Essays on Portfolio Risk Management and Weather Derivatives

Anders Daasvand Sleire

Thesis for the degree of Philosophiae Doctor (PhD)
University of Bergen, Norway
2023

UNIVERSITY OF BERGEN

Essays on Portfolio Risk Management and Weather Derivatives

Anders Daasvand Sleire

Thesis for the degree of Philosophiae Doctor (PhD)
at the University of Bergen

Date of defense: 25.08.2023

Year: 2023

Title: Essays on Portfolio Risk Management and Weather Derivatives

Name: Anders Daasvand Sleire

Print: Skipnes Kommunikasjon / University of Bergen

Preface

This dissertation is submitted as a partial fulfillment of the requirements for the degree Doctor of Philosophy (PhD) at the Department of Mathematics, University of Bergen. The research has recieved support from the Research Council of Norway via the industrial PhD scheme. The study was supervised by Prof. Bård Støve and Prof. Jan Bulla. The present work has been carried out during my time as an industrial PhD student. First at *Kinect Energy Markets* (2015-2017), second at *Tryg Forsikring* (2017-2021), and finally at the FinTech start up *Horde*, from January 2022 until present day.

This dissertation pursues several objectives. The first one is to develop a method for performing asset allocation when there are asymmetric dependence structures between instruments in investment portfolios using the theory of local Gaussian correlation. The second aim is to provide market practitioners and researchers with open source software tools for risk management and commodity price hedging in energy portfolios. The third aim is to develop a framework for designing and pricing derivatives contracts for managing financial risks arising from severe urban air pollution.

This thesis consists of two parts. The first part provide a brief introduction to dependence modelling, energy market risk management and weather derivatives. The second part consists of three papers:

Paper A Anders D. Sleire, Bård Støve, Håkon Otneim, Geir Drage Berentsen, Dag Tjøstheim and Sverre Hauso Haugen (2021), *Portfolio Allocation under Asymmetric Dependence in Asset Returns using Local Gaussian Correlations*, Finance Research Letters, p. 102475.

Paper B Anders D. Sleire (2022), *etmr: Energy Trading and Risk Management in R*, The R-Journal, 14/1, p. 320.

Paper C Anders D. Sleire (2023), *Modelling and Pricing Air Pollution Derivatives*, working paper.

Acknowledgements

First of all, I would like to thank my supervisor, Professor Bård Støve, for his kind assistance and support throughout the project. Støve was also the supervisor for my Master's project, and it has been a pleasure working together during all these years. When I started on the Master's programme while holding a full time position on the trading floor of an energy company, I was a somewhat atypical student at the department. Operational tasks and deadlines from industry work had to be carefully managed, and Støve's flexible approach and guidance proved invaluable. As an industrial PhD student, I have continued to balance day-to-day industry work with the PhD project, and Støve has supported me in an excellent manner.

In addition, I want to express gratitude to my coauthors and the staff and colleagues at the Department of Mathematics. It has been a pleasure to be a part of this department.

I would also like to thank my former colleagues in Kinect Energy Markets for 11 years with interesting discussions and learning opportunities. It has been a true privilege to participate in a market undergoing such fundamental change and development. I am also grateful to Morten Hegna at Montel for providing me access to their market data platform. I also wish to thank to my former and current colleagues at Tryg and Horde.

I also would like to acknowledge NVIDIA Corporation for providing their GPUs in the academic GPU Grant Program.

Finally, I could not have done this without the support from my friends and family. In particular my parents, Tone and Arne Henrik, the love of my life Kristine, and our children Bjørnar, Ebba and Hennie. Without you, Kristine, this would not be possible.

Bergen, March 2023
Anders D. Sleire

Abstract

This thesis is concerned with the development and practical implementation of risk management methods for investment portfolios, energy portfolios, and weather and pollution risk. The thesis includes three scientific papers that each address different aspects of financial risk management. The first paper focuses on portfolio allocation in the presence of asymmetric dependence between asset returns. The second paper examines energy price risk management, and introduces an open source toolkit for energy portfolio management which has been developed as a part of the PhD project. The final paper present a theoretical framework for managing pollution risk using financial derivatives contracts, which builds upon the existing theory of weather derivatives. These papers all contribute to the overall theme, which is the development of risk management methods for various types of portfolios and the exploration of the role of financial derivatives in managing risks related to market prices, weather and pollution.

In order to provide a theoretical context, we have included a brief chapter exploring alternative methods for dependence modelling and how these may be utilized when managing investment portfolios. One of these measures, the local Gaussian correlation, is used to extend the classical mean-variance framework for asset allocation in the first paper. Thereafter, a short introduction to spot and forward energy markets is provided. The primary focus here is commodity market price risk, and how this can be managed with financial derivatives contracts. We demonstrate how portfolio management may be performed with our open source toolkit using European energy market data. Finally, we include a chapter on weather derivatives. This contains a introduction to weather related risk, a brief introduction to the weather markets, frequently used contract types and pricing methods.

To ensure reproducibility, we have also added a chapter on computer code, where the interested reader may find links to Git repositories with all data and the R code needed to run the analysis presented in the thesis.

Sammendrag

Denne avhandlingen handler om utvikling og praktisk implementering av risikostyringsmetoder for investeringsporteføljer, energiporteføljer, og håndtering av vær- og forurensningsrisiko. Avhandlingen inkluderer tre vitenskapelige artikler som hver tar for seg ulike aspekter av finansiell risikostyring. Den første fokuserer på metoder for aktivaallokering når det eksisterer asymmetrisk avhengighet mellom avkastningene for eiendelene i en investeringsportefølje. Den andre artikkelen omhandler energiprisrisikostyring, og introduserer et åpen kildekodeverktøy for energiporteføljeforvaltning som er utviklet som en del av doktorgradsprosjektet. Den siste artikkelen presenterer et teoretisk rammeverk for håndtering av forurensningsrisiko ved hjelp av finansielle derivatkontrakter, som bygger på den eksisterende teorien om værderivater. Disse arbeidene bidrar alle til det overordnede temaet for avhandlingen, som er utvikling av risikostyringsmetoder for ulike typer porteføljer og utforskingen av rollen til finansielle derivater i håndtering av risiko knyttet til markedspriser, vær og forurensning.

For å sette bidragene inn i en teoretisk kontekst har vi inkludert et kort kapittel som presenterer alternative metoder for avhengighetsmodellering, og hvordan disse kan utnyttes når man forvalter investeringsporteføljer. Ett av disse målene, lokal gaussisk korrelasjon, brukes til å utvide det klassiske mean-variance-rammeverket for aktivaallokering i den første artikkelen. Deretter følger et kort introduksjonskapittel til spot- og forwardmarkeder for energi. Hovedfokuset her er råvareprisrisiko, og hvordan denne kan håndteres med finansielle derivatkontrakter. Vi demonstrerer hvordan forvaltning av energiporteføljer kan gjennomføres med vårt åpen kildekodeverktøy ved bruk av data fra det europeiske kraftmarkedet. Til slutt inkluderes et kapittel om værderivater. Dette inneholder en introduksjon til værrelatert risiko, en kort introduksjon til værmarkedet, vanlige kontraktstyper og alternative metoder for prising.

For å sikre reproduserbarhet har vi også lagt til et kapittel om programkode. Her finnes lenker til Git-repositorier med alle data og R-kode for å gjennomføre analysene som presenteres i avhandlingen.

Contents

Chapter 1

Introduction

The practice of financial risk management has a long history dating back to the early days of commerce and finance. The widespread use of financial derivatives, however, is a fairly recent phenomenon. Financial derivatives, e.g. contracts that derive their value from an underlying asset or index, was first applied in commodities markets to insure farmers against crop failure and facilitate trade [Swan, 2000]. Some of the first organised marketplaces include the Antwerp exchange for trading of commodities and bills of exchange in the Netherlands (16^{th} century), and the Dojima rice exchange in Osaka, Japan (17^{th} century). There are numerous examples of uses prior to this, where transaction were executed over-the-counter [Kummer and Pauletto, 2012]. In modern financial markets, there is a plethora of contract types available, and these may be written on stocks, interest rates, commodities, currencies, weather indices and a number of other settlement references.

The utilization of derivatives as a risk management tool has become increasingly popular and sophisticated over time, and they continue to be widely used to manage various types of risks, [Dionne, 2013]. Financial and non-financial institutions have different approaches to these activities. Non-financial institutions, such as corporations and governments, may utilize financial products to manage treasuries or to hedge risks, such as commodity price risk or interest rate risk. Financial institutions, on the other hand, actively assume risks to make a profit, either for their own account or as trustees for third parties. Derivatives contracts may be used as potent tools for speculation due to the gearing effect, and it is not always trivial to determine if a transaction should be evaluated as hedging or trading (speculative) [Goss et al., 1976]. The financial crisis of 2008 highlighted the risks associated with complex financial products. The risk assessments and lax oversight of these instruments prior to the crisis were insufficient, which contributed to significant losses and a global economic downturn [Bordo, 2008].

The intricacies of products such as the Collaterized Debt Obligations also offered an opportunity to hide risks or to pass them on to less vigilant market participants. This underscores the importance of proper risk assessment and transparent risk management practices in finance. In order for to these to be established, the properties of traded contracts and related markets needs to be properly understood. This requires a solid theoretical foundation and models that are able to describe and explain the empirical observations from the marketplace reasonably well.

One of the early mathematical theories of finance can be found in *Théorie de la spéculation* [Bachelier, 1900], where the concept of Brownian motion was first introduced and used to model the behavior of stocks and option prices [Schachermayer and Teichmann, 2008]. The contribution was not well received at time of publication, but has later been recognised as influential on both the study of stochastic processes, and the later *Black-Scholes option pricing model* [Black and Scholes, 1973]. Other significant contributions include the *mean-variance model of portfolio selection* developed by Harry Markowitz [Markowitz, 1952] and the *Capital Asset Pricing Model* [Treynor, 1961], [Sharpe, 1964], [Lintner, 1965], [Mossin, 1966]. There are numerous other authors who have played an important part in building the foundations of modern finance theory over the last decades. While it may be tempting to elaborate further on the history of this discipline, we will constrain ourselves to maintain focus on the topic at hand.

Options, and other derivatives such as futures and swaps may be used as effective hedging tools when they are included in a portfolio within pre-specified boundaries. This is typically formalized in an investment mandate with clear goals, benchmarks and risk limits, which is monitored by by risk managers. The status of the portfolio is often evaluated by calculating risk measures, such as *Value at Risk* and *Expected Shortfall*, see [Duffie and Pan, 1997], [McNeil et al., 2015], [Tasche, 2002]. The use of derivatives contracts have often been motivated by the desire to create insurance-like effects, e.g. attempts to provide guarantees that an investment portfolio is unlikely to fall below a specific value. Such *portfolio insurance strategies* [Black and Jones, 1987], [Leland, 1980] may also be implemented without the use of derivatives contracts, by dynamically allocating capital between risky and non-risky assets in a manner that replicate the effects of a financial derivative. These are the principles we utilize in the **etrm** package in **Paper B**, which is used for energy price hedging in Chapter 3 of the thesis. Not all derivatives are possible to replicate using these arbitrage arguments. If the underlying settlement reference is non-tradable, the value of the contract cannot be calculated using this approach. Weather related risk is a significant concern for many market participants, and the pricing of these instruments require other methods, since the underlying weather indices are not traded in any market. Further discussions of this subject and alternative methods for valuation are presented in Chapter 4 and **Paper C**.

Another topic to be investigated in this thesis is the discipline of managing investment portfolio risk via *diversification*, within the classical mean-variance framework which was pioneered by Harry Markowitz [Markowitz, 1952]. In such a setting, the objective is not to replicate the payoff structure of a derivative, but to optimize the portfolio composition by allocating funds across a set of assets that are not exposed to the same risk dynamics. By investing in mix of instruments that have low correlation with each other, the performance of one asset should not have a significant impact on the overall portfolio. In [Markowitz, 1952], a risk-averse investor facing the portfolio selection problem considers both expected return and risk (variance of portfolio returns) as criteria for making a choice regarding portfolio weights. This dual problem might be formulated in two ways. Either as a wish to i) maximize expected return subject to a constraint on portfolio variance, or as a desire to ii) minimize portfolio variance subject to a constraint on expected return. Under the assumed constant dependency structure between the assets included in the portfolio, the calculated weights is expected to be optimal, given the investors preferences.

However, the dependence structures have not proven to be stable over time. It is well known that asymmetries are present in financial returns, both in the form of skewness in distributions for the individual instruments, and in the dependencies between them [Patton, 2004]. This is particularly visible in times of turmoil. Stocks belonging to different industries might have weak correlation during normal market conditions, but move towards unity in a scenario when the market is under stress. This has important implications for portfolio construction and risk management, as the diversification effect may be significantly reduced in time of crisis, when it is most needed. The evidence of asymmetric dependence between asset returns highlights the need for more advanced risk management and portfolio construction techniques. These should be based on methods that are able to model the dependencies observed in the market during differing conditions.

In Chapter 2 and **Paper A** we suggest using a new nonparametric measure of local dependence, the local Gaussian correlation, to improve portfolio allocation. With this approach, we are able to build easy-interpretable representations of the dependency structures (local correlation matrices) that fits well into the familiar mean-variance framework. Asset allocation can also be improved by using competing methods, such as copulas. These are also presented and discussed further in the following chapter.

Chapter 2

Modelling dependence in finance

The Pearson product-moment correlation coefficient, also known as Pearson's ρ, is a widely used measure of statistical dependence. It measures the linear association between random variables, but it has limitations when the distribution of the variables deviates from a normal distribution or the relationship between them is non-linear. In such cases, the results obtained from the correlation coefficient may be misleading. In finance, this can have serious implications as the Pearson's ρ may not accurately capture the dependence between financial assets. In a mean-variance asset allocation setting, this implies the global covariance matrix would not be a sufficient representation of the dependence structure, and the investor may experience reduced effects from diversification. The shortcomings of Pearson's ρ has led to the development of alternative measures of dependence. Copulas [Sklar, 1959], in particular, are popular tools in finance as they can model the dependence structure between variables also when the variables have different marginal distributions. A more recent developed measure is the local Gaussian correlation, which was introduced in [Tjøstheim and Hufthammer, 2013]. This local characterization of dependence has also proven to be an effective instrument in many applications, including modeling asymmetric dependence patterns in financial returns [Støve and Tjøstheim, 2014] and testing for financial contagion [Støve et al., 2014].

In the following, we will provide short introductions to Pearson's ρ, copulas and local Gaussian correlation, inspired by the overview given in [Tjøstheim et al., 2021]. Thereafter, we look closer at some challenges arising in portfolio management due to the asymmetries often observed in asset returns data. The chapter is concluded with a short motivation for using the local Gaussian correlation to improve asset allocation in investment portfolios.

2.1 The Pearson product-moment correlation

The product-moment correlation was first introduced in Francis Galton's hereditary studies [Galton, 1889] as a measure of association in a regression setting. Pearson [Pearson, 1896] later provided a more precise mathematical treatment and introduced the notation ρ for the population value, and r for the empirical estimate. The correlation coefficient is a measure of linear association between two random variables, defined as the ratio of the covariance of the variables over the product of their standard deviations. Consider two random variables X and Y with finite second moments and covariance $\sigma(X,Y) = E(X - E(X))(Y - E(Y))$. The Pearson product-moment correlation is

$$\rho = \rho(X,Y) = \frac{\sigma(X,Y)}{\sigma_X \sigma_Y} \tag{2.1}$$

where $\sigma_X = \sqrt{E(X - E(X))^2}$ is the standard deviation for X and similarly for σ_Y. The correlation coefficient is $-1 \leq \rho \leq 1$. Given some observations $(X_1, Y_1), ..., (X_n, Y_n)$ the empirical estimate r can be obtained by

$$r = \hat{\rho} = \frac{\sum_{j=1}^{n}(X_j - \bar{X})(Y_j - \bar{Y})}{\sqrt{\sum_{j=1}^{n}(X_j - \bar{X})^2}\sqrt{\sum_{j=1}^{n}(Y_j - \bar{Y})^2}} \tag{2.2}$$

with $\bar{X} = \frac{1}{n}\sum_{j=1}^{n} X_j$ and similarly for $\bar{Y}$. Pearson's ρ take values between -1 and 1, and $|\rho| = 1$ imply the relationship between X and Y is described perfectly with a linear equation. $\rho > 0$ show a positive linear association between the variables, i.e. as X increase, so will Y. Similarly, $\rho < 0$ describe a negative one. When the correlation take the value 0, there is no linear dependency between the variables.

Pearson's ρ is widely used despite it's known shortcomings. There are several possible reasons for this. First, the empirical correlation coefficient r is straightforward to calculate, making it accessible to a large audience. It also has a close relationship to the widely used linear regression models. The slope of a regression line describing the relationship between two variables X and Y is proportional to ρ, hence the empirical estimate r appears as a part of linear least squares. Furthermore, in cases where the variables actually *can* be modeled by a bivariate Gaussian distribution, ρ give a complete characterization of the dependence structure between X and Y. In such cases, the variables are deemed *independent* if and only if $\rho = 0$. In many situations, the Gaussian is also a reasonable approximation. Moreover, the concept of correlation is straightforward to extend to the multivariate case, where the entire dependence structure is determined by pairwise dependencies in the correlation matrix. Finally, we can also without much effort expand the concept of correlation to a time series $\{X_t\}$. Assuming stationarity and existence

of second moments, the *autocorrelation* function is given by $\rho(t) = \rho(X_{t+s}, X_t)$ for arbitrary integers s and t. The dependence structure is completely determined by $\rho(t)$ for Gaussian time series.

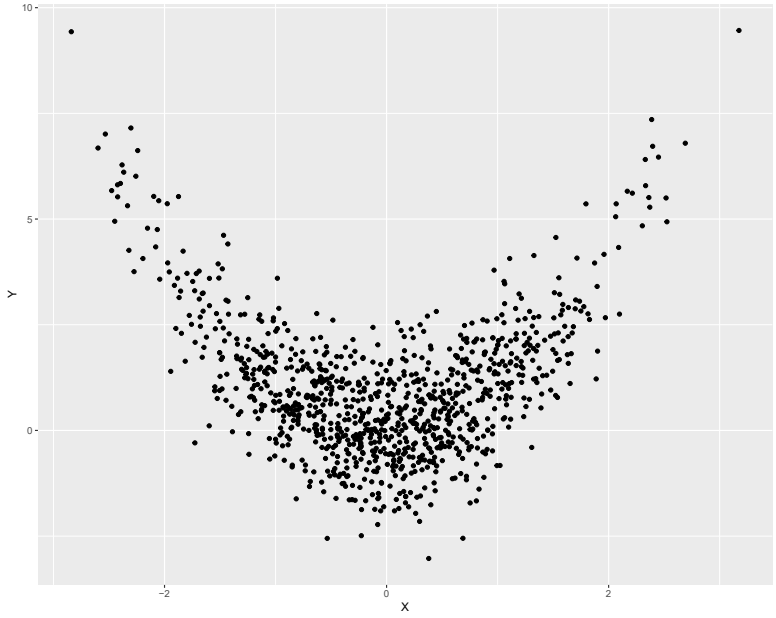

Figure 2.1: Example of nonlinear dependency $Y = X^2 + \epsilon$

There are some serious weaknesses of Pearson's ρ. When moving beyond the Gaussian distribution, the equivalence between uncorrelatedness and independence is generally not true. Uncorrelated variables may still be dependent. Moreover, the empirical correlation estimate r is not robust. If outliers are present in the sample used for estimation, they can have a large impact on the perceived level of dependence in the data. Even a few outliers could dramatically change the outcome. There are attempts to remedy this by relying on ranks when calculating the correlation. Both Spearman's rank correlation [Spearman, 1904] and Kendall's τ rank correlation [Kendall, 1938] provide more robust results, but they do share the vulnerability of Pearson's ρ when dependencies in the data are *non-linear*. To illustrate, we have included the classical example of $Y = X^2 + \epsilon$, where X and ϵ are both standard normal. The plot in Figure 2.1 is created with 1000 random samples. We calculate the three measures of dependence using the simulated data. For reference, the sample estimate for Spearman's rank correlation is given by

$$\hat{\rho}_S = \frac{n^{-1} \sum_{i=1}^{n} R_{i,X}^{(n)} R_{i,Y}^{(n)} - (n+1)^2/4}{(n^2-1)/12} \tag{2.3}$$

where $R_{i,X}^{(n)}$ is the rank of X_i amongst the n observations in $\{X_1, ..., X_n\}$, and similarly for $R_{i,Y}^{(n)}$. Hence Spearman's $\hat{\rho}_S$ is merely the Pearson correlation coefficient between the rank variables, and $-1 \leq \hat{\rho}_S \leq 1$.

