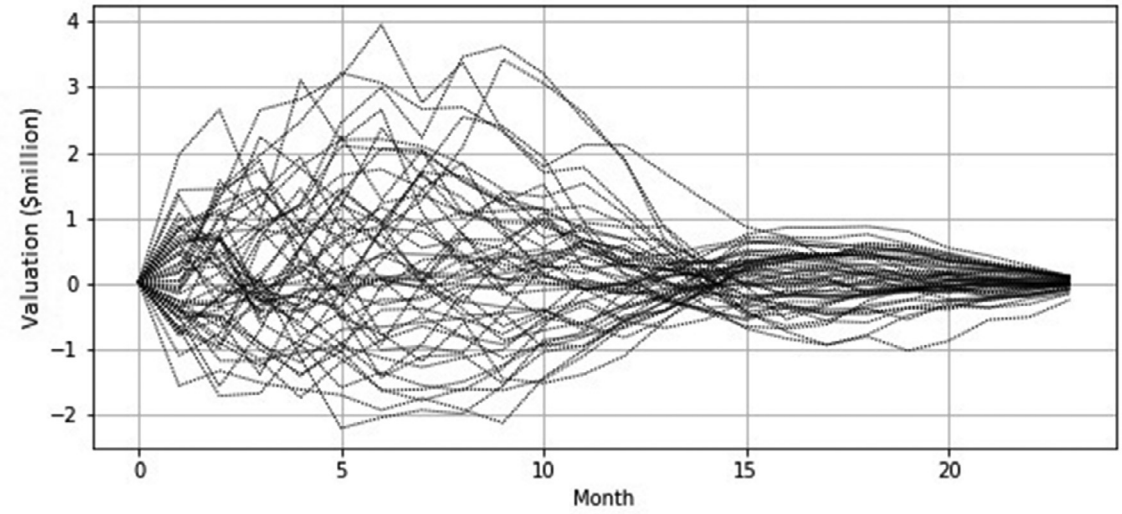

Fig. 9. Sample portfolio valuation paths. This figure plots a selection of fifty valuation paths for the portfolio comprising Brent Crude and Low Sulphur Gasoil. Simulations are based on the simulation framework Eq. (40) and the required Q-perturbation Eq. (42). Valuations are in millions of dollars.

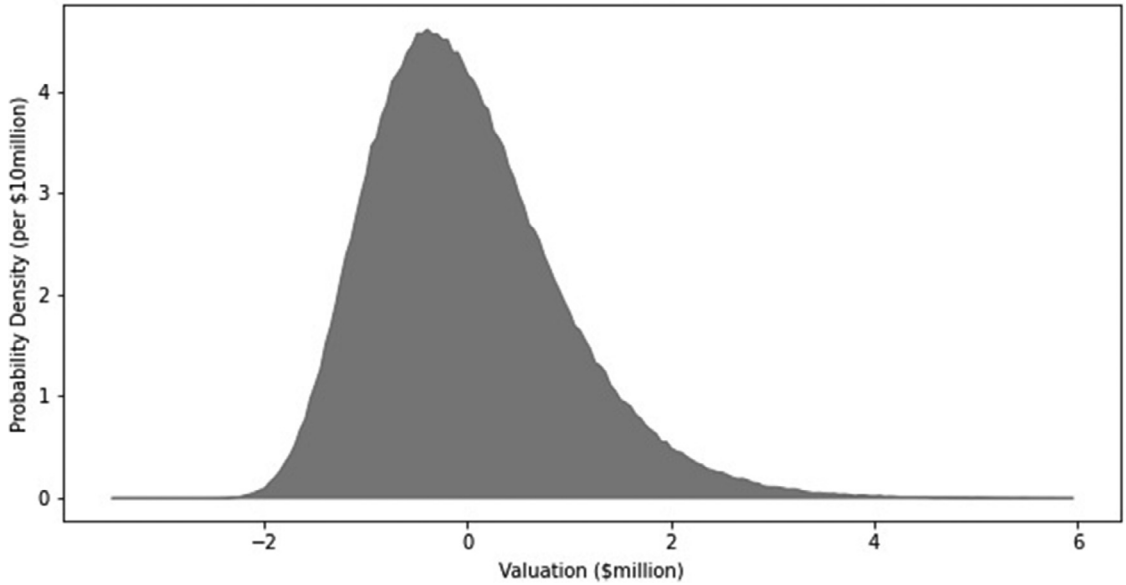

Fig. 10. Probability density of month-10 portfolio valuations. This figure plots the distribution of valuations for simulation month 10. Simulations are based on the simulation framework Eq. (40) and the required Q-perturbation Eq. (42). Valuations are in millions of dollars.

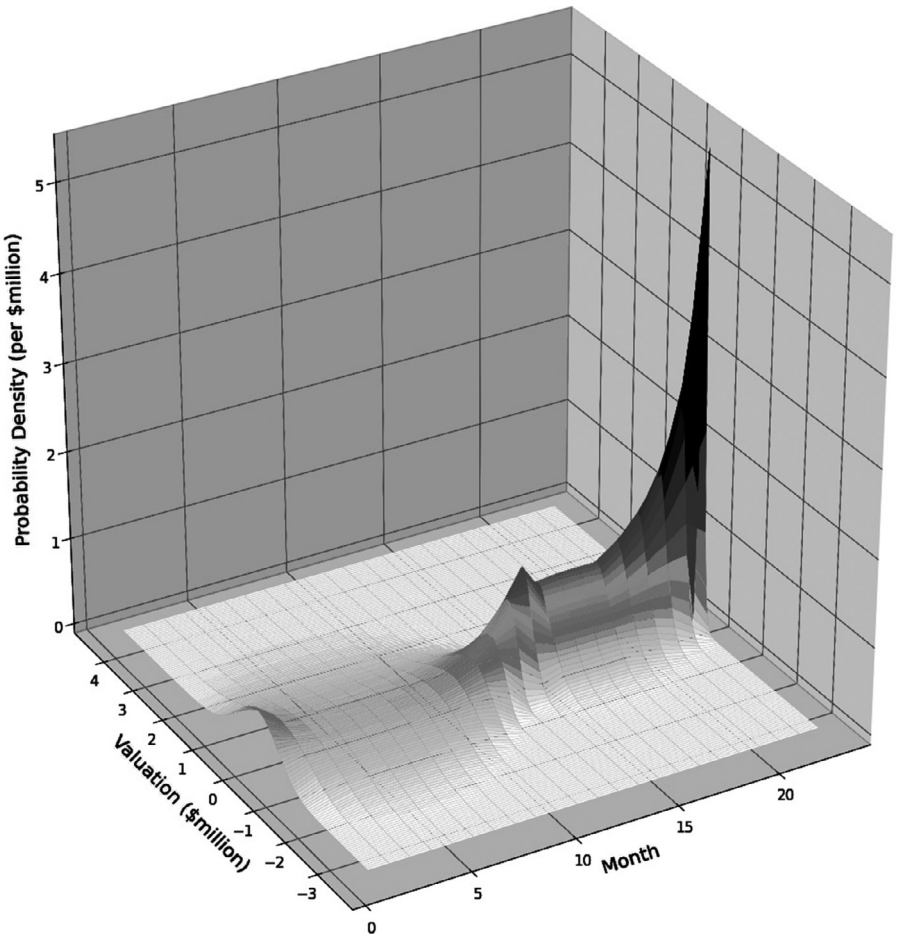

Fig. 11. Portfolio valuation probability densities for months 1 to 24. This figure plots the probability density distribution of valuations for simulation months 1 to 24. Simulations are based on the simulation framework Eq. (40) and the required Q-perturbation Eq. (42). Valuations are in millions of dollars.

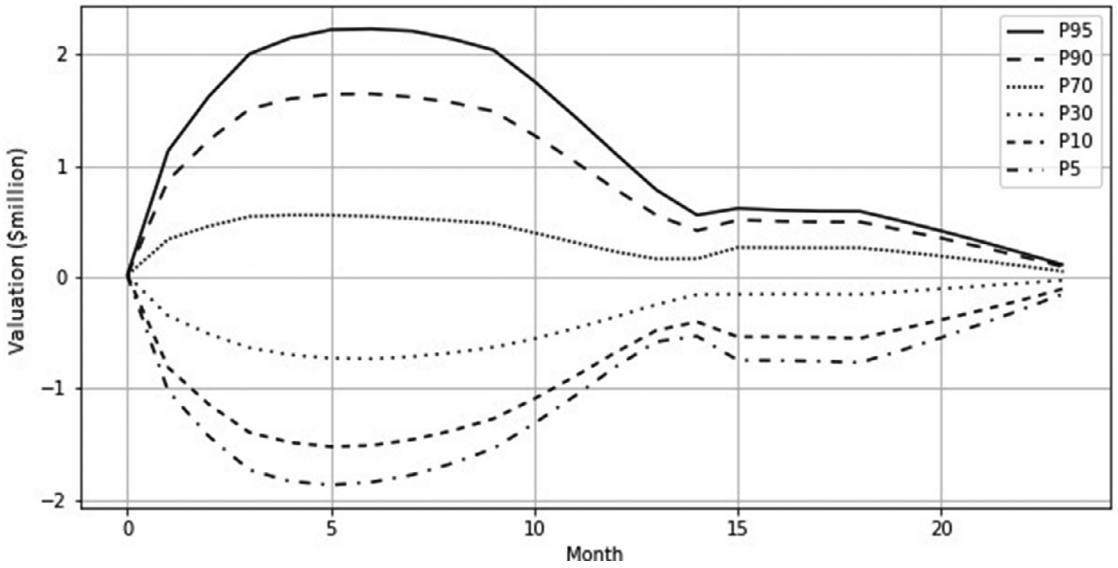

Fig. 12. Portfolio valuation percentiles. This figure plots a range of percentiles for the portfolio valuations across all simulated periods. The percentile range includes the 95th (P95), 90th (P90), 70th (P70), 30th (P30), 10th (P10), and 5th (P5) percentiles. Simulations are based on the simulation framework Eq. (40) and the required Q-perturbation Eq. (42). Valuations are in millions of dollars.

Table 12
Percentage error in portfolio valuation percentiles relative to full PCA.

Pct	Method	1	3	6	9	12	15	18	21	24
95	PCA^2	0.0	0.1	0.0	0.0	−0.1	−0.6	−0.1	−0.6	−0.2
	$PCA(6)$	0.0	0.2	0.1	0.0	0.1	−0.5	−0.3	−0.3	0.1
90	PCA^2	0.0	0.3	−0.1	0.1	−0.2	−0.6	−0.5	−0.2	−0.6
	$PCA(6)$	0.0	0.3	0.0	0.0	0.1	−0.8	−0.2	0.0	−0.1
10	PCA^2	0.0	−0.5	−0.2	−0.1	−0.1	−0.2	0.3	0.0	−0.2
	$PCA(6)$	0.0	−0.4	−0.2	−0.1	−0.3	−0.2	0.2	−0.2	−0.4
5	PCA^2	0.0	−0.5	−0.4	0.0	−0.5	0.0	0.1	−0.2	−0.3
	$PCA(6)$	0.0	−0.4	−0.4	−0.1	−0.6	−0.2	0.1	−0.5	−0.4

This table gives the percentage error in portfolio valuation percentiles for (i) the PCA^2 methodology relative to the full 48-risk factor PCA, and (ii) the six-component PCA methodology relative to the full 48-risk factor PCA. The errors are given for the 95th, 90th, 10th, and 5th percentiles, across the 1st, 3rd, 6th, 9th, 12th, 15th, 18th, 21st and 24th months.