Kendall's rank-based measure τ is calculated in a different manner. Consider the set of n pairs (X_i, Y_i) of the variables X and Y. If the ranks of both elements in two selected pairs (X_i, Y_i) and (X_j, Y_j), $i \neq j$ agree, these two pairs of observations are defined as *concordant*. Concordance require that either both $X_i < X_j$ and $Y_i < Y_j$, or $X_i > X_j$ and $Y_i > Y_j$. In the case that $X_i < X_j$ and $Y_i > Y_j$, or $X_i > X_j$ and $Y_i < Y_j$, the pairs are *discordant*. Now, the estimate for Kendall´s rank correlation τ is given by

$$\begin{aligned} \hat{\tau} &= \frac{\text{(concordant pairs) - (discordant pairs)}}{\text{all pairs}} \\ &= \frac{2}{n(n-1)} \sum_{i<j} \operatorname{sgn}(X_j - X_i) \operatorname{sgn}(Y_j - Y_i) \end{aligned} \tag{2.4}$$

where $-1 \leq \hat{\tau}_S \leq 1$. The measures calculated from the simulated data are $\hat{\rho} = -0.0305$, $\hat{\rho}_S = -0.0068$ and $\hat{\tau} = -0.0081$ for Pearson, Spearman and Kendall, respectively. None of these could be said to be significantly different from zero when performing the hypothesis testing, despite the imposed relationship $Y = X^2 + \epsilon$. Asymmetric dependence is often present in data, in particular in the domain of finance [Campbell et al., 2002], [Garcia and Tsafack, 2011]. Returns on assets in a stock portfolio may have a certain dependence structure under normal market conditions, but this can change fundamentally in a period of turmoil, such as the 2008 financial crisis. In the following we will present some alternative approaches to modeling the association between random variables, which are better suited for our purpose.

2.2 Some alternative dependence measures

2.2.1 The copula

A copula is a multivariate distribution whose marginals are all uniform over $[0, 1]$. For a p-dimensional random vector U on the unit cube, a copula C is

$$C(u_1, ..., u_p) = P(U_1 \leq u_1, ..., U_p \leq u_p) \tag{2.5}$$

Sklar's theorem [Sklar, 1959] states that a multivariate model $F(x_1, ..., x_p)$ with marginal cumulative distributions $F_i(x_i), i = 1, ..., p$ can be written as

$$F(x_1, ..., x_p) = C\big(F_1(x_1), ..., F_p(x_p)\big) \tag{2.6}$$

where $C(u_1, ...u_p)$ is a distribution function over the unit cube $[0, 1]^p$ that decodes all the interdependence structure of $F(x_1, ..., x_p)$. The joint distribution is decomposed into univariate marginal distribution functions $F_i(x_i)$ and the copula $C(u_1, ...u_p)$, and these can be estimated separately. This structure provide a high degree of flexibility. Competing models for the marginals may be evaluated against each other, and there are a variety of parametric copula families available, each with parameters that govern the strength of dependence. Copulas have been widely used in applications such as risk management [Embrechts et al., 2001], portfolio optimization, and insurance [Bouyé et al., 2000], [Frees and Wang, 2005], [Frees et al., 1996]. The choice of $C(u_1, ...u_p)$ relies on the specific problem and the type of dependence that needs to be modeled. Elliptical copulas and Archimedean copulas are two commonly used families of copulas in finance.

An elliptical copula is a copula that corresponds to an elliptical distribution, such as the Gaussian or Student's t-distribution, via Sklar's theorem. As an example, consider the bivariate Gaussian case. The Gaussian copula can be written as

$$C(u_1, u_2) = \Phi_\rho \left(\Phi^{-1}(u_1), \Phi^{-1}(u_2)\right) \tag{2.7}$$

where Φ^{-1} is the inverse cumulative distribution function of a standard Gaussian and Φ_ρ is the joint cumulative Gaussian distribution with mean vector equal to zero and correlation $-1 \leq \rho \leq 1$. A benefit of this specification is that different degrees of correlation between the marginals may easily be specified via ρ. On the negative side, elliptical copulas does typically not have simple closed-form solutions, and they are limited by the restriction to have radial symmetry. Still, due to its ease of implementation, it has been used to solve many tasks within finance, such as asset allocation and risk management. Several studies have pointed out a potentially large weakness, though: It can be prone to underestimate dependencies in the tail during extreme events.

See for example [Zimmer, 2012] for an analysis of the effectiveness of the Gaussian copula when estimating ratings for Collaterized Debt Obligations during the financial crisis of 2008. The study by [Malevergne and Sornette, 2003] report to have found that dependence structure between most pairs of currencies and major stocks are compatible with the Gaussian copula. The authors cannot, however, find support for the Gaussian copula hypothesis when they analyse commodities, and issue a general warning against

trusting the approach blindly. The t copula [Demarta and McNeil, 2005] is also part of the family of elliptical copulas. The bivariate distribution may be specified as

$$C(u_1, u_2) = t_{\rho,\nu}\left(t_\nu^{-1}(u_1), t_\nu^{-1}(u_2)\right) \tag{2.8}$$

where $t_{\rho,\nu}$ is the joint cumulative distribution function, t_ν^{-1} is the inverse of a standard Student-t, ν the degrees of freedom and ρ is a correlation parameter. As the degrees of freedom ν increases, the t copula converges towards the normal. For smaller values (under 30, approximately) it has more mass in the tails, which may be preferable when modeling financial data. It does, however, also suffer from the limitation of being symmetric. A further treatment of elliptical distributions and copulas can be found in [Tjostheim et al., 2021] and [Fang et al., 2018].

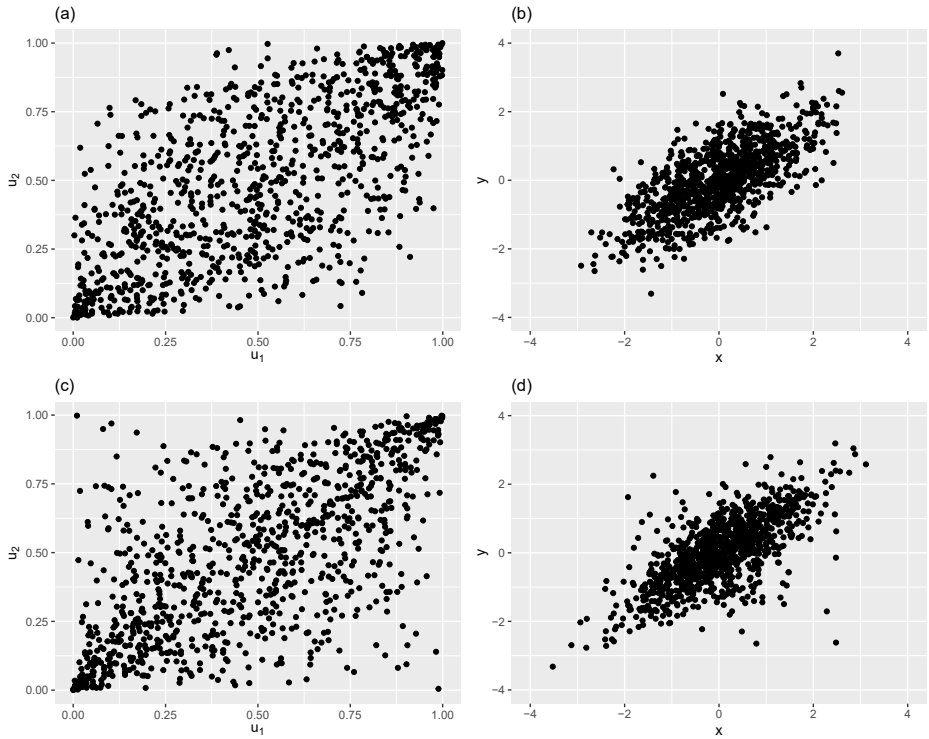

Figure 2.2: 1000 observations simulated from (a) Gaussian copula and (c) t copula with $\rho = 0.7$ and $\nu = 3$ using standard normal for marginals. Right column shows the corresponding bivariate distributions for (b) Gaussian and (d) Student t.

As shown in the simulations in Figure 2.2, the t copula has a slightly tighter formation in the two corners compared to the Gaussian, due to the stronger tail dependency.

In contrast to the elliptical copulas that are constructed implicitly from well-known distributions, models in the Archimedean family [Genest and MacKay, 1986], can be constructed explicitly using a generator function. These copulas take the form

$$C(u_1, ..., u_p) = \varphi^{-1}\big(\varphi(u_1) + ... + \varphi(u_p)\big) \tag{2.9}$$

where φ^{-1} is the inverse of the generator function φ, which maps the unit interval to itself. In order for ((2.9)) to be a copula, the generator needs to be continuous, strictly decreasing and convex function, which is p-monotone. See [Nelsen, 2007] for a deeper treatment of various copulas and their generators. Archimedean copulas have closed-form expressions for their densities and cumulative distribution functions, and they can be easily estimated and simulated from. See for example [Genest and Rivest, 1993] for a thorough introduction to statistical inference procedures for bivariate models of this class. Archimedean copulas are flexible and can model a wide range of dependence structures. Some popular examples are the Gumbel [Gumbel, 1960], Clayton [Clayton, 1978], Frank [Frank, 1979] copulas. For the sake of illustration, we move on with a presentation of the bivariate specifications of these. They can be be extended to multivariate versions via the generator function.

The Gumbel copula [Gumbel, 1960] is an asymmetric copula, which means that it exhibits different levels of dependence in the positive and negative tails. For this model, the degree of dependence is stronger in the upper tail, making it a candidate for modelling extreme events. The copula is given by

$$C_\theta^{Gu}(u_1, u_2) = \exp\Big[- \big[(-\ln(u_1))^\theta + (-\ln(u_2))^\theta\big]^{\frac{1}{\theta}}\Big] \tag{2.10}$$

where the generator function is $\varphi(u) = (-\ln(u))^\theta$. The parameter $\theta \in [1, \infty)$ controls the strength of the dependence, and positive values of θ indicate greater positive tail dependence. As the value of θ increases, the positive tail dependence increases and the negative tail dependency is reduced.

The Clayton copula [Clayton, 1978] is also asymmetric, and often used to model situations where there is greater dependence in the negative tail than in the positive. The bivariate version is defined as:

$$C_\theta^{Cl}(u_1, u_2) = \big(u_1^{-\theta} + u_2^{-\theta} - 1\big)^{-\frac{1}{\theta}} \tag{2.11}$$

where the generator function is $\varphi(u) = \frac{1}{\theta}(u^{-1} - 1)$. The parameter $\theta \in [-1, \infty)\backslash\{0\}$ controls the degree of dependence. The Clayton copula is commonly used in economics and finance to model asymmetric dependence structure between financial asset returns.

Finally, the bivariate Frank copula [Frank, 1979] is a symmetric copula, given by:

$$C_\theta^{Fr}(u_1, u_2) = -\frac{1}{\theta} \ln \left[1 + \frac{(e^{-\theta u_1} - 1)(e^{-\theta u_2} - 1)}{e^{-\theta} - 1} \right] \tag{2.12}$$

with generator function is $\varphi(u) = -ln\left(\frac{e^{-\theta u}-1}{e^{-\theta}-1}\right)$ and $\theta \in (-\infty, \infty) \backslash \{0\}$. As the value of θ increases, the Frank copula has more prominent tails and higher density along the diagonal from the lower left to the upper right corner. Simulation results illustrating the properties of the Gumbel, Clayton and Frank copulas can be found in Figure 2.3.

When modelling these bivariate Archimedeans, only one copula parameter was needed to capture the dependence structure. The choice of the copula and θ determines the shape of the distribution and the strength and type of dependence between the variables. In contrast, the general case with multiple variables requires a multivariate copula, which models the dependence structure between all the variables. Archimedean copulas, while they can be constructed in any dimension, are limited by the few parameters they have, making them too rigid to be useful in many cases. Elliptical copulas can be utilized when the observations follow an elliptical distribution, but they are not suitable when the data departs from symmetry or has prominent one-sided tail dependency.

Vine copulas [Bedford and Cooke, 2001], [Bedford and Cooke, 2002] are flexible and powerful tools for modeling multivariate dependencies in a more sophisticated manner, as they allow complex dependence structures to be represented by combining multiple bivariate copulas in a tree-like structure. This structure, known as a vine copula, can capture complex interdependencies between variables by modeling them in stages, starting from pairwise relationships and then working towards more complex relationships. In [Aas et al., 2009], pair-copula constructions are used to model multivariate asset returns exhibiting complex tail dependencies, by decomposing the distribution into a set of pair copulas that are applied to the original variables and to their conditional and unconditional distribution functions. [Low et al., 2018] investigate potential benefits of applying such *canonical vines* in the context of portfolio management. Using minimization of the Conditional Value-at-Risk (CVaR) as main objective, they construct portfolios ranging from 3-12 constituents, and allow weights to be determined without any short-sales constraints. Such an approach may be complex to implement, but they find that the selected Clayton canonical vine produces consistently better results, and conclude the extra effort is "worth it". But is the extra complexity *required*? Could the classical methods be improved by using simpler, more easy-interpretable techniques, such as the local Gaussian correlation?

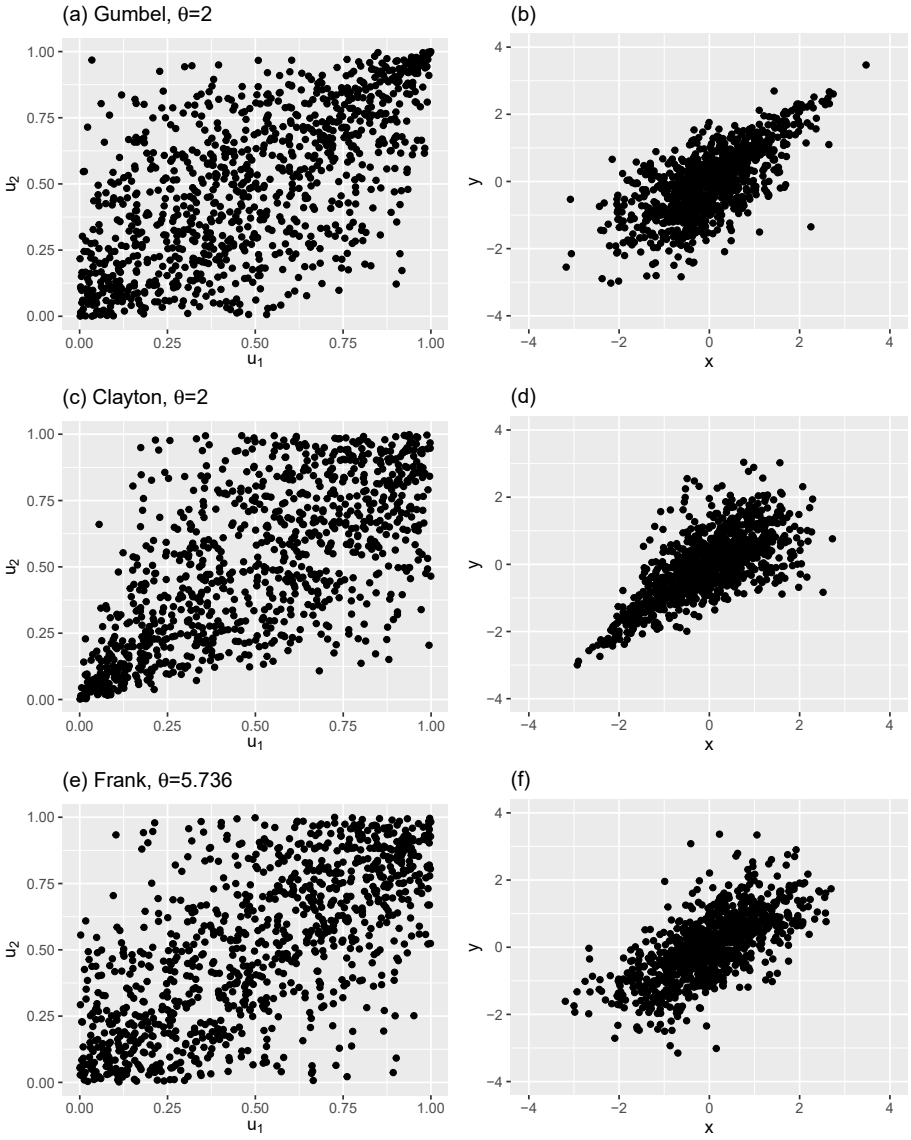

Figure 2.3: 1000 observations simulated for archimedean copulas with θ corresponding to $\rho = 0.7$ Left column shows the (a) Gumbel, (c) Clayton and (e) Frank copulas. The corresponding bivariate distributions are displayed in the right column for (b) Gumbel, (d) Clayton and (f) Frank. Standard normal distributions were used for the marginals.

2.2.2 Local Gaussian correlation

The local Gaussian correlation was introduced in [Tjøstheim and Hufthammer, 2013]. We will briefly present the core concepts here, and refer the interested reader to **Paper A** with supplementary materials. We would also like to note that this section provide a standard introduction to local Gaussian correlation, and will to some extent follow the presentation in Tjøstheim et al. [2021]. Further details regarding how the method can be implemented in the R programming language can be found in [Berentsen et al., 2014] and [Otneim, 2021].

For illustrative purposes, we will focus on the bivariate case. Let $X = (X_1, X_2)$ be a random variable with density $f(x) = f(x_1, x_2)$. The local Gaussian correlation can be calculated by approximating f locally in each point $x = (x_1, x_2)$ using a Gaussian bivariate density $\psi_x(v)$ where where $v = (v_1, v_2)$ is the running variable. The approximating density at this specific point x is given by

$$\begin{aligned}\psi_x(v, \theta(x)) =& \frac{1}{2\pi\sigma_1(x)\sigma_2(x)\sqrt{1-\rho^2(x)}} \times \exp\Big[-\frac{1}{2}\frac{1}{1-\rho^2(x)}\Big(\frac{(v_1-\mu_1(x))^2}{\sigma_1^2(x)} \\ & - 2\rho(x)\frac{(v_1-\mu_1(x))(v_2-\mu_2(x))}{\sigma_1(x)\sigma_2(x)} + \frac{(v_2-\mu_2(x))^2}{\sigma_2^2(x)}\Big)\Big]\end{aligned}$$

where the parameter vector

$$\theta(x) = \big(\mu_1(x), \mu_2(x), \sigma_1^2(x), \sigma_2^2(x), \rho(x)\big)$$

contains the local means $\mu_1(x), \mu_2(x)$, the local standard deviations $\sigma_1(x), \sigma_2(x)$ and the local correlation $\rho(x)$. When moving to another point x^*, a new local density with parameters $\mu_1(x^*), \mu_2(x^*), \sigma_1^2(x^*), \sigma_2^2(x^*), \rho(x^*)$ is estimated. The population parameter vector is obtained by minimizing the local penalty function measuring the difference between f and ψ

$$q = \int K_b(v-x)[\psi(v, \theta(x)) - \ln\{\psi(v, \theta(x))\}f(v)]\mathrm{d}v \tag{2.13}$$

where $K_b(v-x) = (b_1 b_2)^{-1} K_1(b_1^{-1}(v_1 - x_1))K_2(b_2^{-1}(v_2 - x_2))$ is a product kernel with bandwidths $b = (b_1, b_2)$. As in [Hjort and Jones, 1996], we interpret q as a locally weighted Kullback-Leibler distance from f to $\psi(\cdot, \theta(x))$, and therefore require that the minimizer $\theta_b(x)$ (also depending on K) needs to satisfy

$$\int K_b(v-x)\frac{\partial}{\partial\theta_j}[\ln\{\psi(v, \theta(x))\}f(v) - \psi(v, \theta(x))]\mathrm{d}v = 0, \quad j = 1, \ldots, 5. \tag{2.14}$$

We move further in two steps. First, we define the population value $\theta_b(x)$ as the minimizer of (2.13), assuming that there is a unique solution to (2.14). Second, we let $b \to 0$ and consider the limiting value $\theta(x) = \lim_{b\to 0} \theta_b(x)$. In [Tjøstheim and Hufthammer, 2013], this is done via Taylor expansion arguments. Here, we will proceed under the assumption that a limiting value $\theta(x)$ independent of b and K exists. We need to use the neighbourhood of x and a finite bandwith b when estimating $\theta(x)$ and $\theta_b(x)$, as in nonparametric density estimation. The parameters are then obtained via maximum likelihood. Given observations $\{X_1, \ldots, X_n\}$, the local likelihood function is

$$L(X_1, \ldots, X_n, \theta(x)) = n^{-1} \sum_i K_b(X_i - x) \log \psi(X_i, \theta(x)) - \int K_b(v - x)\psi(v, \theta(x))\mathrm{d}v. \tag{2.15}$$

The likelihood reduces to the ordinary global likelihood when $b \to \infty$, since the last term has 1 as its limiting value. This last term is of high importance, as it ensures that $\psi(x, \theta_b(x))$ is not allowed to stray far away from $f(x)$ as $b \to 0$. Further details regarding asymptotic behaviour and estimation can be found in [Tjøstheim and Hufthammer, 2013] and the supplementary material of **Paper A**.