Table 13
Scaling benefit of PCA^2 relative to PCA.

		PCA			PCA^2		Ratios	
c	n	Ω	Δ	m	Ω'	Δ'	Ω/Ω'	Δ/Δ'
2	5	55	65	2	40	26	1.4	2.5
10	10	5,050	1,055	3	1,015	151	5.0	7.0
10	100	500,500	105,050	3	50,965	5,146	9.8	20.4
10	100	500,500	105,050	10	55,550	6,105	9.0	17.2
100	10	500,500	10,055	3	50,650	961	9.9	10.5
100	100	50,005,000	1,005,050	3	550,150	5,956	90.9	168.7
100	100	50,005,000	1,005,050	10	1,005,500	15,105	49.7	66.5
1,000	10	50,005,000	100,055	3	4,556,500	9,061	11.0	11.0
1,000	100	5,000,050,000	10,005,050	3	9,551,500	14,056	523.5	711.8
1,000	100	5,000,050,000	10,005,050	10	55,055,000	105,105	90.8	95.2

This table outlines of the scaling benefit of PCA^2 relative to PCA. c is the number of forward curves, each consisting of n maturities. The number of unique covariances in the PCA scenario is given by Ω as per Eq. (43). Δ is the number of unique covariances added by adding one more curves, as outlined in Eq. (44). m is the number of components retained after dimension reduction in the first step of PCA^2. Ω' is the total number of unique covariances in the PCA^2 scenario as per Eq. (45). Δ' is the total number of unique covariances added as a result of adding another curve, as outlined in Eq. (46). The ratios Ω/Ω' and Δ/Δ' are measures of the scaling benefit, as defined in Eqs. (47) and (48).

Table 14
PC covariances with proxying.

$\mathbf{C}_{\gamma\delta\epsilon}$	$PCa1$	$PCa2$	$PCb1$	$PCb2$	$PCc1$	$PCc2$
$PCa1$	1.000	0.000	0.745	−0.075	0.745	−0.075
$PCa2$	0.000	1.000	0.031	0.285	0.031	0.285
$PCb1$	0.745	0.031	1.000	0.000	0.745	−0.075
$PCb2$	−0.075	0.285	0.000	1.000	0.031	0.285
$PCc1$	0.745	0.031	0.745	0.031	1.000	0.000
$PCc2$	−0.075	0.285	−0.075	0.285	0.000	1.000

This table replicates the covariances of the principal components reported in Table 6, and extends this with the covariances related to some third asset that is proxied by means of copying the principal component covariances observed between Brent Crude and Low Sulphur Gasoil. $PCa1 - PCa2$ and $PCb1 - PCb2$ are respectively the principal components linked to Brent Crude and Low Sulphur Gasoil as noted in Tables 3 and 4. $PCc1 - PCc2$ are the principal components linked to this third asset.

Table 15
Second step of PCA^2 with proxying.

$\mathbf{S}_{\gamma\delta\epsilon}$	$PCa1$	$PCa2$	$PCb1$	$PCb2$	$PCc1$	$PCc2$
$PC1$	0.912	0.011	0.911	−0.044	0.907	−0.098
$PC2$	0.024	−0.723	−0.035	−0.724	−0.093	−0.718
$PC3$	0.056	0.346	−0.033	−0.686	−0.022	0.350
$PC4$	0.018	−0.597	−0.022	0.001	0.075	0.593
$PC5$	0.205	0.014	−0.409	0.020	0.202	−0.033
$PC6$	0.349	−0.011	0.002	0.045	−0.349	0.022

This table reports the result of the second step of PCA, akin to Table 7 but extended for some third asset for which a proxy is sought. $PCa1 - PCa2$ and $PCb1 - PCb2$ are respectively the principal components linked to Brent Crude and Low Sulphur Gasoil as noted in Tables 3 and 4, while $PCc1 - PCc2$ are the principal components linked to this third asset. $PC1 - PC6$ are the principal components derived from the second step of PCA^2.

sues both in terms of memory utilisation and the ability to parallelise computation.

Suppose we want to model c forward curves each consisting of n maturities, then we have cn risk factors overall, and a $cn \times cn$ covariance matrix. The number of unique covariances in this matrix is:

$$\Omega = cn(cn+1)/2 \tag{43}$$

This equation is quadratic in both c and n, so adding curves or extending the number of maturities for each curve becomes more and more expensive. The number of unique covariances added by adding one more curve is:

$$\Delta = cn^2 + n(n+1)/2 \tag{44}$$

PCA^2 greatly improves scalability by reducing the number of factors on a curve-by-curve basis before the overall covariance matrix is derived. For step one we will have $cn(n+1)/2$ unique covariances, and adding another curve will add $n(n+1)/2$ more.

After reducing to m components per curve in the first step, we will then have $cm(cm+1)/2$ unique covariances, and adding another curve will add $cm^2 + m(m+1)/2$ more. Consequently, the total number of unique covariances is:

$$\Omega' = c[m(cm+1) + n(n+1)]/2 \tag{45}$$

The total number of unique covariances added as a result of adding another curve is:

$$\Delta' = cm^2 + [n(n+1) + m(m+1)]/2 \tag{46}$$

The ratios Ω/Ω' and Δ/Δ' are measures of the scaling benefit:

$$\Omega/\Omega' = \frac{n(cn+1)}{m(cm+1) + n(n+1)} \tag{47}$$

Table 16
Final PCA^2 components with proxied curve $\times 1000$.

$\mathbf{S}_{abc}$	*CO*1	*CO*6	*CO*12	*CO*18	*CO*24	*QS*1	*QS*6	*QS*12	*QS*18	*QS*24	*X*1	*X*6	*X*12	*X*18	*X*24
*PC*1	20.1	17.0	14.5	12.8	11.4	16.5	14.5	12.5	11.1	9.9	19.3	17.0	14.7	13.1	11.7
*PC*2	−4.1	1.0	2.1	2.6	2.9	−4.8	0.2	1.1	1.4	1.6	−6.7	−1.2	0.3	0.9	1.4
*PC*3	3.5	0.8	0.0	−0.3	−0.5	−4.5	0.2	1.0	1.3	1.5	1.8	−0.7	−1.2	−1.4	−1.5
*PC*4	−3.4	0.8	1.7	2.1	2.4	−0.4	−0.4	−0.3	−0.3	−0.2	5.5	0.9	−0.2	−0.8	−1.2
*PC*5	4.6	3.8	3.2	2.8	2.5	−7.4	−6.5	−5.6	−5.0	−4.5	4.2	3.8	3.3	3.0	2.7
*PC*6	7.6	6.5	5.6	5.0	4.4	0.3	0.0	−0.1	−0.1	−0.1	−7.5	−6.5	−5.6	−5.0	−4.5

This table reports the final PCA^2 components spanning the original risk factors, akin to Table 8 but extended for some third asset for which a proxy is sought. *PC*1 – *PC*6 are the principal components derived from the second step of PCA^2, as noted in Table 15. Brent Crude is given by the tickers CO1, CO6, CO12, CO18 and CO24, and Low Sulphur Gasoil is given by the tickers QS1, QS6, QS12, QS18 and QS24, corresponding to the futures contracts spanning the first-, sixth-, twelfth-, eighteenth- and twenty-fourth-month maturities. For the third asset, the generic label used is "X", and so X1, X6, X12, X18 and X24, correspond to the futures contracts spanning the first-, sixth-, twelfth-, eighteenth- and twenty-fourth-month maturities.

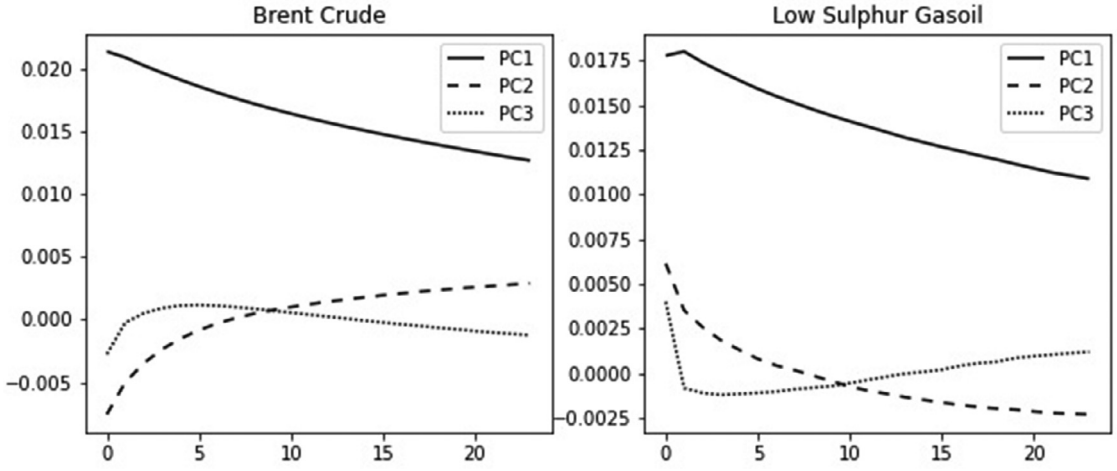

Fig. 13. PCA^2 first-step PCs. This figure shows the first three principal components (PC1-PC3) derived from the first step of the PCA^2 application to the Brent Crude and Low Sulphur Gasoil data as set out in Section 5.

$$\Delta/\Delta' = \frac{cn^2 + n(n+1)}{cm^2 + (n(n+1) + m(m+1))} \tag{48}$$

Table 13 illustrates the scaling benefit of PCA^2 for a range of c, n, and m. Taking, for example, the second last entry of the table, it can be seen that when one faces ($c =$)1,000 forward price curves with ($n =$)100 forward contract maturities each, and one reduces the number of components to ($m =$)3 in step one of PCA^2, dramatic efficiency improvements are observed. Indeed, the total number of unique covariances required sees efficiency improvements over 500 times, while for every price curve added, calculating the incremental covariances is over 700 times more efficient. The last entry of the table, for example, moves from 3 components to 10, under the same c and n settings, but efficiency improvements are still over 90 times.

Moreover, from Eqs. (47) and (48) we see that:

$$\lim_{c\to\infty} \Omega/\Omega' = \lim_{c\to\infty} \Delta/\Delta' = (n/m)^2 \tag{49}$$

Eq. (49) shows that when dealing with a large number of curves, the scaling benefit of PCA^2 is approximately $(n/m)^2$. Even if n and/or m vary by curve, we still know that the scaling benefit will be approximately the squared average $(\overline{n/m})^2$.

6.2. Identifiable components at the curve level

To the trained eye the shift-tilt-bend components of Litterman & Sheinkman (1991) give a good indication of curve behaviour. Data quality issues are often indicated by deviations from these usual observed shapes. However, the principal components for multiple curves have no such interpretation or diagnostic use. The benefit of PCA^2 is that it gives us the benefit of the familiar 'shift', 'tilt' and 'bend' components at the individual curve level. Figure 13 shows, for example, the first-step principal components derived from the PCA^2 application to our Brent Crude and Low Sulphur Gasoil data.

Arguably, however, the biggest benefit of having identifiable first-step principal components is that it facilitates curve proxying, which we describe in the following section.