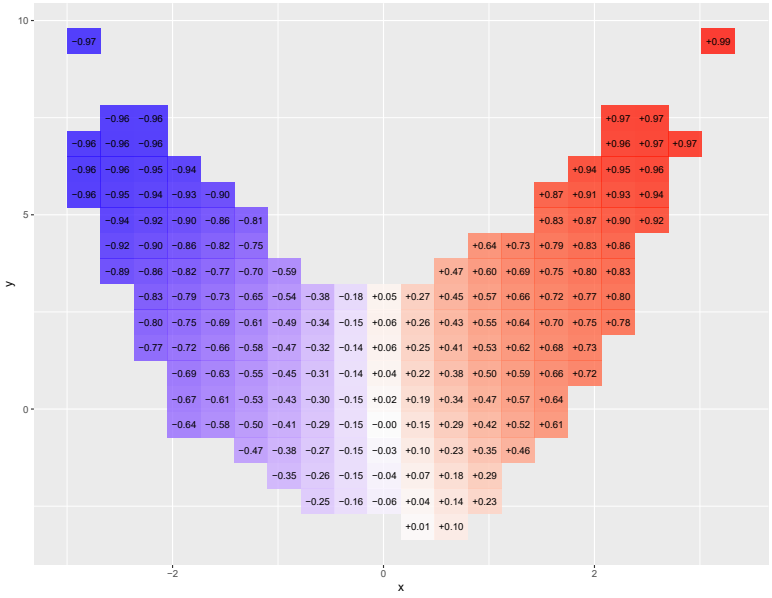

Figure 2.4: Local correlation plot $Y = X^2 + \epsilon$

In Figure 2.4 we revisit the bivariate nonlinear example from Figure 2.1. Here, we have defined a grid of 20×20 equally sized cells covering the range of X and Y, and calculated

the local correlations for cells containing observations. A bandwidth of $b = (\frac{11}{10}\sigma_X, \frac{11}{10}\sigma_Y)$ was applied in the estimation procedure using the R package **lg** described in [Otneim, 2021]. The observed estimates, such as $\hat{\rho}(-2, 5) = -0.94$ and $\hat{\rho}(2, 5) = 0.95$, are simple to interpret, and the chart provides a clear summary of the dependence structure.

In a setting with a sample $X_i = \{X_{1i}, \ldots, X_{pi}\}$, $i = 1, \ldots, n$ with dimensions $p > 2$, it is possible to extend the method and construct local multivariate distributions. This would provide local $(p \times p)$ correlation matrices and $(p \times 1)$ vectors for mean and variance at each evaluation point, via maximum likelihood estimation. However, as the number of dimensions p in the sample increases, the precision of the estimates produced by the estimation procedure will decline, due to the curse of dimensionality. A pragmatic approach to solving this challenge is to break it down into smaller, bivariate problems. Instead of considering all coordinates at once in a p-variate problem, each local correlation is estimated by only considering the corresponding pair of observation vectors. This results in a series of bivariate problems that estimate pairwise local correlations based on their respective pairs of coordinates. The estimated local correlations are gathered in a $(p \times p)$ matrix. If the corresponding local covariance matrix obtained from the bivariate problem is not positive definite, the method described in [Higham, 2002] is used to adjust it. This can easily be implemented, for example with the `nearPD()` function from the R package **Matrix**.

The local Gaussian correlation has been applied in various studies related to finance, such as the analysis of dependence structure between asset returns [Støve and Tjøstheim, 2014] and in testing for financial contagion [Støve et al., 2014], [Bampinas and Panagiotidis, 2017], [Nguyen et al., 2020]. While it has been successful in understanding the relationship between assets, it has not yet been applied to the process of making investment decisions and constructing portfolios. In the following, we will explore how the local correlation matrices may be utilized within the classical mean-variance framework introduced by Harry Markowitz [Markowitz, 1952].

2.3 Portfolio allocation under asymmetric dependence in returns

Investors use portfolios as a strategy to spread their investments across multiple assets in order to minimize risk. By allocating wealth among different assets, the negative impact of any single asset's poor performance on the overall portfolio may be offset by changes in the other constituents. This diversification effect is strongly dependent on the associations between the instruments in the portfolio, and how these change during differing circumstances. Some common asset classes that are considered include equities, fixed

income, commodities, currencies, cash equivalents, real estate and infrastructure. Institutional investors will typically face restrictions regarding what asset classes, companies and countries that qualify for investment. One example is the the *Norwegian Government Pension Fund Global* managed by Norges Bank Investment Management, which needs to comply with the exclusion list and policies overseen by Norges Bank's Executive Board, [Chambers et al., 2012]. After considering such restrictions, the portfolio manager is left with an *investment universe* where funds may be allocated.

Investors may use various methods to decide on the most suitable mix for a portfolio with their desired level of risk and return. There is a large research literature on *factor models* for asset returns, dating back to the CAPM [Treynor, 1961], [Sharpe, 1964], [Lintner, 1965], [Mossin, 1966]. The factor models developed by Fama-French [Fama and MacBeth, 1973], [Fama and French, 1993] are other examples of key contributions. In this school of thought, asset returns are explained by focusing on underlying factors, such as company size and stock price compared to fundamental characteristics such as earnings, dividends, and book value. Later studies, such as [Fama and French, 2015], have suggested a number of other factors to be considered. The use of standard deviation for portfolio returns as risk measure has also been questioned on numerous occasions. While it can be a useful tool for measuring the overall volatility of a portfolio's returns, it has some limitations when it comes to capturing the *downside* risk that investors are primarily concerned with. Measures, like as *Value at Risk* and *Expected Shortfall*, [Duffie and Pan, 1997], [McNeil et al., 2015], [Tasche, 2002] may be be used as alternatives in the asset allocation setting. These also have their own strengths and weaknesses.

Here, we will adopt the classical formulation for asset allocation from [Markowitz, 1952], which consists of optimising a target objective, given risk preferences and operational constraints on the asset allocation. We assume there are N risky assets, and use volatility of portfolio returns as risk measure. The returns on the risky assets are denoted by $\mathbf{r}_t \in \mathbb{R}^N$, which are assumed to have expected values $\mu_\mathbf{t} \in \mathbb{R}^N$, and covariance matrix $\mathbf{\Sigma_t} \in \mathbb{R}^N \times \mathbb{R}^N$. Further, let $\mathbf{w_t} \in \mathbb{R}^N$ be the unknown vector of optimal portfolio weights at time t. The investor may choose weights $\mathbf{w_t}$ to either i) maximize expected return subject to a constraint on portfolio variance, or ii) minimize portfolio variance subject to a constraint on expected return. Taking risk preferences into account, the objectives may be expressed as maximisation problems using utility functions

$$\max_{\mathbf{w_t}} U_1 = \mathbf{w_t^T}\mu_\mathbf{t} - \frac{\gamma}{2}\mathbf{w_t^T}\mathbf{\Sigma_t}\mathbf{w_t} \tag{2.16}$$

$$\max_{\mathbf{w_t}} U_2 = -\mathbf{w_t^T}\mathbf{\Sigma_t}\mathbf{w_t} \tag{2.17}$$

where the parameter γ represents the investor's degree of risk aversion under quadratic

utility. Various operational constraints can be applied to the weights. The *full investment* $w_1 + \cdots + w_N = 1$ condition is often used. Another type of constraint is related to *long only* positions, which specify that we can only buy shares and therefore only have positive weights, in contrast to the case of short positions, in which the selling positions would be reflected as negative weights. The full optimisation problem for a long-only investor aiming to maximise returns under a constraint on portfolio volatility σ_p is

$$\begin{aligned} \max_{\mathbf{w_t}} \quad & \mathbf{w_t^T}\mu_\mathbf{t} - \frac{\gamma}{2}\mathbf{w_t^T}\mathbf{\Sigma_t}\mathbf{w_t} \\ \text{s.t.} \quad & \mathbf{w}^T\mathbf{\Sigma}\mathbf{w} \leq \sigma_p^2 \\ & \mathbf{w}^T\mathbf{1} = 1 \\ & \mathbf{w} \geq 0 \end{aligned} \tag{2.18}$$

where $\mathbf{w}^T\mathbf{1} = 1$ also dictate full investment. The minimum-variance investor with similar weight constraints and target return R_p^* for the portfolio needs to solve

$$\begin{aligned} \max_{\mathbf{w_t}} \quad - \quad & \mathbf{w_t^T}\mathbf{\Sigma_t}\mathbf{w_t} \\ \text{s.t.} \quad & \mathbf{w}^T\mu = R_p^* \\ & \mathbf{w}^T\mathbf{1} = 1 \\ & \mathbf{w} \geq 0 \end{aligned} \tag{2.19}$$

For a range of different risk aversion levels, expressed via γ, the optimization is expected to produce corresponding optimal portfolios with a trade-off between return and volatility. There are, however, a number of assumptions made by the Markowitz model that are not met. In financial markets, investors do not necessarily behave in a rational manner, they might have incomplete information, and limited opportunity to borrow an lend at the risk-free rate. The markets can also be inefficient, e.g. prices of financial assets might not fully reflect all publicly available information. Many studies have also shown that the distribution of financial returns is often fat-tailed and asymmetric, meaning that they are not normally distributed, and correlation between returns can change depending on market conditions [Patton, 2004]. In [Cont, 2001], a list of stylized facts and empirical properties of asset returns are summarized. There, and in similar studies, some of the typical findings are that i) returns are not independent identically distributed, ii) distributions are heavy tailed and asymmetric, iii) volatility varies over time and extremes appear in clusters, iv) volatility is mean-reverting, and v) dependencies between asset returns are asymmetric.

This makes risk management via diversification difficult, as the portfolio is more vulnerable to market downturns than it is assumed in the standard mean-variance framework. There have been many attempts to improve the workhorse of modern portfolio theory.

One significant contribution worth mentioning is the Black-Litterman asset allocation model [Black and Litterman, 1992], which combines the traditional Markowitz-approach with Bayesian methods for updating investor views on future asset performance. Other approaches have focused on incorporating more advanced methods for modelling returns and measuring risk into the original framework. Copula models, in particular, have been popular tools for this purpose. In [Patton, 2004], copulas are applied to construct models of time-varying dependence structures that allow for different dependence during bear bull markets. These are used to allocate portfolio weights for an investor with constant relative risk aversion. They find that the method may lead to higher gains when short-sales are permitted. In the study by [Low et al., 2018], canonical vine copulas are used for modelling the dependence structure for an investor aiming to minimize expected shortfall, also without short sales constraints. When comparing elliptical and asymmetric models, they find that the Clayton canonical vine copula consistently produces the highest-ranked outcomes, and conclude that the extra complexity is worthwhile.

The choice of copula to use in modeling the dependencies between financial returns can be overwhelming for non-technical asset managers. Complex models can also introduce additional assumptions and reduce transparency. Could model complexity represent a risk in itself? The use of local Gaussian correlation is a simpler approach compared to many of the copulas suggested in the literature. By modeling the correlation structure of asset returns at a local level, the local Gaussian correlation approach can effectively capture complex and dynamic dependencies between assets and report it in a familiar way. The local correlation matrices can be applied directly in the mean-variance framework, and results may be evaluated and compared with the traditional approach in a transparent way. Further details in this regard can be found in **Paper A**.

Chapter 3

Energy market risk management

The markets for electricity and natural gas are complex and dynamic, where the prices of these commodities are determined by the balance of supply and demand. They are characterized by high volatility and price fluctuations, driven by a variety of factors such as weather conditions, cost of production, supply disruptions, limited transmission and storage capacity, and changes in demand. Trading takes place both through organized exchanges such as the European Energy Exchange and Nasdaq OMX Commodities, and over-the-counter (OTC). The market is divided into spot and forward markets. In the spot market, physical energy is traded for near-term delivery. In contrast, the forward market is a market where electricity and natural gas is traded for delivery at a later date, and prices are based on expectations of future supply and demand. Forward markets allow market participants to hedge against price risk, by buying or selling contracts at a fixed price today, for delivery in the future. The interested reader may find a more detailed overview of energy market structure, available contracts and the related markets for fuel, freight and weather products in [Eydeland and Wolyniec, 2002], [Benth et al., 2008] and [Kirschen and Strbac, 2018].

The large spot price variations leaves both consumers and producers of energy with a significant price risk, and lack of attention and risk mitigation might lead to severe financial stress. When energy prices are high, it can have consequences for consumers and society as a whole. This is certainly the case in Europe at present (2022-2023), where demand is is outstripping energy supply, causing an unprecedented price risk crisis. Some of the main causes are the rebound after the COVID-19 pandemic and changes on the supply side due to the transition to renewable energy production. The problem is magnified by the Russia-Ukraine conflict, which has impacted the flow of Russian natural gas. For individuals, high energy prices can create financial strain, as they may need to spend more money on utility services such as heating and electricity.

This can be especially difficult for low-income households, who may struggle to afford these necessities. High energy prices can also have a negative impact on businesses, as they may need to pay more for energy to power their operations. This can lead to increased costs, which may in some cases be passed on to consumers in the form of higher prices. Additionally, high energy prices can impact the overall economy, as they can increase inflation and slow economic growth.

A corporate market player with a long term strategy for energy procurement can implement measures to protect against financial risk. Since the price risk management need to take place prior to the physical delivery, we will have our main focus on the contracts in the forward market. We will however start out with a brief presentation of some European electricity spot price markets in order to motivate the need for hedging. Next, an introduction to the forward market and the maximum smoothness forward curve is provided. We continue with a demonstration of the forward curve calculation implemented in the **etrm** package. Thereafter, we discuss some key concepts within energy market risk management, such as the energy portfolio, the management authorisation and hedging strategies. The chapter is concluded with a demonstration of price hedging with the **etrm** package using a data set with Nordic electricity futures contracts listed at Nasdaq OMX Commodities.

3.1 The spot market

Energy spot markets are markets where physical quantities are bought and sold in a short-term timeframe prior to delivery. These markets allow buyers and sellers to trade electricity and natural gas on a near-term basis, typically day-ahead or close to real-time. The prices are determined by the balance of supply and demand, and they can fluctuate based on a number of factors, including the availability of generation resources, transmission constraints, and weather conditions. Prices are highly influenced by seasonal temperature variations. During the summer months when temperatures are higher, the demand for electricity for air conditioning will typically increase, leading to an increase in prices. Conversely, during the winter months, when temperatures are lower, the demand for heating increases leading to an increase in prices. During the day, electricity prices tend to be higher during office hours when businesses and households are using more power, and lower during the night and weekends when demand is reduced.

In the following, we will have a closer look at electricity spot price data from three geographical regions in Europe; The Netherlands, Germany/ Austria and the Nordics. The spot prices are calculated for each hour the following day, determined by the inter-

section of the aggregated supply and demand curves representing all bids and offers in the day-ahead auction at the power exchanges EPEX Spot and Nord Pool. The Physical Electricity Index (Phelix) refers to the base load price index published daily on the power spot market for the German/Austrian market area by EPEX SPOT. This is used as the underlying asset for the EEX Phelix Future listed on the European Energy Exchange (EEX). A similarly reference price is calculated on the same exchange for the Netherlands, the Amsterdam Power Exchange Index (APX). Nordic market participants enter their bids at Nord Pool, who is responsible for the calculation of the Nordic System Price (NPSYS). This is used as settlement reference for the derivatives traded at Nasdaq OMX Commodities.

In the absence of grid bottlenecks, the indices published by spot exchanges are the commodity prices large consumers or suppliers of energy will have to pay to procure the volume needed for the next day, if they participate directly in the wholesale market. Similarly, a producer will be able to sell generated volume at these prices. In practise however, grid congestion is often present. The physical market in a specific region may therefore be divided into price zones due to physical limitations with respect to transmission of electric power. This creates a *basis risk* for market participants, as the local power price exposure may not be fully hedged with contracts that are settled against the index published by the spot exchange.

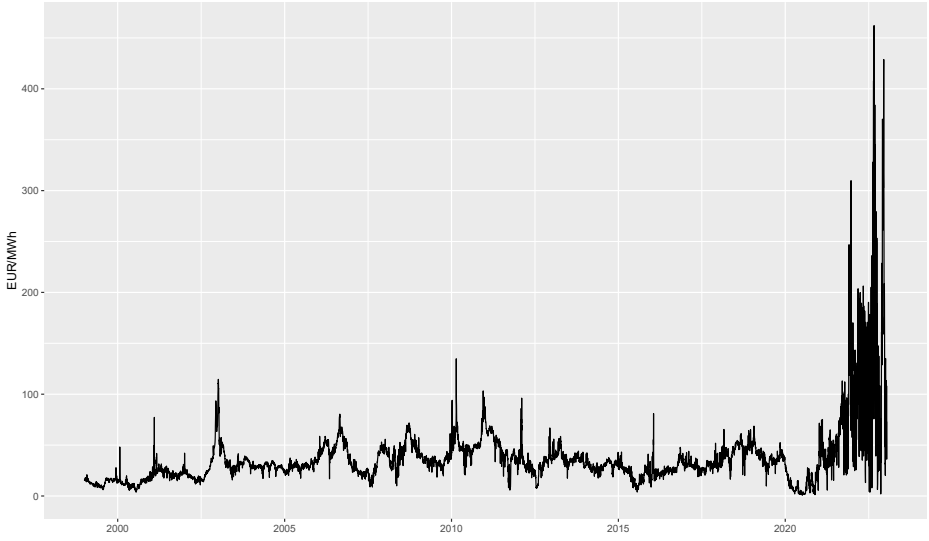

Figure 3.1: Daily average Nordic System Price (NPSYS)

The data presented here consists of daily averages of the indices described above. When we look at the Nordic System Price curve in Figure 3.1, the typical spot price char-

acteristics with high volatility and sudden spikes followed by mean reversion is clearly present. The recent period has seen unprecedented volatility and sharp price increases, which have caused the majority of prices to be concentrated in a small range on the chart. This produces a somewhat distorted view of the volatility prior to the ongoing European energy price risk crisis. In previous years, factors such as low reservoir levels, high fuel costs and cold spells have led to prices that were considered challenging for many companies with large energy consumption.

	APX	Phelix	NPSYS	r_{APX}	r_{Phelix}	r_{NPSYS}
Observations	8,782	8,782	8,782	8,781	8,781	8,781
Mean	55.195	49.178	37.406	0.0001	-0.001	0.0002
Std.dev.	54.227	54.799	32.660	0.290	0.343	0.191
Variance	2,940.530	3,002.956	1,066.668	0.084	0.118	0.036
Skewness	4.880	5.171	5.389	0.744	0.683	0.046
Kurtosis	31.070	34.356	41.977	14.496	12.157	38.267
Min	-5.450	-56.870	0.720	-2.672	-2.948	-2.962
1 Quartile	32.850	27.820	23.363	-0.112	-0.161	-0.045
Median	42.480	37.140	31.330	-0.012	-0.025	-0.006
3 Quartile	54.528	49.810	42.450	0.084	0.115	0.040
Max	693.830	700.220	462.100	3.539	3.569	2.443

Table 3.1: Summary statistics for day-ahead prices and log returns

The summary statistics provide a more detailed view, and here we can identify negative prices both in the APX and Phelix series. The non-storable nature of electricity combined with an increased share of less flexible power generation such as wind contribute to the increased volatility, and in extreme cases negative prices for short periods. When the hourly negative prices are large and/ or prolonged, it will also dominate in the calculation of daily average values. Phelix is the index showing the most extreme values, with daily spot prices of -56.87 EUR/MWh and 700.22 EUR/MWh. For the APX, Phelix and NPSYS we see high standard deviations, positive skew, and fat tails.

We have also calculated the logarithmic price differences in order to explore the price dynamics further. When we did this we had to remove the negative prices from the data set. The fat tails in the price returns are quite striking, and we can see that the log spot prices changes were -295% and $+357\%$ in the most extreme case, which was for the Phelix. These calculations are based on daily data, and longer average spot prices will be smoother. Still, when we look at monthly or even yearly data, there will be considerable variation in the prices. In this thesis we will not focus more on spot markets, but rather used these observations as motivation for hedging the price risk by trading in the forward markets. A rigorous treatment of the spot price dynamics can be found in [Benth et al., 2008].

3.2 The forward market

In this chapter we will have a closer look at the characteristics of some of the contracts that are used in energy forward markets. Thereafter, we show how prices from these standardized products can be used to construct a smooth forward curve able to describe the term structure of the market with daily forward prices. This section also contains a more detailed presentation of the mathematical approach used when generating the *Maximum Smoothness Forward Curve*, see also **Paper B**. Finally we apply the methods and generate forward curves for Nordic electricity contracts listed at Nasdaq OMX Commodities, and comment on how these can be used for non-standard pricing and risk management purposes.

3.2.1 Forward market contracts on a flow delivery

The standardised forwards for electricity and gas are contracts for flow delivery. The underlying commodity is not delivered at a fixed point in time, but over a time interval. Strictly speaking we should think of them as swap contracts. These contracts are without flexibility, with a continuous delivery of electricity or gas over a future time interval. When we describe them, we will follow the procedure in [Bjerksund et al., 2010]. Here, the forward price at time t for one unit of energy delivered at a rate $\frac{1}{(\tau_e-\tau)}$ over the time interval (τ_s, τ_e) is given by $F(t, \tau_s, \tau_e)$, where $t \leq \tau_s \leq \tau_e$. The time distance between τ_s (start) and τ_e (end) for standard contracts in the market may be from a day's length, to week, month, quarter, season or year. A forward contract for a flow delivery may be thought of as the average of hypothetical single-delivery contracts. At time t, each of these would have a unique price $f(t, u)$ for the delivery at u with an infinitesimal delivery period. If we were to receive 1 unit of energy over the interval (τ_s, τ_e), the present value at time t would be

$$V_t\left[\int_{\tau_s}^{\tau_e} f(u,u)du\right] = \int_{\tau_s}^{\tau_e} e^{-r(u-t)} f(t,u)du \tag{3.1}$$

where r represent the risk free rate of interest. The fact that the present value of entering a forward contract should be zero at time t require

$$\int_{\tau_s}^{\tau_e} e^{-r(u-t)} F(t,\tau_s,\tau_e)du = \int_{\tau_s}^{\tau_e} e^{-r(u-t)} f(t,u)du \tag{3.2}$$

This leads to $F(t, \tau_s, \tau_e)$ being the weighted average

$$F(t, \tau_s, \tau_e) = \int_{\tau_s}^{\tau_e} w(r, u) f(t, u) du \tag{3.3}$$

where

$$w(r, u) = \frac{e^{-ru}}{\int_{\tau_s}^{\tau_e} e^{-ru} du} = \frac{re^{-ru}}{e^{-r\tau_s} - e^{-r\tau_e}} \tag{3.4}$$

is the weight function. In the following section we will construct a forward market price curve for the entire horizon T, with all $f(t, u)$ where $t \leq u \leq T$. In the following we will simplify notation and use $f(u)$ for the function describing the forward curve at time t.

3.2.2 The maximum smoothness forward price curve

In this section we will use the methods in [Ollmar, 2003] and [Benth et al., 2007a] to develop a continuous market curve from closing prices observed in the forward market. We will generate the smoothest possible curve able to reproduce the observed closing prices by combining a prior function $\Lambda(u)$ which contain our subjective views on the future prices with an adjustment function $\epsilon(u)$ to take the observed market prices into account. We define a functions smoothness as the mean squared value of its second derivative, and achieve the smoothest possible curve on an interval (τ_s, τ_e) by minimising

$$\int_{\tau_s}^{\tau_e} [e^{''}(u)]^2 du \tag{3.5}$$

The prior function can be a fundamental model or a more simple sinusoidal function to describe the seasonality we observe in energy markets. The smoothing is calculated on $\epsilon(u)$ in order not to disturb the description of the seasonality in $\lambda(u)$. If we choose to exclude the prior, the seasonal price patterns will not be visible in the far end of the curve, where only year or seasonal contracts are available.

Consider a market at time t with m forward contracts available for trading. Let the list

$$S_t = \{(\tau_1^s, \tau_1^e), (\tau_2^s, \tau_2^e), ..., (\tau_m^s, \tau_m^e)\}$$

contain the start and end dates for each of these contracts. The time distance between τ_i^s and τ_i^e for a contract i in $1, .., m$ cover standardized periods such as *week*, *month*, *quarter*, *winter season*, *summer season* and *year*. Some of these settlement intervals

might overlap, and in order to handle this we create a new list of dates $\{t_0, t_1, ..., t_n\}$ to identify each separate sub period, see Figure 3.2. The new list is made by sorting all dates in S_t in ascending order and removing duplicates.

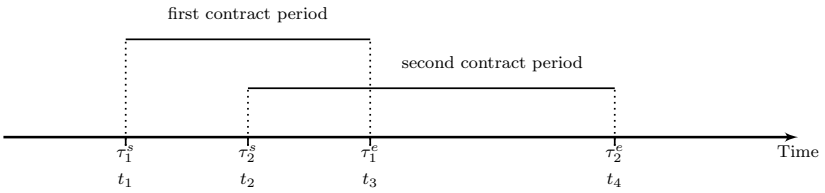

Figure 3.2: Start and end dates for two overlapping contracts

In the following we construct a forward market price curve for the entire horizon using a simplified notation $f(u)$ for the function describing the forward curve at time t. In order to model the strong seasonality in energy markets, the forward curve function is decomposed into two elements:

$$f(u) = \Lambda(u) + \epsilon(u) \qquad u \in [t_0, t_n] \tag{3.6}$$

Following [Ollmar, 2003] we calculate $f(u)$ by combining a prior function $\Lambda(u)$ which contain our subjective views on the future prices with an adjustment function $\epsilon(u)$ to ensure match with the observed closing prices for each of the m contracts. The prior could be generated with a simple sinusoidal function or from a fundamental model more capable of describing the typical seasonality and calendar effects observed in energy markets. Should the prior be excluded, the seasonal price patterns will not be visible in the far end of the curve, where only yearly or seasonal contracts are available. The smoothing is calculated on the adjustment function, we aim to minimize the total curvature of $\Lambda(u)$ while preserving the information from the prior. Here, smoothness is defined as the integral of second-order derivative of the function, and the smoothest possible curve over the interval $[t_0, t_n]$ is achieved by minimising

$$\int_{t_0}^{t_n} [\epsilon''(u)]^2 \, \mathrm{d}u$$

As shown in [Lim and Xiao, 2002], the smoothest possible curve is achieved when each of the n sub periods are modelled by fourth-degree polynomials, and we can write the

adjustment function as a spline

$$\epsilon(u) = \begin{cases} a_1u^4 + b_1u^3 + c_1u^2 + d_1u + e_1, & u \in [t_0, t_1], \\ a_2u^4 + b_2u^3 + c_2u^2 + d_2u + e_2, & u \in [t_1, t_2], \\ . & \\ . & \\ a_nu^4 + b_nu^3 + c_nu^2 + d_nu + e_n, & u \in [t_{n-1}, t_n]. \end{cases}$$

To build the forward curve function, we need to identify the parameters of $\epsilon(u)$

$$\mathbf{x}^\intercal = [a_1\ b_1\ c_1\ d_1\ e_1 \quad a_2\ b_2\ c_2\ d_2\ e_2 \ldots\ a_n\ b_n\ c_n\ d_n\ e_n]$$

by solving the quadratic optimisation problem

$$\min_x \int_{\tau^s}^{\tau^e} [\epsilon''(u;x)]^2\,\mathrm{d}u \tag{3.7}$$

subject to

$$\begin{aligned} (a_{j+1} - a_j)u_j^4 + (b_{j+1} - b_j)u_j^3 + (c_{j+1} - c_j)u_j^2 + (d_{j+1} - d_j)u_j + e_{j+1} - e_j &= 0 \quad \text{(a)} \\ 4(a_{j+1} - a_j)u_j^3 + 3(b_{j+1} - b_j)u_j^2 + 2(c_{j+1} - c_j)u_j + (d_{j+1} - d_j) &= 0 \quad \text{(b)} \\ 12(a_{j+1} - a_j)u_j^2 + 6(b_{j+1} - b_j)u_j + 2(c_{j+1} - c_j) &= 0 \quad \text{(c)} \\ \epsilon'(u_n;x) &= 0 \quad \text{(d)} \\ \frac{1}{\tau_i^e - \tau_i^s} \int_{\tau_i^s}^{\tau_i^e} (\Lambda(u) + \epsilon(u))\,\mathrm{d}u &= F_i^c \quad \text{(e)} \end{aligned}$$

for spline knot $j = 1, ..., n-1$ and contract $i = 1, ..., m$. The constraint in (a) ensures the adjustment function is continuous in the knots, while (b) and (c) imposes this restriction also for the first and second order differentials. The (d) constraint require the adjustment function to be horizontal at time T, and finally (e) also require the average value of the forward price function $f(u)$ over the delivery period for contract i to match the quoted closing price F_i^c. Here, we could take the interest rate effect from r into account and set the present value of the average of the forward price function equal to present value

of the forward contract. Instead of doing that, we will follow [Benth et al., 2008] and [Ollmar, 2003] and assume $r = 0$ such that the weight function in (3.4) is approximated with $w(u, \tau^s, \tau^e) \approx \frac{1}{\tau^e - \tau^s}$. Like [Benth et al., 2008] we will argue that both the prior and the smoothing will outweigh a marginal interest rate effect. In order to illustrate further, we look closer at the first spline where $u \in [\tau_0, \tau_1]$, and we can see that

$$
\begin{aligned}
\epsilon(u) &= a_1u^4 + b_1u^3 + c_1u^2 + d_1u + e_1 \\
\epsilon'(u) &= 4a_1u^3 + 3b_1u^2 + 2c_1u + d_1 \\
\epsilon''(u) &= 12a_1u^2 + 6b_1u + 2c_1 \\
\int_{\tau_0}^{\tau_1} [\epsilon''(u)]^2 \, du &= \int_{\tau_0}^{\tau_1} [(12a_1u^2 + 6b_1u + 2c_1)]^2 \, du \\
&= \frac{144}{5}a_1^2u^5 + 36a_1b_1u^4 + 16a_1c_1u^3 + 12b_1^2u^3 + 12b_1c_1u^2 + 4c_1^2u \Big|_{\tau_0}^{\tau_1} \\
&= [a_1 \ b_1 \ c_1 \ d_1 \ e_1] \begin{bmatrix} \frac{144}{5}\Delta_1^5 & 18\Delta_1^4 & 8\Delta_1^3 & 0 & 0 \\ 18\Delta_1^4 & 12\Delta_1^3 & 6\Delta_1^2 & 0 & 0 \\ 8\Delta_1^3 & 6\Delta_1^2 & 4\Delta_1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} a_1 \\ b_1 \\ c_1 \\ d_1 \\ e_1 \end{bmatrix} \\
&= \mathbf{x_1'h_1x_1}
\end{aligned}
$$

with $\Delta_1^l = \tau_1^l - \tau_0^l, \ l = 1, ..., 5$. We generalize for n sub periods with $\Delta_j^l = \tau_j^l - \tau_{j-1}^l$ and express the minimization problem as

$$\min_{\mathbf{x}} \quad \mathbf{x}^\intercal \mathbf{H} \mathbf{x},$$

where $\mathbf{x}$ is a $(5n \times 1)$ vector and

$$
\mathbf{H} = \begin{bmatrix} h_1 & & 0 \\ & \ddots & \\ 0 & & h_n \end{bmatrix}, \mathbf{h}_j = \begin{bmatrix} \frac{144}{5}\Delta_j^5 & 18\Delta_j^4 & 8\Delta_j^3 & 0 & 0 \\ 18\Delta_j^4 & 12\Delta_j^3 & 6\Delta_j^2 & 0 & 0 \\ 8\Delta_j^3 & 6\Delta_j^2 & 4\Delta_j & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}
$$

where the block diagonal matrix $\mathbf{H}$ has dimensions of $5n \times 5n$. All of the constraints are linear in x, and the optimisation problem may be expressed on the form $\mathbf{Ax} = \mathbf{b}$. The first three constraints apply to the knots, $j = 1, 2, ..., (n-1)$. If we focus on the first two knots $j = 1, 2$, we see that these constraints may be inserted as rows in the matrix $\mathbf{A}$:

$$\begin{bmatrix}
-u_1^4 & -u_1^3 & -u_1^2 & -u_1 & -1 & u_1^4 & u_1^3 & u_1^2 & u_1 & 1 & 0 & 0 & 0 & 0 & 0 & \dots \\
-4u_1^3 & -3u_1^2 & -2u_1 & -1 & 0 & 4u_1^3 & 3u_1^2 & 2u_1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & \dots \\
-12u_1^2 & -6u_1^2 & 2 & 0 & 0 & 12u_1^2 & 6u_1^2 & 2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \dots \\
0 & 0 & 0 & 0 & 0 & -u_2^4 & -u_2^3 & -u_2^2 & -u_2 & -1 & u_2^4 & u_2^3 & u_2^2 & u_2 & 1 & \dots \\
0 & 0 & 0 & 0 & 0 & -4u_2^3 & -3u_2^2 & -2u_2 & -1 & 0 & 4u_2^3 & 3u_2^2 & 2u_2 & 1 & 0 & \dots \\
0 & 0 & 0 & 0 & 0 & -12u_2^2 & -6u_2^2 & 2 & 0 & 0 & 12u_2^2 & 6u_2^2 & 2 & 0 & 0 & \dots \\
\vdots & & & & & & & & & & & & & & &
\end{bmatrix}$$

Constraints for the remaining knots, $j = 3, ..., (n-1)$ are developed in a similar fashion, producing a total of $r \times (n-1)$ rows to be inserted into $\mathbf{A}$. Combined with the end condition and the m price constraints, we get a total of $3n + m - 2$ constraints. These are expressed as $\mathbf{Ax} = \mathbf{b}$, where $\mathbf{A}$ is a $3n + m - 2 \times 5n$- dimensional matrix and $\mathbf{b}$ is a $3n + m - 2$-dimensional vector. The end condition is added as a single row in $\mathbf{A}$:

$$\begin{bmatrix} \dots & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 4u_n^3 & 3u_n^2 & 2u_n & 1 & 0 \end{bmatrix}$$

In order to include the final m rows with closing price constraints, we start out with

$$F_i^C = \frac{1}{\tau_i^e - \tau_i^s} \int_{\tau_i^s}^{\tau_i^e} \epsilon(u) + \Lambda(u)\ du$$

$$\int_{\tau_i^s}^{\tau_i^e} \epsilon(u)\ du = (\tau_i^e - \tau_i^s) F_i^C - \int_{\tau_i^s}^{\tau_i^e} \Lambda(u)\ du$$

to ensure the average value of the smoothed forward curve in the delivery period is equal to the closing price for contract $i = 1, ..., m$. The left side of the above equality is put in matrix $\mathbf{A}$, and the right side is put in the $\mathbf{b}$ vector after the $3n - 2$ zero's from the knot constraints. In the simplest case with no overlapping contracts, the delivery period for each contract i is covered by a single polynomial

$$\begin{aligned}
\int_{\tau_i^s}^{\tau_i^e} \epsilon(u)\ du &= \frac{1}{5} a_i u^5 + \frac{1}{4} b_i u^4 + \frac{1}{3} c_i u^3 + \frac{1}{2} d_i u^2 + e_i u \Big|_{\tau_i^s}^{\tau_i^e} \\
&= \frac{1}{5}[(\tau_i^e)^5 - (\tau_i^s)^5] a_i + \frac{1}{4}[(\tau_i^e)^4 - (\tau_i^s)^4] b_i + \frac{1}{3}[(\tau_i^e)^3 - (\tau_i^s)^3] c_i + \\
&\quad \frac{1}{2}[(\tau_i^e)^2 - (\tau_i^s)^2] d_i + [(\tau_i^e) - (\tau_i^s)] e_i \\
&= \rho_i^1 a_i + \rho_i^2 b_i + \rho_i^3 c_i + \rho_i^4 d_i + \rho_i^5 e_i
\end{aligned}$$

and the m rows to be added to matrix $\mathbf{A}$ are of the form

$$\begin{bmatrix} \rho_1^1 & \rho_1^2 & \rho_1^3 & \rho_1^4 & \rho_1^5 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \dots \\ 0 & 0 & 0 & 0 & 0 & \rho_2^1 & \rho_2^2 & \rho_2^3 & \rho_2^4 & \rho_2^5 & 0 & 0 & 0 & 0 & 0 & \dots \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \rho_3^1 & \rho_3^2 & \rho_3^3 & \rho_3^4 & \rho_3^5 & \dots \\ \vdots & & & & & & & & & & & & & & & \end{bmatrix}$$

However, if we have a case with overlapping delivery periods, the integral needs to be split in the number of sub periods necessary to cover the entire delivery period of the contract. We will use the same example as [Ollmar, 2003] and assume that the first contract has a delivery period equal to $[\tau_0, \tau_1]$ and write the constraint as

$$\begin{bmatrix} \rho_1^1 & \rho_1^2 & \rho_1^3 & \rho_1^4 & \rho_1^5 & \rho_2^1 & \rho_2^2 & \rho_2^3 & \rho_2^4 & \rho_2^5 & 0 & 0 & 0 & 0 & 0 & \dots \\ 0 & 0 & 0 & 0 & 0 & \rho_2^1 & \rho_2^2 & \rho_2^3 & \rho_2^4 & \rho_2^5 & 0 & 0 & 0 & 0 & 0 & \dots \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \rho_3^1 & \rho_3^2 & \rho_3^3 & \rho_3^4 & \rho_3^5 & \dots \\ \vdots & & & & & & & & & & & & & & & \end{bmatrix}$$

Now that we finally have $\mathbf{A}$ and $\mathbf{b}$, we can use the Lagrange multiplier method and reformulate as an unconstrained minimisation problem

$$\min_{\mathbf{x},\boldsymbol{\lambda}} L(x, \lambda) \;=\; \boldsymbol{x}^\intercal \boldsymbol{H} \boldsymbol{x} + \boldsymbol{\lambda}^\intercal (\boldsymbol{A}\boldsymbol{x} - \boldsymbol{B}) \tag{3.8}$$

If $[x, \lambda_0]$ is a solution, we require

$$\begin{aligned} \frac{\partial}{\partial x} L(x_0, \lambda_0) &= \mathbf{2H}x_0 + \mathbf{A}^\intercal \lambda_0 = 0 \\ \frac{\partial}{\partial \lambda} L(x_0, \lambda_0) &= \mathbf{A}x_0 - \mathbf{b} = 0 \end{aligned}$$

such that the solution can be found by solving the set of linear equations

$$\begin{bmatrix} \mathbf{2H} & \mathbf{A}^\intercal \\ \mathbf{A} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{x} \\ \boldsymbol{\lambda} \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \mathbf{B} \end{bmatrix} \tag{3.9}$$

where the left matrix has dimension of $(8n + m - 2) \times (8n + m - 2)$ and both vectors are of dimension $(8n + m - 2)$.

3.2.3 Forward curve calculation with the etrm package

The **etrm** package developed as a part of this PhD-project can be used to specify and solve the optimisation problem described above. The package is available on CRAN, and may be installed and loaded into the R environment by running the following commands:

```
install.packages("etrm")
library(etrm)
```

We demonstrate the **etrm** functionality with a practical example using market data from the power exchange Nasdaq OMX Commodities. The selected contracts are of type Nordic electricity (ENO), average rate (A) futures on base load (FUTBL) with monthly, quarterly and yearly delivery periods. Several other products are available, but at the trading date 2021-11-05, these were the instruments actively traded.

Include	Contract	Start	End	Closing
TRUE	ENOAFUTBLMNOV-21	2021-11-07	2021-11-30	62.50
TRUE	ENOAFUTBLMDEC-21	2021-12-01	2021-12-31	60.25
TRUE	ENOAFUTBLMJAN-22	2022-01-01	2022-01-31	62.50
TRUE	ENOAFUTBLMFEB-22	2022-02-01	2022-02-28	62.00
TRUE	ENOAFUTBLMMAR-22	2022-03-01	2022-03-31	44.00
TRUE	ENOAFUTBLMAPR-22	2022-04-01	2022-04-30	38.40
TRUE	ENOFUTBLQ1-22	2022-01-01	2022-03-31	56.00
TRUE	ENOFUTBLQ2-22	2022-04-01	2022-06-30	31.00
TRUE	ENOFUTBLQ3-22	2022-07-01	2022-09-30	20.90
TRUE	ENOFUTBLQ4-22	2022-10-01	2022-12-31	31.15
TRUE	ENOFUTBLQ1-23	2023-01-01	2023-03-30	41.30
TRUE	ENOFUTBLQ4-23	2023-10-01	2023-12-31	31.40
FALSE	ENOFUTBLYR-22	2022-01-01	2022-12-31	34.50
TRUE	ENOFUTBLYR-23	2023-01-01	2023-12-31	29.35
TRUE	ENOFUTBLYR-24	2024-01-01	2024-12-31	29.50
TRUE	ENOFUTBLYR-25	2025-01-01	2025-12-31	29.83

Table 3.2: Selected base load futures, Nasdaq OMX Commodities 2021-11-05

There will typically be multiple instruments available covering the same calendar period, for example January, February and March versus the Q1 contract. By selecting the ones of shortest duration, we ensure that information of high granularity readily available in market prices is utilized in the forward curve calculation. This is handled via the boolean `include` argument to the `msfc()` constructor function, see Listing 3.1 for a code example. Once provided the remaining data presented in Table 3.2, the constructor returns an instance of the S4 class `MSFC`. Three generic methods are implemented for this type of object: `plot()`, `summary()` and `show()`.

The spot wholesale electricity price in the Nordics was unusually high during the fall of 2021, partly due to low renewable energy supply, and high fuel and carbon prices. This can be seen in the short contracts in the upper part of Table 3.2. Further ahead in time, prices are lower. These contracts cover longer time spans (years). At this horizon, the seasonality typically present in such markets is not reflected in market prices. To compensate, the user may provide a prior function to be included in the calculation. In the code example in Listing 3.1 we have used a simple trigonometric prior, where weekend prices have been adjusted down by a fixed factor. The `msfc()` function is used to create two instances of the `MSFC` class, and the `plot()` method create the charts in Figure 3.3.

```
library(etrm)
library(gridExtra)
load(file = "eno.Rda")
load(file = "prior.Rda")

# forward curve with prior
fwd.fut.wpri <- msfc(tdate = as.Date("2021-11-05"),
                     include = eno$Include,
                     contract = eno$Contract,
                     sdate = eno$Start,
                     edate = eno$End,
                     f = eno$Closing,
                     prior = prior$mod)

# forward curve without prior
fwd.fut.npri <- msfc(tdate = as.Date("2021-11-05"),
                     include = eno$Include,
                     contract = eno$Contract,
                     sdate = eno$Start,
                     edate = eno$End,
                     f = eno$Closing,
                     prior = 0)

grid.arrange(plot(fwd.fut.wpri, ylab = "EUR/MWh", legend = ""),
             plot(fwd.fut.npri, ylab = "EUR/MWh", legend = ""))
```

Listing 3.1: Example of the `msfc()` constructor and the `plot()` method for class `MSFC`

As the forward curve provide a detailed description the forward market and it's term structure, it is a flexible tool for supporting investment- and trading decisions. Since the standardised contracts are average price contracts for delivery at a fixed rate over some time interval, they are not suited for finding the value of a profiled volume. Quite often, the volume consumed by both businesses and households will be subject to a similar type of seasonality as the energy price, creating a positive correlation between energy

demand and price. This makes the forward curve useful when pricing non-standard contracts and fixed price agreements tailor made for a specific consumption profile. The curve has many other practical applications, for example within risk management and production planning.