6.3. Curve proxying

In standard PCA, if we do not have adequate historical data for a forward price curve then this curve either has to be mapped directly to another curve (in which case inter-curve risk is ignored) or excluded entirely. Effectively this means it is impossible to proxy in standard PCA. However, in stark contrast, the two-step approach of PCA^2 allows us to use the first-step principal components from one curve, together with correlation estimates, to create a proxy for another (related) curve. For example, if we had a crude oil curve that could reasonably be expected to behave like Brent, i.e. display similar 'shift', 'tilt' and 'bend' tendencies, then we could use a copy of the first-step eigenvectors S_a of Brent (i.e. *PCa*1 to *PCa*5 shown in Table 3) for our new curve and then determine the necessary extra second-step principal component correlations (i.e. extend Table 6) on a judgement basis. The second PCA step and construction of the PCA^2 components can then proceed as before; leading to components that span all three price curves.

So PCA^2 addresses situations in which the historical data for a curve is insufficient for either full PCA or data imputation meth-

Table 17
Covariances recovered from PCA^2 components with proxied curve $\times 1000$.

$\mathbf{C}_{abc}$	CO1	CO6	CO12	CO18	CO24	QS1	QS6	QS12	QS18	QS24	X1	X6	X12	X18	X24
CO1	0.526	0.404	0.336	0.289	0.253	0.306	0.262	0.226	0.199	0.178	0.366	0.309	0.264	0.233	0.207
CO6	0.404	0.347	0.299	0.264	0.236	0.245	0.221	0.192	0.171	0.153	0.294	0.259	0.225	0.200	0.179
CO12	0.336	0.299	0.261	0.232	0.209	0.206	0.189	0.165	0.147	0.132	0.248	0.221	0.193	0.172	0.155
CO18	0.289	0.264	0.232	0.208	0.188	0.180	0.166	0.146	0.130	0.117	0.216	0.195	0.171	0.152	0.137
CO24	0.253	0.236	0.209	0.188	0.170	0.158	0.148	0.130	0.116	0.105	0.190	0.173	0.152	0.136	0.123
QS1	0.306	0.245	0.206	0.180	0.158	0.370	0.285	0.239	0.207	0.182	0.307	0.258	0.221	0.194	0.173
QS6	0.262	0.221	0.189	0.166	0.148	0.285	0.252	0.218	0.193	0.173	0.249	0.220	0.191	0.170	0.153
QS12	0.226	0.192	0.165	0.146	0.130	0.239	0.218	0.191	0.170	0.153	0.212	0.189	0.165	0.147	0.133
QS18	0.199	0.171	0.147	0.130	0.116	0.207	0.193	0.170	0.152	0.137	0.185	0.167	0.146	0.130	0.117
QS24	0.178	0.153	0.132	0.117	0.105	0.182	0.173	0.153	0.137	0.123	0.164	0.149	0.131	0.117	0.105
X1	0.366	0.294	0.248	0.216	0.190	0.307	0.249	0.212	0.185	0.164	0.526	0.404	0.336	0.289	0.253
X6	0.309	0.259	0.221	0.195	0.173	0.258	0.220	0.189	0.167	0.149	0.404	0.347	0.299	0.264	0.236
X12	0.264	0.225	0.193	0.171	0.152	0.221	0.191	0.165	0.146	0.131	0.336	0.299	0.261	0.232	0.209
X18	0.233	0.200	0.172	0.152	0.136	0.194	0.170	0.147	0.130	0.117	0.289	0.264	0.232	0.208	0.188
X24	0.207	0.179	0.155	0.137	0.123	0.173	0.153	0.133	0.117	0.105	0.253	0.236	0.209	0.188	0.170

This table reports the effective covariance matrix, akin to Table 9, but extended for some third asset for which a proxy is sought. Brent Crude is given by the tickers CO1, CO6, CO12, CO18 and CO24, and Low Sulphur Gasoil is given by the tickers QS1, QS6, QS12, QS18 and QS24, corresponding to the futures contracts spanning the first-, sixth-, twelfth-, eighteenth- and twenty-fourth-month maturities. For the third asset, the generic label used is "X", and so X1, X6, X12, X18 and X24, correspond to the futures contracts spanning the first-, sixth-, twelfth-, eighteenth- and twenty-fourth-month maturities.

ods.[3] To demonstrate this, we add another curve into the example in Section 3 above, by copying the principal component covariances observed between Brent and Gasoil as shown in Table 14. The principal components linked to this third asset are labelled $PCc1 - PCc2$. We can now proceed with the second-step PCA, leading to the six principal components $PC1 - PC6$ detailed in Table 15. We then obtain the PCA^2 components that include the proxied curve as shown in Table 16, where for the third asset we use the generic label "X". Table 17 confirms the covariances recovered from the PCA^2 components with the proxied curve.

The example given above is limited for the purposes of this paper, but in practice we would have a larger number of curves from which to judge the appropriate second-step principal component correlations.

7. Conclusion

Contributing to the limited literature on modelling multiple curves simultaneously, we propose a novel two-step Principal Component Analysis (PCA) approach, which we label PCA^2. We show how readily the PCA^2 methodology can be deployed in the simulation of multiple curves. We then leverage this to demonstrate how the PCA^2 framework can be applied for risk management purposes, outlining the key stages of calibration, simulation, valuation, and risk management. The efficacy of the PCA^2 methodology is established with comparison against the full- and reduced-dimension PCA approaches. We articulate the main benefits of the PCA^2 methodology over PCA: improved scalability allowing for greater computational efficiency and multi-threading, identifiable components at the curve level, and the capability to proxy new curves for which there is little or no data available. The benefits are demonstrated empirically for a two commodity case, but the application is readily extended to any number of forward price curve exposures. The effectiveness and ease of implementation of the PCA^2 methodology will appeal to many practitioners, especially those in the area of risk management.

Future studies that leverage and apply our proposed PCA^2 methodology in alternative market settings offer the potential to provide new insights where modelling multiple curves is required. One particular direction would be to apply PCA^2 to assess the impact on risk measures, such as Value-at-Risk and Expected Shortfall, while benchmarking against a full PCA approach or indeed other possible approaches. A comparative backtesting exercise in a multi-curve setting would provide useful insights for the literature focused on evaluating risk models.