From the closing price constraint in the minimisation problem defined in previous chapter, it is also clear that the standard contracts can be priced with this tool by calculating the average value of the forward curve over the delivery period in question. The `summary()` method demonstrated in Listing 3.2 provide the computed prices along with the quotes used as input for the calculation. The method also deliver a description of the spline used to build the curve and a sample of the initial prior function values.

```
> summary(fwd.fut.wpri)
$Description
[1] "MSFC of length 1518 built with 25 polynomials at trade date
    2021-11-05"

$PriorFunc
[1] 54.43464 54.50272 54.56650 54.62597 49.21299 49.25870

$BenchSheet
   Include          Contract       From         To Price  Comp
1     TRUE ENOAFUTBLMNOV-21 2021-11-07 2021-11-30 62.50 62.50
2     TRUE ENOAFUTBLMDEC-21 2021-12-01 2021-12-31 60.25 60.25
3     TRUE ENOAFUTBLMJAN-22 2022-01-01 2022-01-31 62.50 62.50
4     TRUE ENOAFUTBLMFEB-22 2022-02-01 2022-02-28 62.00 62.00
5     TRUE ENOAFUTBLMMAR-22 2022-03-01 2022-03-31 44.00 44.00
6     TRUE ENOAFUTBLMAPR-22 2022-04-01 2022-04-30 38.40 38.40
8     TRUE    ENOFUTBLQ2-22 2022-04-01 2022-06-30 31.00 31.00
9     TRUE    ENOFUTBLQ3-22 2022-07-01 2022-09-30 20.90 20.90
10    TRUE    ENOFUTBLQ4-22 2022-10-01 2022-12-31 31.15 31.15
11    TRUE    ENOFUTBLQ1-23 2023-01-01 2023-03-30 41.30 41.30
12    TRUE    ENOFUTBLQ4-23 2023-10-01 2023-12-31 31.40 31.40
14    TRUE    ENOFUTBLYR-23 2023-01-01 2023-12-31 29.35 29.35
15    TRUE    ENOFUTBLYR-24 2024-01-01 2024-12-31 29.50 29.50
16    TRUE    ENOFUTBLYR-25 2025-01-01 2025-12-31 29.83 29.83
```

Listing 3.2: Example of the `summary()` method for class `MSFC`

The `MSFC` class contain further details regarding calculation results and methods used:

```
> slotNames(fwd.fut.wpri)
[1] "Name"        "TradeDate"   "BenchSheet"
[4] "Polynomials" "PriorFunc"   "Results"
[7] "SplineCoef"  "KnotPoints"  "CalcDat"
```

The `show()` method can be used to access the data frame in the `Results` slot, where the daily values for the forward curve and the selected contracts are stored. The method is not demonstrated here due to space constraints. The vector with the daily forward curve values can also be accessed directly, in our example with:

```
fwd.fut.wpri@Results$MSFC
```

For further information regarding the maximum smoothness forward curve implementation in **etrm**, the reader is referred to **Paper B**.

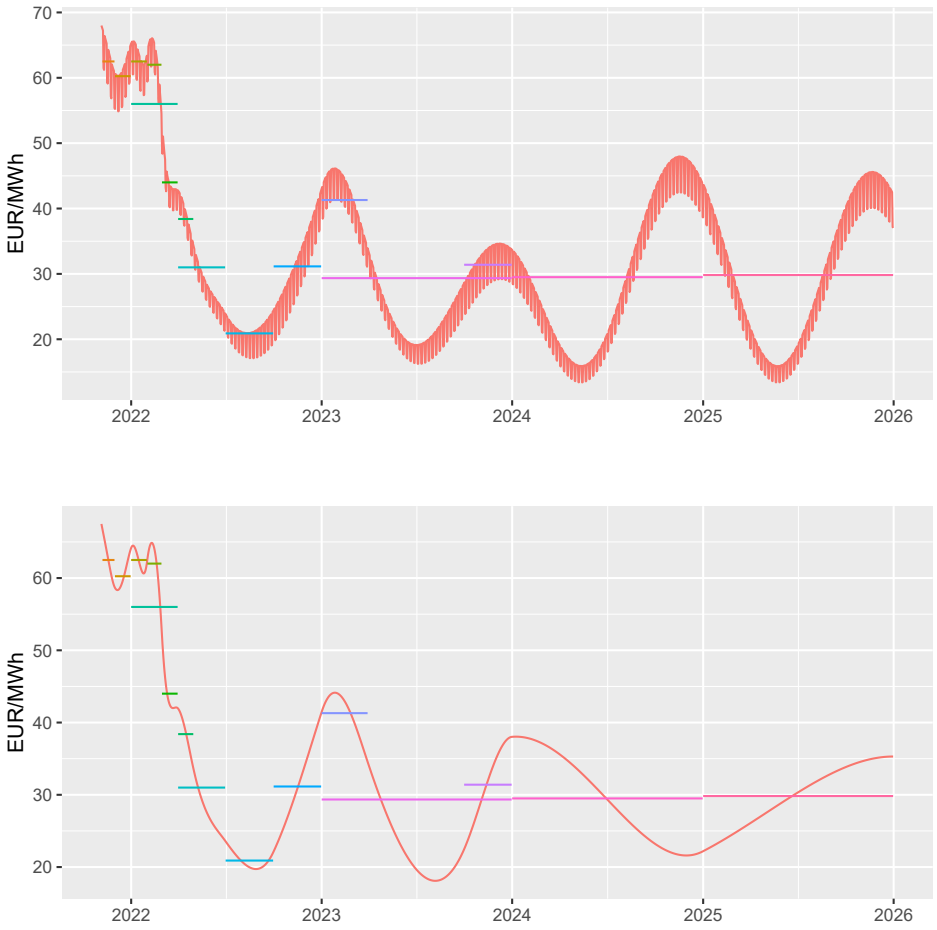

Figure 3.3: MSFC with (top) and without (bottom) prior

3.3 Energy price risk management

Participants in the energy markets are exposed to a number of risks. Some of these have to do with individual characteristics of the participant (volume and profile risk), while others are external (liquidity-, currency-, counterparty- and price risk). The main focus in this thesis is on the price risk of the energy commodity. In this chapter we will start out with a brief treatment of the concept of price risk and hedging as a risk mitigation technique, before we introduce some important terms and definitions such as the *energy portfolio*, the *portfolio price*, the *portfolio management authorisation* and finally the *hedging strategy*. Thereafter, we continue with a presentation of five portfolio insurance strategies that may be used for price risk management: Constant Proportion Portfolio Insurance, Dynamic Proportion Portfolio Insurance, Option Based Portfolio Insurance, Step Hedge Portfolio Insurance and Stop Loss Portfolio Insurance. We conclude with a short demonstration of how these are implemented in the **etrm** package.

3.3.1 The energy portfolio

A fundamental energy market participant (as opposed to a financial investor) has a physical volume to be consumed, produced or delivered over a time interval in the future. There is uncertainty both with respect to how much energy that is needed at each point in time, as well as the total size of the volume over that interval. Typically, a prognosis will be made based on fundamental variables relevant for the company in question. Some common explanatory variables are planned production or activity level, holidays and calendar effects, and weather. The future commodity spot price to be paid for this volume is also unknown, and represent a significant financial risk, as illustrated in previous sections.

Market players can mitigate some of the price risk by entering into forward market contracts, locking in the price of a future delivery today. When we evaluate the values of the future energy volumes, we can think of them as assets (or liabilities) exposed to the same price risk as the traded contracts. The price risk of the energy volumes can be mitigated by taking offsetting positions in the forward contracts. At this point we will introduce the term energy portfolio.

Definition 1 (Energy Portfolio) *Consider a physical energy market participant with a volume prognosis for N future delivery periods with start and end dates (τ_s^i, τ_e^i), $i = 1, 2, ..., N$. The volume prognosis and corresponding risk offsetting positions in the forward market is the energy portfolio of the market participant.*

The delivery periods most commonly used for reporting in the industry are years and seasons. In the remainder of this thesis we will be using the former.

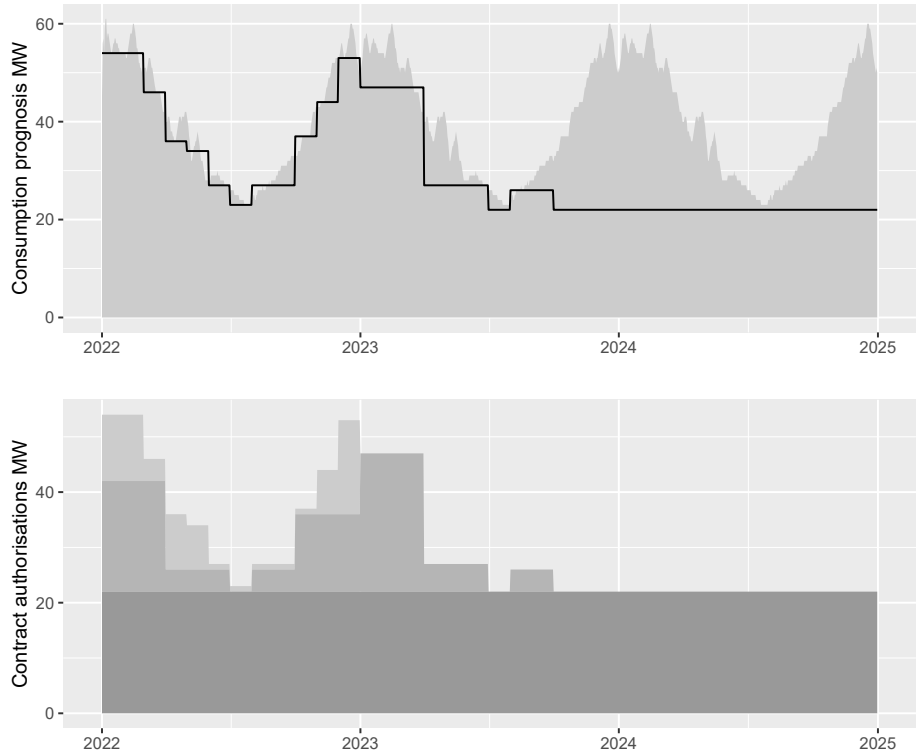

Figure 3.4: Example consumption prognosis with maximum hedge rate (top) and portfolio management authorisations for year, quarter and month contracts (bottom)

In order to sharpen our focus on the commodity price exclusively, we will not go further into other risk factors such as counter-party, market liquidity, currency, and volume and profile risk. The reader is referred to [Eydeland and Wolyniec, 2002] for further details regarding these topics. We also assume that the underlying settlement reference for the derivatives contracts match the commodity price for the volumes, disregarding potential basis risk. The unit price of the portfolio, the portfolio price, is a key metric in commodity price risk management.

Definition 2 (Portfolio Price) *The unit price of energy at time t for a specific delivery period (τ_s^i, τ_e^i), $i = 1, 2, ..., N$ is termed the portfolio price, p_t. It is calculated as the market value of the volume prognosis and the forward market positions, per unit of energy for the period in question.*

In this set up, a portfolio manager is allowed to trade in the forward market to reduce price exposure of the portfolio. This is where the portfolio management authorisation comes into play.

Definition 3 (Portfolio Management Authorisation) *Consider a physical energy market participant with an energy portfolio. The portfolio management authorisation for the energy portfolio is a set of restrictions on trading activity:*

- *Maximum trading horizon (delivery periods available for trading)*
- *Maximum and minimum hedge rate per period, typically in* $h_t \in (0, 1)$
- *Target portfolio price,* p_t^**, per period (cap or floor)*

In practise, one could have more complex authorisations with time varying requirements for the hedge rates, restrictions on re-selling transacted volumes, start and end dates for trading per contract type, and more. It is important that the elements included in the authorisation do not come in conflict with each other, and that the mandate is possible to implement in the current market environment. The authorisation also has to match the portfolio owner's needs, experience and overall risk preferences.

An example portfolio management authorisation for the years 2022-2024 is illustrated in Figure 3.4. For the year 2022, the portfolio manager is allowed to trade year, quarter and month contracts. Further ahead in time, there are no monthly contracts available for trading, or liquidity might be poor. Focus is kept on the yearly and quarterly instruments. As time progresses and contracts of shorted length are made available for the last two years, the mandate may be updated.

Going forward, we will focus on delivery periods and sub periods with constant base load volumes. As illustrated in Figure 3.4, these can be used to approximate the volume prognosis. When operating with base load periods, the price risk can be offset by taking positions in standardised contracts in the forward market. In the remainder of the chapter, we will be focusing exclusively on constant rate delivery and hedging with yearly contracts. The portfolio price p_t can then be calculated with

$$p_t = \big[f_0 h_0 q + \sum_{i=1}^{t} f_i(h_i - h_{i-1})q + f_t(1 - h_t)q\big]/q \tag{3.10}$$

$$= f_0 h_0 + \sum_{i=1}^{t} f_i(h_i - h_{i-1}) + f_t(1 - h_t) \tag{3.11}$$

where q is the required base load volume, $h_t \in (0, 1)$ the hedge rate and f_t the futures price at time t. Prior to trading $p_t = f_t$. The first two terms in (3.11) contain the value of the transacted volumes, while the last term is the current market value of the unhedged part of the portfolio.

In order to manage the risk we need to enter into positions in the forward market, and those positions need to be within the limits set by the portfolio management authorisation. The overall objective is to control risk, and the key metric is the portfolio price. Each day of the hedging period there is a wide array of possible hedge rates for each of the contracts in the energy portfolio. The importance of how h_t is determined is illustrated in Figure 3.5.

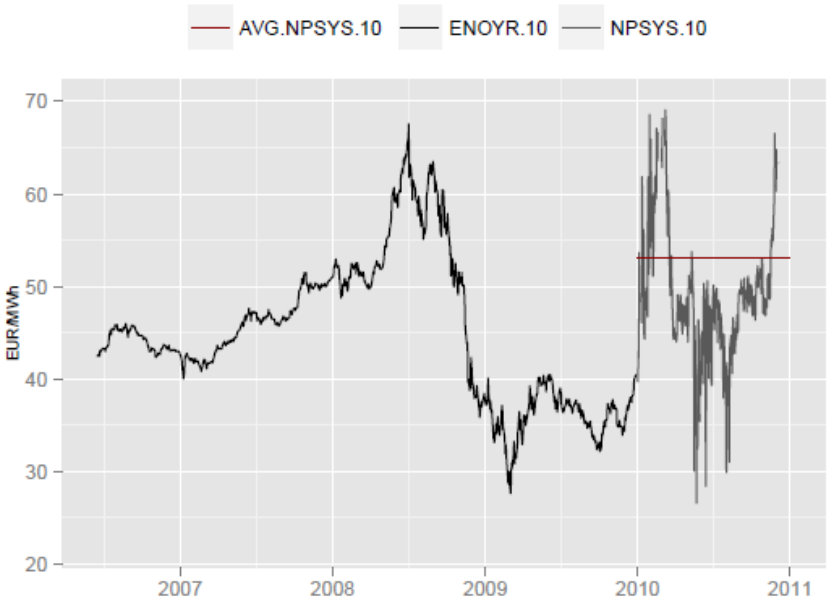

Figure 3.5: Nordic forward and spot price for 2010

The price of the Nordic electricity forward ENOYR-10 varied significantly before it reached expiry in late 2009. After climbing above 65 EUR/MWh during the first half of 2008, the following financial crisis led to a sharp drop below 30 EUR/MWh. The spot price realised during the delivery period was also very volatile. A consumer locking in the price prior to the market fall without exit possibilities could end up with a lost opportunity after the drop and considerable higher cost than the average spot price delivered during 2010. These observations underline the importance of having a good method for determining the hedge rates of the energy portfolio.

3.3.2 Hedging strategies

The strategies to be presented here all aim to provide *portfolio insurance*, e.g. they seek to protect the portfolio from adverse market price movements, in particular breaching a pre-defined *target price* level for the portfolio. We start out with a quite parsimonious definition that captures some of the key characteristics of a hedging strategy.

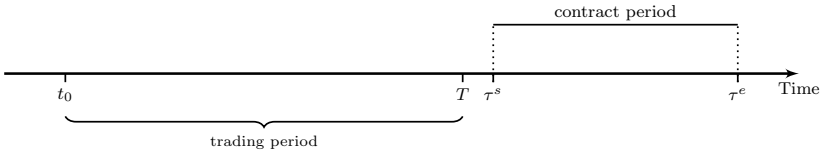

Figure 3.6: Trading and delivery period for an energy forward contract

Definition 4 (Hedging Strategy) *Consider a base load delivery period available for trading in the portfolio management authorisation. A hedging strategy for the forward contract covering the period is a function*

$$F(f_t, p_t, \theta_t) \rightarrow h_t \tag{3.12}$$

where t are the days in the trading period $t = 0, 1, ..., T$, f_t the forward price, p_t the portfolio price, θ_t a parameter vector and the hedge rate is $h_t \in (0, 1)$.

The strategies all evaluate current market and portfolio price and return a desired hedge dependent on model parameters. We will elaborate more on the model specific parameters when we describe the strategies below. From this point we will proceed with a description of energy price risk management, as seen from the perspective of the consumer. This discipline of managing expected future energy costs is relevant for large end users as well as suppliers procuring energy in the wholesale market for further distribution to end users. A similar picture, seen from the viewpoint of the producer side can easily be achieved by reversing the logic. We proceed with a short description of the strategies implemented in **etrm**. For further information regarding the technical implementation, the reader is referred to **Paper B**.

Introduced by [Perold, 1986] and [Black and Jones, 1987], the *Constant Proportion Portfolio Insurance* (CPPI) strategy was initially used as a means for placing a floor on the value of investment portfolios in fixed-income instruments and stocks. The manager of an investment portfolio would adjust the proportion of bonds and stocks and thereby create a leveraged exposure to the risky assets, that would depend on how far away the

portfolio value departed from the floor. In our context, the CPPI strategy is set up in the following way: For a specific future delivery period, a target price p^* equal to the highest acceptable portfolio price is determined. The target price must be set higher than the market price of f_0 upon start of the trading period. The difference c_t between the target price and the portfolio price is termed the cushion. The key idea of CPPI is that the proportion of the portfolio exposed to the market should be calculated as a constant multiple m of the cushion. The multiple is given by $m = \mu^{-1}$, where the risk factor $\mu > 0$ is the maximum daily price change to be handled by the asset allocation scheme. The hedge rate is given by

$$h_t = 1 - c_t \times m \tag{3.13}$$

$$= 1 - (p^* - p_{t-1}) \times \mu^{-1} \tag{3.14}$$

Note that the portfolio management authorisation for the hedging scheme require that $h_t \in (0, 1)$. In order to prevent the consumer to overhedge or going short in the futures positions, we require $h_t = 0$ if $c_t > \mu$ and $h_t = 1$ if $c_t < 0$. When starting the CPPI hedging, the owner of the energy portfolio will set a target price and a risk factor based on risk tolerance. The target price will typically be set above the market in such a way that the cushion is smaller than the risk factor. In this scenario an increase in the price f_t will lead to a smaller cushion due to the effect of the unhedged volume on the portfolio price. This would again dictate an increase in the hedge rate to keep the target price protected. Conversely, when the price f_t falls, the cushion increases, triggering a reduction in the hedge to allow more volume to be exposed to the falling market. In case of a sudden and significant price spike above the target price, the CPPI strategy can run into difficulties with liquidity and gap risk, see for example [Jessen, 2014]. The significance of the gap risk will be determined by the choice of risk factor, and the portfolio owner needs to make a well informed decision prior to the start of the hedging. The strategy can also lead to the portfolio price being locked in at or above the target level, and not following the market back down in the case of a price fall.

Dynamic Proportion Portfolio Insurance (DPPI) resembles the CPPI strategy, but as the name suggests, the proportion exposed to the market is no longer a constant multiple of the cushion. The risk factor is allowed to vary with time, see for example Lee et al. [2008] and [Chen et al., 2008] for two implementations focusing on equity portfolios. By varying the multiple, the DPPI strategy aim to better capture changing market conditions, such as the increase in volatility observed in Figure 3.5. The time dependent μ_t might be determined by using Value-at-risk, Expected Shortfall or some other risk measure, see [McNeil et al., 2015] for a thorough treatment of this topic. Here, we will implement a model also allowing resetting of the target price p_t^*, in order to push the protected price

cap further down, if market conditions give us this opportunity. The hedge rate is now calculated as

$$h_t = 1 - c_t \times m_t \tag{3.15}$$

$$= 1 - (p_t^* - p_{t-1}) \times \mu_t^{-1} \tag{3.16}$$

Where $p_t^* = Min(\lambda p_{t-1}, p_{t-1}^*)$ and $\lambda = \frac{p_0^*}{p_0}$. The target for the portfolio price is evaluated at each point in time t and reset according to the decision rule above, when possible. This may allow the DPPI strategy to lock in a lower cap in a falling market scenario. The DPPI is still vulnerable to gap risk though, as well as the lock-in effect if the target price level is reached, as described above.