[3] In cases where we have sufficient data but the series are of unequal length we would recommend some form of data imputation method, such as the dynamic Cholesky method described in Atkins & Cummins (2022).

Appendix A. Effective covariance matrix $\mathcal{C}_{ab}$

The effective covariance matrix corresponding to $\mathcal{S}_{ab}$ is given by:

$$\begin{aligned}\mathcal{C}_{ab} &= \mathcal{S}_{ab}^{\top}\mathcal{S}_{ab} = \left(\mathbf{S}_a^{\top} \oplus \mathbf{S}_b^{\top}\right)\mathbf{S}_{\alpha\beta}^{\top}\mathbf{S}_{\alpha\beta}\left(\mathbf{S}_a \oplus \mathbf{S}_b\right)\\ &= \left(\mathbf{S}_a^{\top} \oplus \mathbf{S}_b^{\top}\right)\begin{bmatrix}\mathbf{I}_n & \mathbf{C}_{\alpha\times\beta}\\ \mathbf{C}_{\alpha\times\beta}^{\top} & \mathbf{I}_n\end{bmatrix}\left(\mathbf{S}_a \oplus \mathbf{S}_b\right)\\ &= \begin{bmatrix}\mathbf{C}_a & \mathbf{S}_a^{\top}\mathbf{C}_{\alpha\times\beta}\mathbf{S}_b\\ (\mathbf{S}_a^{\top}\mathbf{C}_{\alpha\times\beta}\mathbf{S}_b)^{\top} & \mathbf{C}_b\end{bmatrix}\end{aligned} \tag{A.1}$$

Since we have not reduced the number of principal components at either PCA step, the effective covariance matrix $\mathcal{C}_{ab}$ should be equal to our original covariance matrix $\mathbf{C}_{ab}$. We can prove this if we assume (or ensure) that the constituent risk-factor returns exhibit zero means, in which case we can re-express the covariance matrix in terms of returns histories, e.g.

$$\mathbf{C}_a = \mathbf{H}_a^{\top}\mathbf{H}_a/N, \quad \mathbf{C}_b = \mathbf{H}_b^{\top}\mathbf{H}_b/N \tag{A.2}$$

We can now write:

$$\mathbf{C}_{ab} = \begin{bmatrix}\mathbf{H}_a^{\top}\\ \mathbf{H}_b^{\top}\end{bmatrix}\begin{bmatrix}\mathbf{H}_a & \mathbf{H}_b\end{bmatrix}/N = \begin{bmatrix}\mathbf{H}_a^{\top}\mathbf{H}_a & \mathbf{H}_a^{\top}\mathbf{H}_b\\ \mathbf{H}_b^{\top}\mathbf{H}_a & \mathbf{H}_b^{\top}\mathbf{H}_b\end{bmatrix}/N = \begin{bmatrix}\mathbf{C}_a & \mathbf{C}_{a\times b}\\ \mathbf{C}_{a\times b}^{\top} & \mathbf{C}_b\end{bmatrix} \tag{A.3}$$

$$\mathbf{C}_{\alpha\times\beta} = \mathbf{H}_{\alpha}^{\top}\mathbf{H}_{\beta}/N = (\mathbf{H}_a\mathbf{S}_a^{-1})^{\top}\mathbf{H}_b\mathbf{S}_b^{-1}/N = (\mathbf{S}_a^{\top})^{-1}\mathbf{C}_{a\times b}\mathbf{S}_b^{-1} \tag{A.4}$$

From the equation above we derive:

$$\mathbf{S}_a^{\top}\mathbf{C}_{\alpha\times\beta}\mathbf{S}_b = \mathbf{C}_{a\times b} \tag{A.5}$$

which with Eq. (A.1) yields:

$$\mathcal{C}_{ab} = \mathbf{C}_{ab} \tag{A.6}$$

This proves that without dimension reduction PCA^2 returns the original covariance matrix.

Appendix B. The matrix $\mathbf{J}_{ab}$

Let us look at $\mathbf{J}_{ab}$ in more detail by defining:

$$(\mathbf{J}_a \oplus \mathbf{J}_b) := \mathbf{J}_{ab}$$

and considering the properties of $\mathbf{J}_a = \mathbf{U}_a^\dagger \mathbf{U}_a$. Firstly, we can see that $\mathbf{J}_a$ is symmetric:

$$\begin{aligned}\mathbf{J}_a^\top &= \left[\mathbf{U}_a^\dagger \mathbf{U}_a\right]^\top = \left[\mathbf{U}_a^\top (\mathbf{U}_a\mathbf{U}_a^\top)^{-1}\mathbf{U}_a\right]^\top = \mathbf{U}_a^\top\left[(\mathbf{U}_a\mathbf{U}_a^\top)^{-1}\right]^\top\mathbf{U}_a \\ &= \mathbf{U}_a^\top\left[(\mathbf{U}_a\mathbf{U}_a^\top)^\top\right]^{-1}\mathbf{U}_a \\ &= \mathbf{U}_a^\top(\mathbf{U}_a\mathbf{U}_a^\top)^{-1}\mathbf{U}_a = \mathbf{J}_a \end{aligned} \tag{B.1}$$

We can also see that $\mathbf{J}_a$ is idempotent:

$$\begin{aligned}\mathbf{J}_a^2 &= \left[\mathbf{U}_a^\top (\mathbf{U}_a\mathbf{U}_a^\top)^{-1}\mathbf{U}_a\right]\left[\mathbf{U}_a^\top (\mathbf{U}_a\mathbf{U}_a^\top)^{-1}\mathbf{U}_a\right] \\ &= \mathbf{U}_a^\top (\mathbf{U}_a\mathbf{U}_a^\top)^{-1}\left[(\mathbf{U}_a\mathbf{U}_a^\top)(\mathbf{U}_a\mathbf{U}_a^\top)^{-1}\right]\mathbf{U}_a \\ &= \mathbf{U}_a^\top(\mathbf{U}_a\mathbf{U}_a^\top)^{-1}\mathbf{U}_a = \mathbf{J}_a \end{aligned} \tag{B.2}$$