The *Option Based Portfolio Insurance* (OBPI) aims to protect the energy portfolio from price risk by replicating a European call option on the forward contract. We do not procure the option, but rather synthesize the contingent claim by trading in the underlying asset. This *delta hedging* scheme relies on the option pricing model used. Since the underlying asset is a forward contract, we utilize the Black-76 formula introduced by [Black, 1976]. The price at time t of the European call option with exercise date T and strike price K, on a futures contract with delivery start $\tau^s \geq T$ is given by

$$C(t, f_t, K, \sigma, r) = e^{-r(T-t)}[f_t N(d_1) - KN(d_2)] \tag{3.17}$$

where f_t is the futures price and N is the cumulative distribution function of the standard normal $N(0, 1)$, where

$$d_1 = \frac{\ln(f_t/K) + (\sigma^2/2)(T-t)}{\sigma\sqrt{(T-t)}} \tag{3.18}$$

$$d_2 = d_1 - \sigma\sqrt{(T-t)} \tag{3.19}$$

and r is the risk free rate of interest, σ the volatility of the underlying futures price and $(T-t)$ is the time to exercise. The Black-76 model assumes a frictionless market with no transaction costs or tax implications, and a constant risk free interest rate r both for borrowing and lending money. Furthermore, the futures price is presumed to follow a log-normal process with constant volatility σ. Some of these assumptions are clearly not fulfilled, and there are other models better able to capture empirical observations such as time-varying volatility and breach of the log-normal distribution assumption. Still, we will rely on the Black-76 as this is widely used in the industry and considered sufficient for our purpose.

From the equations above, we see that the price of a call option is sensitive to 1) price changes in the underlying asset and 2) changes in the parameters used in the model.

The sensitivity measures of the option premium with respect to these changes are often called the greeks. Some of the most frequently used greeks are the delta (change in price of underlying asset), vega (change in volatility), theta (change in time to expiration) and rho (change in interest rate). The first sensitivity, the delta, is a measure of risk exposure i.e. how the value of a derivative asset will change given a certain change in the underlying asset price. The other greeks, however, are sensitivities regarding model parameters. It might seem contradictory to evaluate changes in these, as they are assumed constant in the Black-76 model. This issue is also commented in [Bjork, 2009], where it is underlined that these greeks formally should be considered as sensitivities with respect to misspecification of the model, rather than measures of risk exposure.

We shall focus on the delta, denoted Δ^C. For the call option, this sensitivity measure is calculated with

$$\Delta^C = \frac{\partial C(t, f_t, K, \sigma, r)}{\partial f_t} = e^{-r(T-t)} N(d_1) \tag{3.20}$$

If we evaluate Δ^C over over the range of possible values for f_t, we quickly see from the expression for d_1 and the properties of $N(0,1)$ that $\Delta^C \to 0$ as $f_t \to 0$ (option is out-of-the-money). Further, we see that $\Delta^C \to e^{-r(T-t)}$ as $f_t \to +\infty$ (option is in-the-money). When $f_t = K$ (option is at-the-money), $d_1 = \frac{(\sigma^2/2)(T-t)}{\sigma\sqrt{(T-t)}}$, and Δ^C will be closer to $\frac{1}{2}$, depending on the assumed risk free rate of interest, volatility and time to exercise date.

The Δ^C plays an decisive role in portfolio hedging, and it is used in the derivation of the Black-Scholes equation. Using arbitage arguments, the writer of a call can sell one unit of the option and hedge the exposure by taking a position of Δ^C units of the underlying asset. Such a portfolio that is insensitive to small changes in the price of the underlying is said to be delta neutral, see [Bjork, 2009]. The corresponding position in the underlying asset is the delta hedge.

Consider a simple portfolio, consisting of a short position in the European call, and a long position in the underlying futures contract. Formally, at time t we have

$$\begin{aligned} -1 &= \textit{number of units of the call} \\ h_t &= \textit{number of units of the futures contract at time } t \\ f_t &= \textit{futures price at time } t \end{aligned}$$

The portfolio value is $V_t = -1 \times C(t, f_t, K, \sigma, r) + h_t \times f_t$, and delta neutrality require

$$\frac{\partial}{\partial f_t}\big[- C(t, f_t, K, \sigma, r) + h_t \times f_t \big] = 0$$

which is solved by

$$h_t = \frac{\partial C(t, f_t, K, \sigma, r)}{\partial f_t} = e^{-r(T-t)} N(d_1) \tag{3.21}$$

Thus, if we were to issue a call option, the option delta give us the number of units of the underlying futures contract needed to hedge the short position in the option. This delta hedge will only work for small changes in f_t, and the position would have to be evaluated and updated continuously if the portfolio should be perfectly delta neutral. In practice, one would have to come up with some rules for when the futures positions should be updated, and by how much at the time. This is probably what our counterparty would do if we were to implement the OBPI strategy by purchasing a call option. The writer of the call would calculate a daily hedge ratio based on the delta neutral position for the underlying contract. The counterparty will then adjust the hedge ratio as the market moves, aiming for a cost of rebalancing the hedge portfolio which is lower than what is received in option premium. When we implement the Black-76 delta hedge strategy we seek to replicate the characteristics of the option by using exactly the same hedging regime as the seller of the option. For a specific future delivery period we will select the appropriate futures contract and set strike price upon start of hedging. The daily hedge ratio will give us the risk neutral position of the portfolio moving forward. The hedge level will be increased when the market price f_t go up in order to protect the portfolio. A decreasing hedge rate when market price f_t is falling makes it possible for the portfolio price to fall together with the market.

Note that also for this strategy, we are facing some risks associated with the portfolio rebalancing. In a situation with poor liquidity, we might struggle to transact the volumes wee need at reasonable prices. If the price of f_t spikes and we do not succeed to rebalance as planned, the hedging strategy can lose some of its protection. As we may face these issues of liquidity and gap risk, we should expect that the portfolio price most likely will be capped at initial market price plus the premium of the European call option. In practise, one could take these problems into account by adding a buffer on top of the strike + option premium when setting the cap for the portfolio price. This would be the same as making a conservative estimate of the replication costs. Our goal, however, is to achieve the same protection as the option holding OBPI strategy, at a lower cost than the option premium. Note that here, the expected target price to be protected for the portfolio is given implicitly via the selected strike price K and the option replication costs (premium).

$$p^* \approx K + \hat{C}(t, f_t, K, \sigma, r)$$

In order to protect this price level, the level of trading activity and related transaction costs cannot be too high before they make the rebalancing too costly.

The *Step Hedge Proportion Portfolio Insurance* (SHPI) strategy is also included as a simple benchmark strategy. The hedge rate for a buyer at time t is given by

$$h_t = \begin{cases} \frac{t}{T-t_0+1} & \text{if} \quad p_t < p^* \\ 1 & \text{if} \quad p_t \geq p^* \end{cases} \tag{3.22}$$

Consider a situation where the portfolio manager has a mandate allowing trading to start 500 days prior to the contract expiry. I this case, the hedge rate would be gradually increased during the period $h_t = \frac{1}{500}, \frac{2}{500}, ..., 1$. This approach aim to smoothen out the hedging price by mechanically spreading transactions over the entire trading period. If the target level is reached before time T, the portfolio is fully hedged. The strategy is vulnerable to the same gap risk and lock-in effects mentioned above, and it may require large volumes to be locked in rather quickly if the target level is reached early in the trading period.

Finally, the *Stop Loss Portfolio Insurance* (SLPI) is a rudimentary approach to hedging, where the future volume is left unprotected unless the target level p^* is reached. The hedge rate is determined according to the rule

$$h_t = \begin{cases} 0 & \text{if} \quad p_t < p^* \\ 1 & \text{if} \quad p_t \geq p^* \end{cases} \tag{3.23}$$

In the event that the target price level is not reached before contract expiry, the volume is left open. Also this approach is vulnerable to gap risk and target price lock-in in a potentially falling market. It will also be the strategy most exposed to poor market liquidity, as the entire volume needs to be procured in the event that target price level is reached. This may also happen with the other strategies though, if the market price f_t first drop from the initial f_0 level, leading to a low hedge rate, followed by a sharp increase in price.

The strategies described above all implement a forward trading scheme procuring energy prior to the actual consumption to protect against spot prices realised through the delivery period. However, the only approach that will automatically end up at a full hedge rate $h_T = 1$ at the time of contract expiry is the SHPI. In a scenario where the market price evolves in a favourable way, the remaining strategies will lower hedge rates in an effort to achieve lower prices for the portfolio. When approaching expiry, the portfo-

lio manager has the option either to 1) lock in a desired hedge level for the contract, or 2) decide to leave the position as-is and "roll-forward" into shorter duration contracts, see the example mandate in Figure 3.4. It is important to underline that in practical applications, the models presented here are contract-specific strategies to be included into a complex multi-contract strategy managing positions matching the consumption prognosis more closely, all within the limits of the management authorisation given by the portfolio owner. This is a challenging task. The strategies and corresponding **etrm** package portfolio insurance functions should be considered tools that may be used when solving such a problem.

From Definition 4 and the presentation of the individual models above, we also see that there has not been made any attempt to include fundamental factors explaining the changes in f_t when making decisions regarding hedging. Most energy market participants will have a market view based on factors affecting the price, such as temperature, fuel and carbon prices, wind and hydrology. One way to include this could be to treat the portfolio insurance strategies as benchmarks, and allow deviations from the benchmark within pre-defined limits. In a setting where OBPI is used, we could also mimic an option trading scheme by resetting the strike K in the delta hedging based on market price expectations. When expectations are included in the decision making, one can dispute whether the activity performed by the portfolio manager still can be defined as hedging, though. Further exploration of this approach to energy procurement is left for future studies.

Before we move on with some strategy evaluation criteria we should underline one important point: If you enter into positions in the forward market in order to hedge against future price risk, you can think of that activity as procuring insurance, and insurance come at a cost. Even though the expected value of insurance agreements are negative, people still buy them - due to risk aversion. Therefore, it would be reasonable to expect to pay a (preferably low) premium relative to the market for the price risk management, which will protect you in times of turmoil. When comparing the strategies in the empirical analysis below, we will using the following metrics:

1. Difference between portfolio price and market price at contract expiry
2. Maximum difference between portfolio price and market price
3. Target price protection
4. Rebalancing and required trading activity

For the cases where the portfolio price does not reach the target price level, the first criteria identify the added costs associated with hedging activity by comparing the achieved

portfolio price at the contract expiry with the market price f_T. When the market passes the target level, the criteria measures savings achieved with hedging. The second metric evaluate the maximum difference between the price of the portfolio and the market over the entire trading period. A participant in the forward market will have to meet margin requirements and manage cash flows for settlement, and this metric provide insight to be considered in this regard. As mentioned above, the strategies are vulnerable to gap risk. The third metric evaluate how effective the target price level protection has been. As the DPPI might adjust the target during the trading period, the metric is calculated by comparing the initial target price with final portfolio price. We also evaluate the maximum difference between the portfolio price and the target over the entire trading period. Finally, rebalancing and trading activity is added to the evaluation criteria. By calculating the proportion of transacted volumes over the volume q to be delivered, we get the portfolio churn rate. This is an important input when assessing costs of trade.

An example of how the strategies can be implemented using the **etrm** package is presented below.

3.3.3 Price risk management with the etrm package

In this chapter we will explore how the hedging strategies perform when applied to actual forward market data. The data set provided by Montel consists of daily closing prices for 16 yearly contracts for Nordic electricity (ENOFUTBLYR), for delivery in 2006-2021. The first 12 contracts are of type *Delayed Settlement Future* where the mark-to-market settlement is delayed to the delivery period. These contracts resemble traditional forward contracts, but are typically cleared at a clearinghouse to manage counterparty risk. The remaining part of the data set consists of the new products introduced at Nasdaq OMX Commodities, e.g. traditional futures contracts with daily mark-to-market settlement. Both have the Nordic system price as settlement reference.

The **etrm** package has tools for implementing trading strategies for price risk management, for commercial hedgers with both long or short exposure. We will focus on a large consumer in the following example. The strategy functions `cppi()`, `dppi()`, `obpi()`, `shpi()` and `slpi()` all aim to achieve a favorable unit price for the energy portfolio, while preventing it from breaching a predefined target price level, as described in previous chapter. They act as constructor functions for their corresponding S4 classes `CPPI`, `DPPI`, `OBPI`, `CPPI`, `SHPI` and `SLPI`. Each of these classes have implemented the methods `plot()`, `summary()` and `show()`. This design makes it convenient to test alternative assumptions regarding the model parameters and trading costs without too much effort.

```
library(etrm)
load(file = "enocal.Rda")

# back test settings
q <- 30                # 30 mw base load delivery
rper <- 0.2            # constant risk factor 0.2 * f_0 for cppi
vol <- 0.14            # simple fixed annualized volatility
daysleft <- 750        # approximately three years horizon
r <- 0                 # zero risk free interest rate
tcost <- 0.1           # 0.1 EUR/MWh fixed transaction cost
tdate <- enocal$Day    # day number 1,..,750 for time vector

# lists for storing strategy objects
cppi_list <- list()
dppi_list <- list()
obpi_list <- list()
shpi_list <- list()
slpi_list <- list()

for (i in 1:(dim(enocal)[2]-1)){

  # price, contract name and rf_t for contract i
  f <- enocal[, (i + 1)]
  cn <-names(enocal)[i + 1]
  rper_t <- rep(rper, length(f))

  # run strategies on contract i
  obpi_list[[cn]] <- obpi(q = q, tdate = tdate, f = f, vol=vol,
                          r = r, daysleft=daysleft, tcost = tcost)

  # implicitly defined tper from option strategy
  tper <- obpi_list[[cn]]@Results$Target[1]/f[1] - 1

  cppi_list[[cn]] <- cppi(q = q, tdate = tdate, f = f, tper = tper,
                          rper = rper, tcost = tcost)

  dppi_list[[cn]] <- dppi(q = q, tdate = tdate, f = f, tper = tper,
                          rper_t = rper, tcost = tcost)

  shpi_list[[cn]] <- shpi(q = q, tdate = tdate, f = f, daysleft =
                         daysleft, tper = tper, tcost = tcost)

  slpi_list[[cn]] <- slpi(q = q, tdate = tdate, f = f, tper = tper,
                          tcost = tcost)
}
```

Listing 3.3: **etrm** strategies example

In the R code in Listing 3.3, we run all strategies for the 16 contracts. The results will naturally depend on the assumptions used in the test. Here we have assumed an annual base load consumption equivalent to 30 MW, and a constant risk factor μ for the CPPI calculated with $\mu = f_0 \times rper$ for each of the contracts. For simplicity and illustrative purposes, we use a constant annualized volatility of 14 percent, and trading starts approximately three years (750 days) prior to contract expiry. We set the option strike price K at-the-money first day of trading, and use the implicit target price from OBPI to calculate target prices for the remaining strategies. We assume a risk free rate of interest equal to zero, and constant trading costs equal to 0.1 EUR/MWh. The strategy functions handle both dates and numeric time vectors, an integer vector is used here. We also would like to underline that the inputs used here are not to be considered a best attempt to achieve a specific goal with respect to price risk management, but merely for illustrative purposes. The 80 strategy objects that have been generated are stored in lists for analysis and evaluation. The metrics needed for the strategy evaluation can be collected from the list elements. Some are also calculated by the `summary()` method:

```
> summary(dppi_list$`Cal-07`)
$Description
[1] "Hedging strategy of type DPPI and length 750"

$Volume
[1] 30

$Target
[1] 31.32706 31.32413 31.19073 31.07011 30.92757 30.41440 30.34167
    30.21557

$ChurnRate
[1] 2.6

$Stats
       Market Trade Exposed Position      Hedge    Target Portfolio
First   28.57    16      14       16 0.5333333 31.32706  28.62333
Max     59.00    16      15       30 1.0000000 31.32706  30.10367
Min     26.27    -2       0       15 0.5000000 30.21557  27.55633
Last    37.04     0       0       30 1.0000000 30.21557  30.10367
}
```

Listing 3.4: DPPI strategy `summary()` method example

As we can see in Listing 3.4, the DPPI strategy managed to lower the target price level a number of times when applied to the Cal-07 contract. The reported churn rate tell us the underlying volume has been traded 2.6 times during the three years. The largest

trade was made at the very beginning, where the DPPI needed to transact 16 MW in order to reach initial hedge rate of 53.3 percent. This is a consequence of the selected target price and risk factor. The initial portfolio price deviate from market price due to the transaction costs. Final portfolio price achieved is approximately 30.10 EUR/MWh, almost 7 EUR/MWh lower than the market, which trended upwards most of the period. We continue with some charts created with the `plot()` method to illustrate 2007 and 2017 results for all strategies.

Figure 3.7: Market (red) and CPPI and DPPI results (blue) for Cal-07 and Cal-17

In the rising Cal-07 contract on the left side of Figure 3.7, there are minor differences between the CPPI and the DPPI. For Cal-17 however, the DPPI successfully move the target level down and secure a lower price for the portfolio. In Figure 3.8, all strategies have a reasonable target level protection for the increasing Cal-07 contract. For the Cal-17, the OBPI breach the target level, and the SLPI shows an example of the undesired target level lock-in. In this test, the SHPI achieved the lowest portfolio price for Cal-17 by steadily locking in volumes in a falling market. When f_t increased towards the end of trading, SHPI was not strongly affected due to the already high hedge rate.

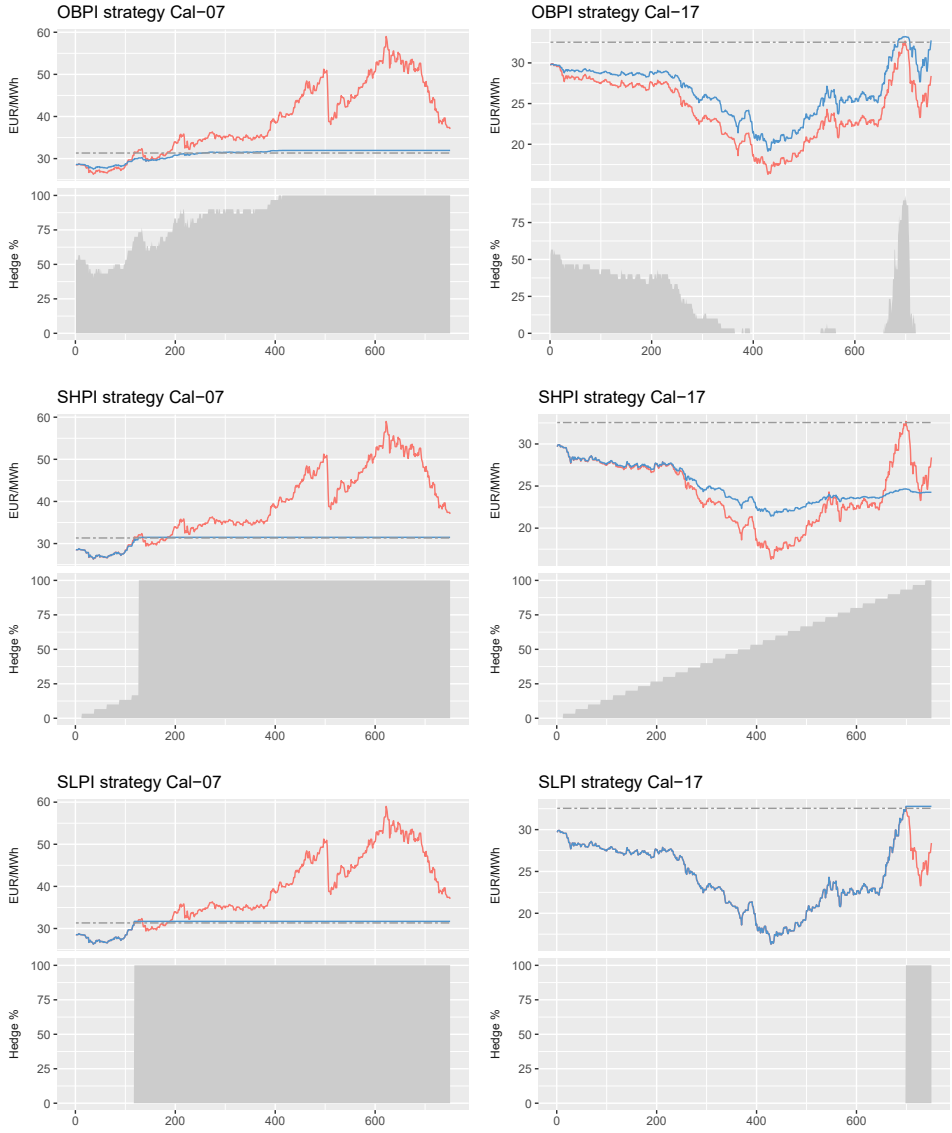

Figure 3.8: Market (red) and OBPI, SHPI and SLPI results (blue) for Cal-07 and Cal-17

Most large buyers of energy have a strategy for handling market price risk. They might all run some hedging scheme, but the underlying reasons and motivation for doing so can vary. Furthermore, their risk preferences are not the same. An energy supplier about to sign a fixed price agreement with an industrial client will require a high degree of price certainty in order not to get exposed to the risk associated with future price development. A large real estate owner able to pass on some costs to tenants might not need the same degree of predictability. The key point in this respect is that the nature of the price risk can depend on individual characteristics of the owner of the energy portfolio, who should select a strategy that matches the owners needs and risk preferences. The evaluation metrics we defined above can serve as inputs for such decisions. The calculated test results can be found in Table 3.3 and Table 3.4.

Strategy	Contract	p_T/f_T	Max(p_t/f_t)	p_T/p_0^*	Max(p_t/p_t^*)	Churn
CPPI	Cal-06	0.821	1.086	0.994	0.994	4.8
CPPI	Cal-07	0.844	1.043	0.998	0.998	3.1
CPPI	Cal-08	0.758	1.006	0.998	0.998	2.4
CPPI	Cal-09	1.213	1.29	0.998	0.998	3.9
CPPI	Cal-10	1.328	1.701	0.979	0.988	9.4
CPPI	Cal-11	0.713	1.199	0.998	0.998	7.7
CPPI	Cal-12	1.153	1.184	0.971	0.992	4.9
CPPI	Cal-13	1.158	1.179	0.901	0.957	8.8
CPPI	Cal-14	1.217	1.217	0.822	0.926	9.9
CPPI	Cal-15	1.155	1.155	0.847	0.913	8.9
CPPI	Cal-16	1.232	1.266	0.66	0.929	7.3
CPPI	Cal-17	1.097	1.205	0.957	0.965	6.7
CPPI	Cal-18	0.869	1.132	0.998	0.998	5.3
CPPI	Cal-19	0.461	1.034	0.997	0.997	2.5
CPPI	Cal-20	0.815	1.057	0.997	0.997	3.4
CPPI	Cal-21	1.379	1.752	0.852	0.948	10
DPPI	Cal-06	0.79	1.087	0.957	0.994	3.4
DPPI	Cal-07	0.813	1.049	0.961	0.996	2.6
DPPI	Cal-08	0.756	1.006	0.995	0.997	2.2
DPPI	Cal-09	1.163	1.237	0.956	0.999	2.5
DPPI	Cal-10	1.304	1.748	0.961	0.998	10.1
DPPI	Cal-11	0.649	1.233	0.908	0.998	4.7
DPPI	Cal-12	1.142	1.172	0.962	0.991	5.8
DPPI	Cal-13	1.165	1.198	0.906	0.96	7.9
DPPI	Cal-14	1.27	1.27	0.858	0.96	8.5
DPPI	Cal-15	1.183	1.219	0.868	0.968	7.4
DPPI	Cal-16	1.422	1.533	0.761	0.97	8.4
DPPI	Cal-17	0.924	1.473	0.807	0.997	7.3
DPPI	Cal-18	0.814	1.165	0.934	0.997	3.6
DPPI	Cal-19	0.451	1.034	0.975	0.997	2.5
DPPI	Cal-20	0.777	1.065	0.952	0.997	2.3
DPPI	Cal-21	1.371	2.705	0.847	1.002	7.3

Table 3.3: Strategy evaluation CPPI and DPPI

Strategy	Contract	p_T/f_T	Max(p_t/f_t)	p_T/p_0^*	Max(p_t/p_t^*)	Churn
OBPI	Cal-06	0.834	1.072	1.01	1.01	5.5
OBPI	Cal-07	0.862	1.047	1.02	1.02	3.8
OBPI	Cal-08	0.775	1.007	1.02	1.02	3.9
OBPI	Cal-09	1.218	1.232	1.002	1.056	7.7
OBPI	Cal-10	1.207	1.294	0.89	1.006	7.1
OBPI	Cal-11	0.782	1.152	1.094	1.094	12.2
OBPI	Cal-12	1.145	1.15	0.965	1.022	9.3
OBPI	Cal-13	1.111	1.117	0.864	0.955	5.5
OBPI	Cal-14	1.13	1.13	0.764	0.925	5.6
OBPI	Cal-15	1.109	1.12	0.814	0.913	4
OBPI	Cal-16	1.159	1.182	0.62	0.929	4.9
OBPI	Cal-17	1.154	1.188	1.007	1.022	6.1
OBPI	Cal-18	0.958	1.115	1.099	1.099	9
OBPI	Cal-19	0.475	1.033	1.027	1.027	4.1
OBPI	Cal-20	0.838	1.057	1.026	1.026	4.1
OBPI	Cal-21	1.174	1.346	0.726	0.949	5.5
SHPI	Cal-06	0.827	1.06	1.002	1.002	1
SHPI	Cal-07	0.85	1.071	1.005	1.005	1
SHPI	Cal-08	0.767	1.017	1.009	1.009	1
SHPI	Cal-09	1.221	1.299	1.005	1.005	1
SHPI	Cal-10	1.363	2	1.005	1.005	1
SHPI	Cal-11	0.68	1.028	0.951	0.964	1
SHPI	Cal-12	1.192	1.23	1.005	1.005	1
SHPI	Cal-13	1.107	1.168	0.861	0.996	1
SHPI	Cal-14	1.196	1.196	0.808	0.918	1
SHPI	Cal-15	1.07	1.127	0.785	0.912	1
SHPI	Cal-16	1.51	1.729	0.808	0.939	1
SHPI	Cal-17	0.854	1.313	0.746	0.919	1
SHPI	Cal-18	0.879	1.145	1.009	1.009	1
SHPI	Cal-19	0.473	1.115	1.023	1.023	1
SHPI	Cal-20	0.826	1.037	1.012	1.012	1
SHPI	Cal-21	1.215	2.389	0.75	0.954	1
SLPI	Cal-06	0.849	1.047	1.028	1.028	1
SLPI	Cal-07	0.856	1.078	1.012	1.012	1
SLPI	Cal-08	0.77	1.021	1.013	1.013	1
SLPI	Cal-09	1.225	1.303	1.008	1.008	1
SLPI	Cal-10	1.365	2.002	1.006	1.006	1
SLPI	Cal-11	0.724	1.079	1.013	1.013	1
SLPI	Cal-12	1.19	1.227	1.003	1.003	1
SLPI	Cal-13	1.303	1.389	1.013	1.013	1
SLPI	Cal-14	1	1	0.676	0.919	0
SLPI	Cal-15	1	1	0.734	0.912	0
SLPI	Cal-16	1	1	0.535	0.941	0
SLPI	Cal-17	1.154	1.408	1.007	1.007	1
SLPI	Cal-18	0.875	1.139	1.005	1.005	1
SLPI	Cal-19	0.475	1.119	1.026	1.026	1
SLPI	Cal-20	0.824	1.034	1.009	1.009	1
SLPI	Cal-21	1	1	0.618	0.954	0

Table 3.4: Strategy evaluation OBPI, SHPI and SLPI

Market participants needing a high degree of price predictability could start by inspecting the column p_T/p_0^*. From Table 3.3 we can verify that none of the initial target price levels were breached for the CPPI and the DPPI. When looking closer at the column $\text{Max}(p_t/p_t^*)$, we do see one minor breach of 0.2 percent of the downwards adjusted target price for DPPI in Cal-17, but the portfolio price is well below initial target p_0^*. Churn rates are fairly high for both, and the portfolio manager will depend on sufficient market liquidity in order to execute the strategy. We also keep in mind that calculated prices in the analysis are including transaction costs. When comparing the two, we see that DPPI achieved a lower p_T/p_0^* in 12 out of 16 contracts. It also achieved a lower portfolio price relative to the market p_T/f_T in 12 out of 16 cases. The premium paid via hedging can in some years be large for both. For example, in the downwards trending Cal-16 contract, DPPI maintain a fairly high hedge rate due to the lowering of the target level, thus ending up with a portfolio price 42.2 percent above the market. As the CPPI does not readjust the target price level, the hedge rate is lowered during the trading period, creating a smaller premium versus the market price. Finally, by studying $\text{Max}(p_t/f_t)$, we see that there will be periods were ability to meet margin requirements and manage cash flows might be subject to a stress test. For the Cal-21 contract, the DPPI strategy had a trading day where the locked-in portfolio price was 2.705 times higher than the market price, possibly putting a strain on liquidity and cash management.

Similar comparisons can be made for OBPI, SHPI and SLPI by studying the results in Table 3.4. Here, we see that these strategies have not offered the same level of target price protection in this test. The largest breach of target level can be found for OBPI in the Cal-11 contract, where p_T ended 9.9 percent above p_0^*. In cases where the market price oscillate close to the strike price level K, the OBPI will make large adjustments to the hedge rate, and accumulate costs as well as exposing the portfolio to gap risk. Formally this effect can be studied by calculating the second order greek *Gamma*, which give the rate of change for the option delta based on changes in f_t. As the Cal-11 varied close to the 40 EUR/MWh level a significant part of the trading period, large adjustments were made in hedge rate and the churn rate ended at 12.2. Since the SHPI and SLPI strategies do not resell volumes, their churn rates are capped at 1. Looking at the portfolio price versus the market at expiry, the worst performance was achieved by the SHPI for Cal-16, where the portfolio price ended 51 percent above the market. The largest strain on margin requirements and cash management would also be given by the SHPI, where the largest difference between portfolio and market for Cal-21 was by a factor of 2.389.

This concludes our illustrative example of how the **etrm** package may be used for managing energy price risk. The case we have been studying has been fairly simple. In practical applications, the strategies would have to be used on multiple contracts match-

ing the volume prognosis more closely. The portfolio manager would also need to take other risk factors into account, such as basis risk. In the case for Nordic electricity, this can be handled via *Electricity Price Area Differentials (EPAD)* contracts. There will also be additional focus on issues such as volume, profile and currency risk, and regulatory requirements. Lack of market liquidity for certain types of instruments is also a challenge practitioners often face. Energy markets are complex. It is our ambition that the **etrm** package provide some useful open source tools that may help decision makers and analysts to implement, evaluate and use alternative strategies for handling price risk management.

Chapter 4

Weather derivatives

In this chapter, we start with a brief discussion of weather risk and alternative tools for risk mitigation. In the following section, we provide a short overview of temperature-based derivatives, including common temperature indices, active markets and contract types. We proceed with a literature review of alternative approaches to pricing these instruments. Finally, we introduce the idea of using such contracts for managing ambient air pollution risk, which is the topic for **Paper C** in the thesis.

4.1 Weather risk

The volatile and unpredictability nature of weather variations creates weather risks, and this may have a significant financial impact on firms. Some industries are more exposed than others. According to the Weather Risk Management Association, the majority of inquiries regarding weather derivatives come from the energy, agricultural, retail, construction, and transportation sectors, [WRMA, 2011]. Unfavorable weather conditions such as abnormally mild winters, cool summers, and periods of protracted drought or low wind can affect customer demand, and a producers ability to deliver the desired quantity of goods to be sold. For an electric power plant, the mild winter would typically lead to a reduction in energy to be delivered on the grid. In general, weather is primarily considered a volume risk amongst commercial enterprises.

While insurance has long been used by corporations and governments to control the risk of catastrophic weather occurrences [Mills, 2005], weather derivatives have gained in importance as a risk management tool for uncertainty in cash flows caused by non-catastrophic events. The instruments are written on indices linked to the sources of weather exposure, such as temperature, rainfall, snowfall and wind. Unlike insurance,

the holder of the contract does not need to file a claim and provide documentation of a loss, the payoff is automatically paid out according to contract specifications. The instruments are freely bought and sold in the financial markets, where fundamental players can transact with each other. Insurance companies and participants from capital markets, such as hedge funds and investment banks are also active. Weather derivatives may be traded on exchanges, or over-the-counter (OTC). In OTC trading the parties in the transaction also have flexibility to define structures that are tailor-made to match a specific company's exposure.

4.2 Temperature derivatives

Temperature-based contracts are by far the most popular and actively traded of all the weather derivatives currently available [WRMA, 2011]. A possible explanation for this might be that temperature is the weather phenomenon that represent a risk for most companies, irrespective of industry. Temperature can also be closely related to other sources of weather exposure, like rainfall, snowfall and wind. According to [Thind, 2014], the first transaction containing a weather component was made in 1996 between Aquila Energy and Consolidated Edison. This dual-commodity trade combined a purchase of electric power with a temperature clause. The first pure weather trade was also based on temperature. It took place in Milwaukee for the winter 1997-1998, with Enron and Koch Industries. The Chicago Mercantile Exchange (CME) introduced exchange-traded weather derivatives contracts later that year, see https://www.cmegroup.com/trading/weather/.

4.2.1 Settlement indices

The most frequently used temperature indices for derivatives contracts are the cumulative average temperature (CAT), the cooling degree-day (CDD) and the heating degree-day (HDD). These are based on the daily average temperature.

Definition 5 (Daily average temperature) *Let $T(t)^{max}$ and $T(t)^{min}$ denote the maximum and minimum temperatures measured at a specific location for day t. The daily average temperature for day t is*

$$T(t) = \frac{T(t)^{max} - T(t)^{min}}{2} \tag{4.1}$$

In practical applications, it is of vital importance that the temperature measurements

are of high quality, and trusted by both sides of the transaction. The time series most commonly used are published by neutral third parties, such as meteorology institutions. For a given measurement period $[\tau_1, \tau_2]$, the settlement indices are calculated with

$$CAT(\tau_1, \tau_2) = \sum_{t=\tau_1}^{\tau_2} T(t) \tag{4.2}$$

$$CDD(\tau_1, \tau_2) = \sum_{t=\tau_1}^{\tau_2} max\{T(t) - b, 0\} \tag{4.3}$$

$$HDD(\tau_1, \tau_2) = \sum_{t=\tau_1}^{\tau_2} max\{b - T(t), 0\} \tag{4.4}$$

where $T(t)$ is the average daily temperature for day t, and b is a base temperature level, normally set at $18^oC/\ 65^oF$. The CAT index considers the accumulated average temperature during a specific period. The degree-days indices focus on temperatures breaching a pre-specified threshold. By accumulating the days above the base level, the CDD index quantifies the need for cooling during the period. Similarly, the HDD reflects the need for heating. HDD instruments can be traded in the winter season to hedge against risks related to unusually low temperatures. CDD contracts may mitigate risks during the summer season due to abnormal heat. Winter contracts typically cover November 1 through March 31, and May 1 through September 30 is often used to define the summer season. Contracts covering other time spans such as week, month and quarter are also used. Instruments built on these indices may also be combined to construct more complex payoff structures.

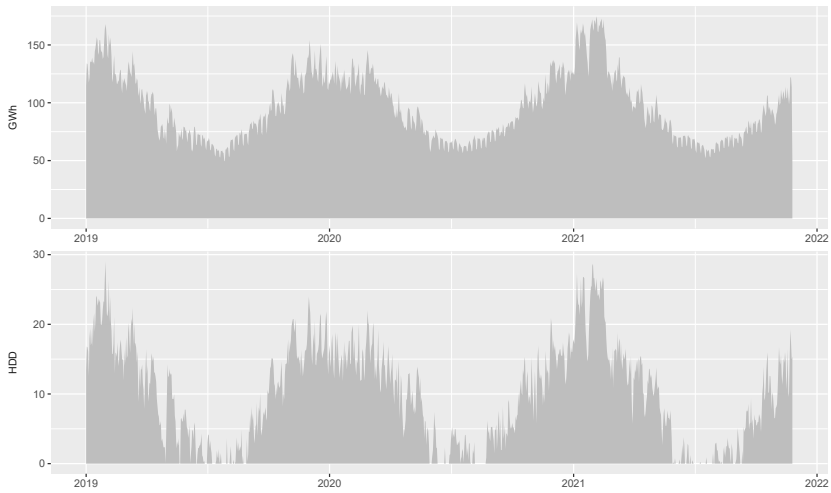

Figure 4.1: Electricity consumption and HDD index NO1 area

To illustrate a use case for temperature derivatives, we look closer at some data from the Nordic power market. The upper part of Figure 4.1 present the daily total power consumption in GWh for the Nord Pool power exchange bidding area NO1 (central eastern Norway). The bottom panel contain the HDD index calculated in Celsius with (4.4) on temperature observations from the metering station at Blindern, Oslo. The relationship between energy demand and temperature is quite strong, and a consumer could hedge against events such as the early 2021 cold spell with a HDD-based contract.

It is worth noting that contracts written on the indices presented here are different from the traditional derivatives used in equities and interest rate markets. The underlying settlement reference, the temperature index, is non-tradable. This makes the weather derivatives market *incomplete*. We cannot replicate the instruments by trading in the underlying, as we did when delta hedging the call option in Chapter 3. We will discuss the implications of this when exploring strategies for valuation below.

4.2.2 Markets and contract types

After the Chicago Mercantile Exchange established the first exchange-traded weather contracts in 1999, the market for these products increased. In the following years, the CME introduced derivatives linked to metering stations in a number of cities in the US, Europe, Canada, Australia, and Japan. The majority of these financial instruments measured cooling or heating degree days. Weather derivatives were also listed on other market venues such as the The Intercontinental Exchange (ICE). At the time of writing, the number of locations available at the CME has declined. The exchange currently offer contracts for nine US cities, two European cities, and one Japanese city [CME, 2021]. The market for weather derivatives is relatively illiquid, and bid/ ask spreads are large. The weather services marketed by ICE at present date are not standardized exchange traded futures, but rather access to the OTC derivatives quotes in the weather data feed on the ICE Chat. Even though weather represent a risk to many participants in the modern economy, the need for customization can be quite high, making the standardized contracts listed at exchanges less relevant for end users. The two European cities available for temperature trading at the CME are London and Amsterdam. A Norwegian power producer fearing a mild winter will not be able to implement a satisfactory hedging scheme with instruments based on indices from these locations. This *basis risk* is simply to great, and market players resort to bilateral agreements traded OTC with specialized institutions, such as large energy traders and reinsurance companies. In [Till, 2014] it is argued that this type of market participants are the ones best suited to manage weather risk, due to their ability to tailor contracts to the needs of the counterparty.

A disadvantage of trading off-exchange is the exposure to counterparty risk. In the event that the seller of the temperature derivative is not able to meet contractual obligations, the buyer might have to face a significant loss. This risk can be mitigated with bilateral collateral requirements. Some of the significant players in the OTC weather derivatives market can be found amongst the participants in the Weather Risk Derivative Survey conducted by the Weather Risk Management Association [WRMA, 2011]. The list include major reinsurance companies and energy traders, such as Swiss Re, Centrica, EDF Trading and Munich Re.

The most commonly used contract types in temperature markets are futures, swaps and options based on the indices defined above. The futures contracts are agreements to buy or sell the value of the selected index at a specific date in the future. For exchange traded instruments, there is also daily cash settlement based on the mark-to-market value of the position. In swap contracts, the two parties agree to exchange risk over a predefined time period, where one side pays a fixed price, while the other pays a variable price. Swaps can often be multi-period, involving a series of future payments. Finally, when using an option contract, the holder pays a premium to the issuer up front. If the index breaches an agreed strike level, the holder receive a payoff. The contracts normally include a maximum payoff level. Even though most trades are non-standard OTC agreements, they share some common attributes, such as the contract type, temperature index, and measuring station.

We will illustrate with an example. Consider a large power consumer located in south-east Norway who wants to protect itself against a spike in energy demand due to an unusually cold winter. The consumer decides to buy a HDD call option for the period December - February:

- Contract type: Call option
- Contract period: 01.12.2021 - 28.02.2022
- Measuring station: Blindern, Oslo
- Weather variable: Temperature (degrees in $°C$)
- Settlement index: HDD
- Strike level: 1.890 HDD
- Tick size: 10.000 EUR
- Maximum payout: 5M EUR

If the index passes the strike level, the option will give a payout. The 1.890 HDD strike specified in the contract is equivalent to an average temperature of $18 - 1.890/90 = -3$ degrees Celsius during the 90 day period. One tick corresponds to one heating-degree-

day, and the tick size determine the amount to be paid out by the contract. The profit function for the power consumer can now be specified as:

$$\text{Payoff} = \text{Min}\big[(\text{Maximum payout - Premium}), (\text{Tick size} \times \text{Max}(\text{HDD} - \text{Strike}, 0) - \text{Premium})\big]$$

Assuming the option premium paid is 835.930 EUR and the index value turn out to be 2.300 HDD, the consumer will receive:

$$\begin{aligned}\text{Payoff} &= \text{Min}\big[(5.000.000 - 835.930), (10.000 \times \text{Max}(2.300 - 1.890, 0) - 835.930)\big] \\ &= 3.264.070 \text{ EUR}\end{aligned}$$

We can also calculate a break-even level for the HDD index:

$$\begin{aligned}10.000 \times (\text{HDD}^* - 1890) - 835.930 &= 0 \\ \text{HDD}^* = (835.930/10.000 + 1.089) &= 1.973,59\end{aligned}$$

This value corresponds to an average temperature of $18 - 1973.59/90 = -3,93$ degrees Celsius. If the average temperature over the three months falls below this level, the consumer will obtain a positive financial result from the temperature hedge. In the event that the temperature average falls in between the break-even level and -3 degrees, the option will be exercised, but the payout will not cover the option premium. If the average temperature stays above -3 degrees, the contract will expire worthless. The consumer can now hedge volume risk by entering into HDD options. The commodity price may also be hedged, for example in the electricity futures market.

4.3 Valuation of temperature derivatives

Since the weather indices presented above are non-tradable, the market for weather derivatives is said to be *incomplete*. The no-arbitrage approach often used when pricing traditional financial derivatives cannot be utilized. In a complete market, we can replicate the payoff structure of a derivative by establishing a self-financing portfolio consisting of the underlying and a riskless asset. Via arbitrage arguments, one can arrive at a price for the instrument by evaluating the cost of establishing the replicating portfolio, see for example [Bjork, 2009]. In weather markets, this is not possible.

In the literature review by [Schiller et al., 2012], the methods for valuation of weather derivatives are divided into three main categories: *burn analysis*, *index modelling* and *daily temperature modelling*. There are however a number of initiatives that does not fit into this grouping. Therefore, we follow [Benth and Saltyte-Benth, 2012] and add a

fourth category: *utility-based approaches*. These four will be described further below.

The simplest method for pricing weather derivatives is the so-called historical *burn analysis*. This is an actuarial approach, where the price is estimated by simply calculating the payoff of the instrument on historical temperature data. The time series may require detrending, due to global warming and the heating effects from increased urbanization described in [Alaton et al., 2002]. The derivative price is calculated as the expected value of historical payoffs, often adjusted upwards with a risk premium [Jewson and Brix, 2005]. A downside of this approach is the loss of of information when only relying on the historical payoff calculations. It also require a fairly long temperature history, and prices obtained typically have large standard deviation. If the weather has been subject to change for example due to global warming, parts of the historical data may also be less relevant.