It is clear that $\mathbf{J}_b$ must also be symmetric and idempotent; and consequently so must $\mathbf{J}_{ab}$. This allows us to simplify Eq. (33) to:

$$\begin{aligned}\mathbb{C}_{ab} &= \mathbf{J}_{ab}\mathbf{C}_{ab}\mathbf{J}_{ab} = (\mathbf{J}_{ab}^{-1}\mathbf{J}_{ab})\mathbf{J}_{ab}\mathbf{C}_{ab}\mathbf{J}_{ab} = \mathbf{J}_{ab}^{-1}\mathbf{J}_{ab}^2\mathbf{C}_{ab}\mathbf{J}_{ab} \\ &= (\mathbf{J}_{ab}^{-1}\mathbf{J}_{ab})\mathbf{C}_{ab}\mathbf{J}_{ab} \\ &= \mathbf{C}_{ab}\mathbf{J}_{ab} \end{aligned} \tag{B.3}$$

We can now compute the "error" introduced by applying PCA^2 with dimension reduction at the first step as:

$$\mathbb{C}_{ab} - \mathbf{C}_{ab} = \mathbf{C}_{ab}(\mathbf{J}_{ab} - \mathbf{I}_{2n}) \tag{B.4}$$

If we consider the *Frobenius Norm* of $(\mathbf{J}_{ab} - \mathbf{I}_{2n})$:

$$\begin{aligned}\left\|\mathbf{J}_{ab} - \mathbf{I}_{2n}\right\|_F^2 &= tr\left([\mathbf{J}_{ab} - \mathbf{I}_{2n}]^\top[\mathbf{J}_{ab} - \mathbf{I}_{2n}]\right) \\ &= tr\left(\mathbf{J}_{ab}^\top\mathbf{J}_{ab} - 2\mathbf{J}_{ab} + \mathbf{I}_{2n}\right) = tr(\mathbf{I}_{2n} - \mathbf{J}_{ab}) \\ &= 2n - tr(\mathbf{J}_a) - tr(\mathbf{J}_b) = 2n - 2m \end{aligned}$$

The final step above relies on the fact that:

$$tr(\mathbf{J}_a) = rank(\mathbf{J}_a) = m$$

In other words:[4]

$$\left\|\mathbf{J}_{ab} - \mathbf{I}_{2n}\right\|_F = \sqrt{2(n-m)} \tag{B.5}$$

This equation is a general and surprisingly simple measure of the extent to which dimension reduction at the first step affects the modelled covariance matrix, although in practice it may be more informative to compute the error directly using Eq. (B.4).

Appendix C. PCA^2 with dimension reduction at both steps

In principle, dimension reduction can be applied at the second step of PCA. As an analogue of the components partition given in Eq. (9), we define as follows:

$$\mathbf{S}_{\gamma\delta} = \begin{bmatrix}\mathbf{X}_{\gamma\delta} \\ \mathbf{Y}_{\gamma\delta}\end{bmatrix} \tag{C.1}$$

were $\mathbf{X}_{\gamma\delta}$ are the selected components. The block view from Eq. (30) can then be extended:

$$\mathbf{S}_{\gamma\delta} = \begin{bmatrix}\mathbf{X}_{\gamma|\gamma\delta} & \mathbf{X}_{\delta|\gamma\delta} \\ \mathbf{Y}_{\gamma|\gamma\delta} & \mathbf{Y}_{\delta|\gamma\delta}\end{bmatrix} \tag{C.2}$$

The analogue of Eq. (31) with dimension reduction at both steps can then be defined as follows, where we use the notation $\check{\mathbb{S}}_{ab}$ to distinguish from the earlier notation $\mathbb{S}_{ab}$:

$$\check{\mathbb{S}}_{ab} = \mathbf{X}_{\gamma\delta}(\mathbf{U}_a \oplus \mathbf{U}_b) = \begin{bmatrix}\mathbf{X}_{\gamma|\gamma\delta}\mathbf{U}_a & \mathbf{X}_{\delta|\gamma\delta}\mathbf{U}_b\end{bmatrix} \tag{C.3}$$

[4] Note that for general n's and m's this becomes $\left\|\mathbf{J}_{all} - \mathbf{I}_{\Sigma n}\right\|_F = \sqrt{\Sigma n - \Sigma m}$.

It follows that the effective covariance matrix corresponding to $\check{\mathbb{S}}_{ab}$ is:

$$\check{\mathbb{C}}_{ab} = \quad \check{\mathbb{S}}_{ab}^\top\check{\mathbb{S}}_{ab} = \left(\mathbf{U}_a^\top \oplus \mathbf{U}_b^\top\right)\mathbf{X}_{\gamma\delta}^\top\mathbf{X}_{\gamma\delta}\left(\mathbf{U}_a \oplus \mathbf{U}_b\right) \tag{C.4}$$

Comparing Eqs. (C.4) to (32) we see that additional differences will be introduced to the extent that $\mathbf{X}_{\gamma\delta}^\top\mathbf{X}_{\gamma\delta}$ differs from $\mathbf{S}_{\gamma\delta}^\top\mathbf{S}_{\gamma\delta}$.

References

Alexander, C. (2008). *Market risk analysis, practical financial econometrics*: vol. 2. John Wiley & Sons.
Alexander, C. (2009). *Market risk analysis, value at risk models*: vol. 4. John Wiley & Sons.
Atkins, P. J., & Cummins, M. (2022). A dynamic Cholesky data imputation method for correlation structure consistency. *Applied Economics Letters, 29*(4), 311–315.
Ballotta, L., Fsai, G., Lregian, A., & Prez, M. F. (2019). Estimation of multivariate asset models with jumps. *Journal of Financial and Quantitative Analysis, 54*(5), 2053–2083.
Barber, J. R., & Cpper, M. L. (1996). Immunization using principal component analysis. *The Journal of Portfolio Management, 23*(1), 99–105.
Blanco, C., Sronow, D., & Seszyn, P. (2002). Multi-factor models of the forward price curve. In *Commodities now* (pp. 80–83). September
Buhler, W., Urig-Homburg, M., Wlter, U., & Wber, T. (1999). An empirical comparison of forward-rate and spot-rate models for valuing interest-rate options. *The Journal of Finance, 54*(1), 269–305.
Canbas, S., Cbuk, A., & Klic, S. B. (2005). Prediction of commercial bank failure via multivariate statistical analysis of financial structures: The Turkish case. *European Journal of Operational Research, 166*(2), 528–546.
Carmona, R., & Culon, M. (2014). A survey of commodity markets and structural models for electricity prices. In *Quantitative energy finance* (pp. 41–83). Springer.
Chantziara, T., & Siadopoulos, G. (2008). Can the dynamics of the term structure of petroleum futures be forecasted? Evidence from major markets. *Energy Economics, 30*(3), 962–985.
Clewlow, L., & Srickland, C., et al. (1999). A multi-factor model for energy derivatives. Technical report
Cortazar, G., & Shwartz, E. S. (1994). The valuation of commodity contingent claims. *Journal of Derivatives, 1*(4), 27–39.
Crepey, S., Gbac, Z., Nor, N., & Sovmand, D. (2015). A levy HJM multiple-curve model with application to CVA computation. *Quantitative Finance, 15*(3), 401–419.
Crepey, S., Gbac, Z., & Nuyen, H. N. (2011). A defaultable HJM multiple-curve term structure model. Preprint, U. Évry.
Cuchiero, C., Fntana, C., & Goatto, A. (2016). A general HJM framework for multiple yield curve modelling. *Finance and Stochastics, 20*(2), 267–320.
Ellefsen, P. E., & Slavounos, P. D. (2009). *Multi-factor model of correlated commodity-forward curves for crude oil and shipping markets*. Center for Energy and Environmental Policy Research.
Falkenstein, E., & Hnweck, J. (1997). Minimizing basis risk from non-parallel shifts in the yield curve part ii: Principal components. *Journal of Fixed Income, 7*, 85–90.
Fanelli, V. (2016). A defaultable HJM modelling of the Libor rate for pricing basis swaps after the credit crunch. *European Journal of Operational Research, 249*(1), 238–244.
Flury, B. N. (1984). Common principal components in k groups. *Journal of the American Statistical Association, 79*(388), 892–898.
Gerhart, C., & Ltkebohmert, E. (2020). Empirical analysis and forecasting of multiple yield curves. *Insurance: Mathematics and Economics, 95*, 59–78.
Grasselli, M., & Mglietta, G. (2016). A flexible spot multiple-curve model. *Quantitative Finance, 16*(10), 1465–1477.
Heath, D., Jrrow, R., & Mrton, A. (1992). Bond pricing and the term structure of interest rates: A new methodology for contingent claims valuation. *Econometrica: Journal of the Econometric Society, 60*(1), 77–105.
Koekebakker, S., & Olmar, F. (2005). Forward curve dynamics in the nordic electricity market. *Managerial Finance, 31*(6), 73–94.
Laurini, M. P., & Oashi, A. (2015). A noisy principal component analysis for forward rate curves. *European Journal of Operational Research, 246*(1), 140–153.
Lekkos, I. (2000). A critique of factor analysis of interest rates. *The Journal of Derivatives, 8*(1), 72–83.
Litterman, R., & Sheinkman, J. (1991). Common factors a affecting bond returns. *Journal of Fixed Income, 1*(1), 54–61.
Lord, R., & Plsser, A. (2007). Level slope curvature fact or artefact?Applied Mathematical Finance, 14(2) 105–130.
Mercurio, F. (2010). Modern LIBOR market models: Using different curves for projecting rates and for discounting. *International Journal of Theoretical and Applied Finance, 13*(01), 113–137.
Moreni, N., & Pllavicini, A. (2014). Parsimonious HJM modelling for multiple yield curve dynamics. *Quantitative Finance, 14*(2), 199–210.
Morino, L., & Rnggaldier, W. J. (2014). On multicurve models for the term structure. In *Nonlinear economic dynamics and financial modelling* (pp. 275–290). Springer.
Sabelli, C., Poppi, M., Stzia, L., & Brmetti, G. (2018). Multi-curve HJM modelling for risk management. *Quantitative Finance, 18*(4), 563–590.
Schmidt, W. M. (2011). Interest rate term structure modelling. *European Journal of Operational Research, 214*(1), 1–14.

Shie, F. S., Chen, M.-Y., & Liu, Y. S. (2012). Prediction of corporate financial distress: An application of the america banking industry. *Neural Computing and Applications, 21*(7), 1687–1696.
Singh, M. K. (1997). Value at risk using principal components analysis. *Journal of Portfolio Management, 24*(1), 101.
Soto, G. (2004). Using principal component analysis to explain term structure movements: Performance and stability. In V. Tavidze (Ed.), *Progress in economics research: vol. 8*.
Stambaugh, R. F. (1988). The information in forward rates: Implications for models of the term structure. *Journal of Financial Economics, 21*(1), 41–70.
Steeley, J. M. (1990). Modelling the dynamics of the term structure of interest rates. *Working Paper*.
Tolmasky, C., & Hindanov, D. (2002). Principal components analysis for correlated curves and seasonal commodities: The case of the petroleum market. *Journal of Futures Markets, 22*(11), 1019–1035.
Trendalov, N. T. (2010). Stepwise estimation of common principal components. *Computational Statistics & Data Analysis, 54*(12), 3446–3457.
West, R. C. (1985). A factor-analytic approach to bank condition. *Journal of Banking & Finance, 9*(2), 253–266.