Index modelling is yet another actuarial approach, where historical data is used to estimate distributions for the relevant settlement index. Rather than being limited to the historical data, we can perform a large number of simulations to estimate the average payout, [Jewson and Brix, 2005]. Assuming a good estimate can be achieved, this approach may provide more stable price estimation compared to burn analysis, [Schiller et al., 2012]. A benefit shared by both methods is the conceptual simplicity, but it may be challenging to fit a suitable distribution for the index.

The work presented in [Cao and Wei, 2004] is an example of *utility-based approaches*. They extend the equilibrium pricing framework from [Lucas, 1978], and include temperature as another source of uncertainty in the economic environment. Weather is related to the aggregated output in the economy, and numerical analysis finds the market price of risk for temperature derivatives to be significant. A strength of this approach is that it enables estimation of the market price of risk, but the framework relies on a host of assumptions and would be complex to implement. Another example is [Davis, 2001], where pricing is investigated using a marginal-utility approach, where participants in the economy have different exposures to weather, and will only procure hedging instruments if they gain an increase in utility. The HDD index and commodity prices are modeled as geometric Brownian motions, allowing closed form expressions for swaps and options. Also this model relies heavily on model assumptions, such as specification of utility functions. Several other contributions can be placed in this category. [Platen and West, 2004] suggest an approach to pricing which is based on the existence of an optimal benchmark portfolio, and [Brockett et al., 2006], proposes an indifference pricing approach focusing on portfolio effects and possible hedging strategies.

The final category of pricing methods is based on *daily temperature modelling*. By creating a model for the daily average temperature (DAT), a large number of temperature series may be simulated and used for valuation of the derivatives contracts. In some situations, closed form pricing formulas may be derived. A considerable benefit of using this approach is that all contracts at a single location may be priced using a single model. This ensures consistency in contract prices. DAT modelling has gained popularity in the weather derivatives literature, and there is a large number of studies and alternative model specifications, see for example [Benth and Saltyte-Benth, 2012], [Alaton et al., 2002] and [Dornier and Queruel, 2000].

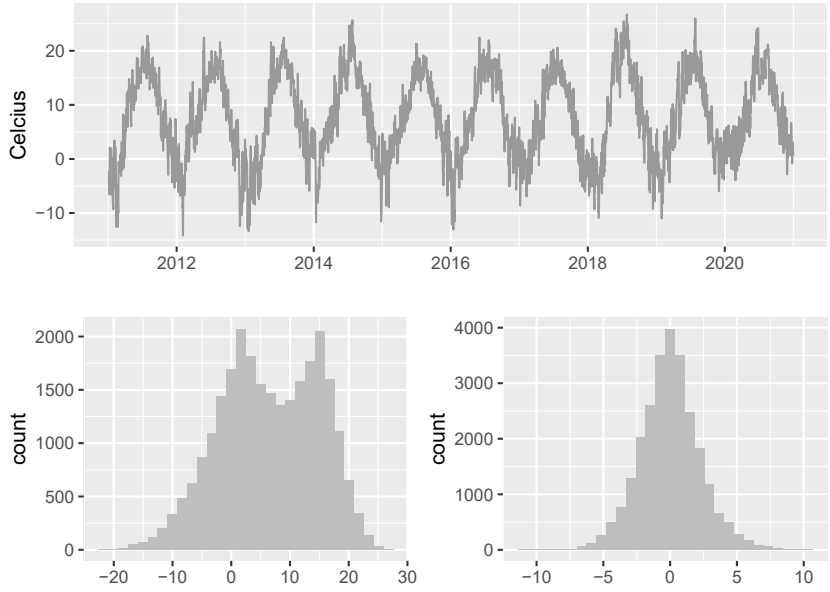

Figure 4.2: Daily average temperature, Blindern

As illustrated in Figure 4.2, the DAT is highly seasonal. The upper part of the chart show daily temperatures the last ten years from the Blindern station in Oslo. The bottom left histogram reveal the bimodality of temperature, whereas the bottom right chart can be used to evaluate the distribution for daily temperature changes. These have been made with data from the period 1950-2020. It's worth noting that the temperature data naturally will vary depending on the location of the measuring station. The temperature *level* in Norway will be quite different from readings from Morocco, but also the magnitude of fluctuations may be different.

Based on the empirical properties of temperature, it is typically modelled as a mean reverting process with seasonality. In [Dornier and Queruel, 2000], the model consist of

two parts: a triginometric seasonal function and a random component. The seasonal mean function is on the form

$$T_t^m = A + Bt + Csin(\omega t + \varphi) \tag{4.5}$$

where t denotes time and $\omega = 2\pi/365$. This specification allows a trend in temperature, and the phase angle φ account for the fact that the coldest period of the year is not necessarily at January 1. The mean-reverting process is modeled with an Ornstein-Uhlenbeck process, which fluctuate around the seasonal mean. The volatility of temperature is assumed constant, and the innovations are assumed to follow a Brownian motion. [Alaton et al., 2002] improve this model by allowing non-constant volatility. Here, the DAT process is defined as

$$dT_t = \left[\frac{dT_t^m}{dt} + \alpha(T_t^m - T_t)\right]dt + \sigma_t dW_t \tag{4.6}$$

where σ_t is a piecewise constant function with a fixed value for each month. The term $\frac{dT_t^m}{dt}$ ensures the process reverts to the mean in (4.5), and a solution is obtained via Ito's Lemma. The solution when starting at temperature T_s is

$$T_t = (T_s - T_t^m)e^{-\alpha(t-s)} + T_t^m + \int_s^t e^{-\alpha(t-\tau)}\sigma_\tau dW_\tau \tag{4.7}$$

[Alaton et al., 2002] use an approximation for the HDD index and the fact that (4.6) is a Gaussian process to derive a pricing formula for HDD options. They argue that this approach holds given that the probability of $\text{Max}(18 - T_t, 0) = 0$ is very low in the winter season. They also price the HDD options using monte carlo simulation, and find that results are similar, but sensitive to the assumed market price of risk.

In [Brody et al., 2002] it is pointed out that temperature fluctuations around the seasonal mean are Gaussian, with a long-range dependence. Therefore, they argue a fractional Brownian motion is better suited for the noise driving process in the Ornstein-Uhlenbeck process:

$$dT_t = \alpha(T_t^m - T_t)dt + \sigma_t dW_t^H \tag{4.8}$$

where W^H is a fractional Brownian motion. The solution is again found using Ito's Lemma, and pricing methods are developed for derivatives written on the HDD and CAT indices. One weakness in this approach pointed out by [Benth and Šaltytė-Benth, 2005] is that (4.8) lack the $\frac{dT_t^m}{dt}$ term which is needed to ensure reversion to the seasonal mean. In the last mentioned paper, a structure similar to (4.6) is used, but with Levy processes for the noise term. In [Schiller et al., 2012] a spline model is proposed for modelling the DAT from a number of weather stations across the United States,

whereas [Caballero et al., 2002] rely on ARMA and ARFIMA models. Campbell and Diebold [2005] suggest an autoregressive time series model with linear trend and a sum of trigonometric functions for describing the seasonality. By regressing the deseasonalized and de-trended temperature against a large number of lagged observations, it differs from the widely used Ornstein-Uhlenbeck approach. They find a large seasonal variation in the squared residuals of the model, which they suggest to handle with an ARCH approach.

[Benth et al., 2007b] suggest a continuous-time autoregressive model (CAR) with seasonal variation for modelling the daily average temperature. They let $\mathbf{X}(t)$ be a stochastic process in $\mathbb{R}^{\mathrm{p}}$ for $p \geq 1$ defined by the vectorial Ornstein-Uhlenbeck equation

$$d\mathbf{X}(t) = A\mathbf{X}(t)dt + \mathbf{e}_p\sigma(t)dW(t) \tag{4.9}$$

where $\mathbf{e}_k$ is the kth unit vector in $\mathbb{R}^{\mathrm{p}}$, $k = 1, ..., p$ and $\sigma(t) > 0$ is a real-valued and square integrable function. A is the $p \times p$ matrix

$$A = \begin{bmatrix} 0 & 1 & 0 & \dots & 0 \\ 0 & 0 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & 1 \\ -\alpha_p & -\alpha_{p-1} & -\alpha_{p-2} & \dots & -\alpha_1 \end{bmatrix} \tag{4.10}$$

where $\alpha_k, k = 1, ..., p$ are constants. In order to represent $\mathbf{X}(t)$ explicitly, the multidimensional Ito's Lemma is applied to (4.9), and we get:

$$\mathbf{X}(s) = \exp(A(s-t))\mathbf{X}(t) + \int_t^s \exp(A(s-u))\mathbf{e}_p\sigma(u)dW(u) \tag{4.11}$$

for $s \geq t \geq 0$ and $\mathbf{X}(t) \in \mathbb{R}^{\mathrm{p}}$. Then, the model for the daily average temperature dynamics can be specified with

$$T(t) = \Lambda(t) + X_1(t) \tag{4.12}$$

where $\Lambda(t)$ is a seasonal function similar to (4.5) and the stochastic component $X_1(t)$, represented by a CAR(p) process is

$$X_1(t) = \mathbf{e_1}^T\mathbf{X}(t) \tag{4.13}$$

where X_q is the qth coordinate of the vector $\mathbf{X}$, $q = 1, ..., p$ and $\mathbf{e_1}^T$ denote the transpose of $\mathbf{e_1}$. The seasonal volatility in the temperature residual process is modeled with a

truncated Fourier series, which is also similar to (4.5), but without a trend. [Benth et al., 2007b] estimate CAR(3) models on temperature data from Stockholm, Sweden. They find a pronounced variation in volatility during the year, and the highest levels are reached during the winter season and early summer. The modelling framework also allow them to provide derivations of explicit prices for several futures and options contracts.

The weather, in particular temperature, represent a source of risk for many participants in the economy. This exposure is likely to increase as the global warming trend lead to more volatile weather patterns. In order to hedge against the financial risk, we need to be able to design and price instruments at the local level, where a specific economic activity is being affected. There is another phenomenon closely linked to the weather that exhibit even stronger spatial variation. The level of *ambient air pollution* can vary greatly within a country, but the severe cases are often found in highly urbanized areas, where a significant proportion of value creation is taking place. The magnitude of the issue is closely monitored by public agencies, and typically reported to the public using a standardized scale, such as the *Air Quality Index (AQI)*, [EPA, 2009], [EPA, 2021]. The problem is particularly urgent in the mega cities of Asia, where rapid industrialization and urbanization have led to a dramatic increase in air pollution. Many of these cities have some of the highest pollution levels in the world, and the problem is compounded by factors such as population density and lack of green spaces. These cities play an important part in the global value chain, and they have millions of citizens. There are numerous examples of extreme pollution incidents, such as the lock-downs and serious pollution episodes in Beijing [An et al., 2007] and Delhi [Kumar et al., 2015].

What are the financial consequences of a pollution alert lock-down in Beijing, a city with more that 22 million inhabitants? What if the financial risks arising from extreme urban air pollution could be managed with instruments similar to the temperature derivatives? This is the topic of **Paper C**.

Chapter 5

Computer code

The R programming language [R Core Team, 2021] has been used to write all of the computer code in this thesis. Most programming effort has been put into the development of the package **etrm** presented in **Paper B**. The package is available on CRAN, and can be installed and loaded into the R environment by running the following commands:

```
install.packages("etrm")
library(etrm)
```

There are six functions in the package. These are `msfc()`, `cppi()`, `dppi()`, `obpi()`, `shpi()` and `slpi()`.

`msfc()`: Calculates the *Maximum Smoothness Forward Curve* (MSFC) for energy forward contracts with flow delivery, as described in [Ollmar, 2003] and [Benth et al., 2007a]. The function arguments are: A date for when to calculate the curve, a Boolean vector indication if a available contract should be included in the calculation, two vectors with start and end dates for the contracts, a price vector, and finally a prior for the price curve as an optional argument. It returns an S4 object of type `MSFC` which contains details regarding the calculation method as well as calculation results stored in a data frame. There methods `plot()`, `summary()` and `show()` are has been implemented for the class for visualization, analysis and reporting.

The remaining five functions offer alternative forward trading strategies for hedging purposes. They share a number of arguments, such as the volume to be hedged, a date vector with trading dates, the futures price vector, assumed trading cost, and a Boolean argument deciding whether the lot size of traded volume should be restricted to whole numbers. Sharing of properties and functionality is handled with inheritance in the

etrm implementation. The strategy classes inherit most of their fields from a parent class `GenericStrat`, which has the methods `plot()`, `summary()` and `show()` associated with it. The strategy functions are:

`cppi()`: Implements the *Constant Proportion Portfolio Insurance* (CPPI) strategy for setting the portfolio hedge rate, see [Perold, 1986] and [Black and Jones, 1987]. Requires the additional arguments `tper` and `rper` to set the forward market allocation. The former is used to calculate a target price level to be protected, the latter for calculating the risk factors, which again affects the gearing of the strategy. The function returns an instance of the `CPPI` class.

`dppi()`: Implements the *Dynamic Proportion Portfolio Insurance* (DPPI) strategy for price risk management [Lee et al., 2008]. Arguments are similar to `cppi()`, but the target price can be subject to adjustments by the DPPI model, and `rper` is may be vector with (non-constant) risk factors. An instance of the `DPPI` is returned as output.

`obpi()`: Implements a synthetic *Option Based Portfolio Insurance* (OBPI) strategy via delta hedging, see [Black, 1976] and [Bjork, 2009]. Required additional arguments are the option strike price level, assumed volatility, risk free rate of interest, trading days per year and days to expiry. The `obpi()` aim to protect the strike price level by entering into positions in the underlying futures contract using the Black-76 option pricing model to determine hedge rates. The expected target level is hence set implicitly for this strategy, where the portfolio price is not expected to breach a threshold given by the strike, adjusted for option premium. The `obpi()` act as a constructor for the `OBPI` class.

`shpi()`: Implements the *Step Hedge Portfolio Insurance* (SHPI) strategy, a mechanical benchmark approach where hedging positions are gradually built throughout the trading period, aiming to smoothen out the price of the hedged volume. The `shpi()` function require the `tper` argument to calculate a target price level. If this is reached, positions are locked and not opened again. The function returns an instance of class `SHPI`.

`slpi()`: Implements the *Stop Loss Portfolio Insurance* (SLPI) strategy for hedging, a simple benchmark where no positions are taken unless the target price level to be protected is reached. In this situation a full hedge is entered, and it is not opened again. An instance of the `SLPI` is returned as output.

The **etrm** package makes us of some other R packages, these are `ggplot2` [Wickham, 2016], `reshape2` [Wickham, 2007], and `testthat` [Wickham, 2011]. The project is actively maintained, and the code can be found at:

https://github.com/sleire/etrm.

In order to make results reproducible, we have put the R code used in this PhD-project in open source code repositories. These have been organized in accordance with the structure in the thesis. The interested reader can recreate results in Chapter 3 by running the code available at:

https://gitlab.com/sleire/energyrisk

If further details regarding results in **Paper A** or **Paper C** are of interest, the code can be found at:

https://gitlab.com/sleire/lgportf
https://gitlab.com/sleire/aqider

These Git repositories can be used to replicate results in the papers, or to run a similar analysis on new data sets.

Chapter 6

Summary of papers

6.1 Portfolio Allocation under Asymmetric Dependence in Asset Returns using Local Gaussian Correlations

Anders D. Sleire, Bård Støve, Håkon Otneim, Geir Drage Berentsen, Dag Tjøstheim, Sverre Hauso Haugen, (2021), Finance Research Letters, p. 102475.

It is well documented that there are asymmetric dependence structures between financial returns, and the correlations between certain asset classes may often strengthen during crisis periods. This asymmetry can weaken the diversification effect achieved when constructing a portfolio under the assumption that returns follow a joint-Gaussian distribution. In this paper we improve portfolio allocation by using a new non-parametric measure of local dependence, the local Gaussian correlation. We extend the classical mean-variance framework, and show that the portfolio optimization is straightforward using our new approach. Local covariance matrices are estimated using local Gaussian correlations for all pairs of assets in evaluation points that are selected via simple moving averages of the returns series. Using a data set consisting of monthly index returns from commodities, stocks and interest rate markets, the new method is shown to outperform the equally weighted ("1/N") portfolio and the classical Markowitz portfolio. Results are evaluated with standard performance metrics such as the Sharpe ratio, the Certainty Equivalent, the Sortino ratio and the Omega ratio. The need for portfolio re-balancing is also taken into consideration and performance is assessed both with and without transaction costs. The results in the paper indicate that standard methods for portfolio construction can be improved with this approach, and we suggest further analysis of alternative local covariance matrix calculation approaches.

6.2 etrm: Energy Trading and Risk Management in R

Anders D. Sleire (2022), The R-Journal, 14/1, p. 320.

This paper introduces the R package **etrm**, which was developed during the PhD project. Prices for commodities such as electricity and natural gas are typically subject to large volatility and seasonal variation. These markets have some unique characteristics that need to be considered when performing risk management tasks such as forward curve calculation and price risk hedging. The relevant tools are often bundled into proprietary Energy Trading Risk Management (ETRM) systems that are not openly available to the public. The **etrm** package offers a transparent solution for building forward price curves which is consistent with methods widely used in the industry. The package also provide implementations of five portfolio insurance trading strategies that may be used for price risk management with energy derivatives contracts. This paper serves as a general introduction of the package into the R ecosystem. It provides the necessary background from theory of finance, an overview of the package structure, available functions and data sets, and examples of use. The tools included in **etrm** can be used for tasks such as risk evaluations, analysis, back testing and decision support for trading.

6.3 Modelling and Pricing Air Pollution Derivatives

Anders D. Sleire, working paper.

An important part of the financial risk management discipline consist of handling market price risk, often using derivatives contracts. Many organisations also aim to manage *volumetric* risk. In industries such as energy, tourism, agriculture and construction, businesses manage volume risk by trading in financial instruments that are linked to the underlying source of this risk - the weather. By holding a weather derivatives contract protecting against abnormally low temperatures, a farmer can secure income lost due to failed crops. In this paper, we build upon weather derivatives theory and design contracts whose payoff depend on measured air pollution levels in Asian mega cities. A significant part of the global value chain is located in heavily polluted cities with large populations. We argue that many of the adverse financial effects of pollution incidents such as those in Beijing [An et al., 2007] and Delhi [Kumar et al., 2015] can be managed with instruments resembling the contracts traded in temperature markets. Air pollution data reported in accordance with the Air Quality Index (AQI) standard introduced by

the US Environmental Protection Agency are reported by government agencies and made available to the public. We develop stochastic models to describe the pollution dynamics for a selection of cities in China, and show how these can be used to price derivatives contracts written on AQI-based indices. The contracts may be used to manage financial risk related to public health, government imposed lock-downs and changes in consumer behaviour, to name a few examples. Some practical use cases are also presented and discussed.

Bibliography

K. Aas, C. Czado, A. Frigessi, and H. Bakken. Pair-copula constructions of multiple dependence. *Insurance: Mathematics and economics*, 44(2):182–198, 2009.

P. Alaton, B. Djehiche, and D. Stillberger. On modelling and pricing weather derivatives. *Applied mathematical finance*, 9(1):1–20, 2002.

X. An, T. Zhu, Z. Wang, C. Li, and Y. Wang. A modeling analysis of a heavy air pollution episode occurred in beijing. *Atmospheric Chemistry and Physics*, 7(12): 3103–3114, 2007. doi: 10.5194/acp-7-3103-2007. URL `https://acp.copernicus.org/articles/7/3103/2007/`.

L. Bachelier. Théorie de la spéculation. In *Annales scientifiques de l'École normale supérieure*, volume 17, pages 21–86, 1900.

G. Bampinas and T. Panagiotidis. Oil and stock markets before and after financial crises: A local gaussian correlation approach. *Journal of Futures Markets*, 37(12):1179–1204, 2017.

T. Bedford and R. M. Cooke. Probability density decomposition for conditionally dependent random variables modeled by vines. *Annals of Mathematics and Artificial intelligence*, 32:245–268, 2001.

T. Bedford and R. M. Cooke. Vines–a new graphical model for dependent random variables. *The Annals of Statistics*, 30(4):1031–1068, 2002.

F. E. Benth and J. Šaltytė-Benth. Stochastic modelling of temperature variations with a view towards weather derivatives. *Applied Mathematical Finance*, 12(1):53–85, 2005.

F. E. Benth and J. Saltyte-Benth. *Modeling and pricing in financial markets for weather derivatives*, volume 17. World Scientific, 2012.

F. E. Benth, S. Koekkebakker, and F. Ollmar. Extracting and applying smooth forward curves from average-based commodity contracts with seasonal variation. *The Journal of Derivatives*, 15(1):52–66, 2007a. URL `https://doi.org/10.3905/jod.2007.694791`.

F. E. Benth, J. Šaltytė Benth, and S. Koekebakker. Putting a price on temperature. *Scandinavian Journal of Statistics*, 34(4):746–767, 2007b.

F. E. Benth, J. S. Benth, and S. Koekebakker. *Stochastic modelling of electricity and related markets*, volume 11. World Scientific, 2008. URL `https://doi.org/10.1142/6811`.

G. D. Berentsen, T. S. Kleppe, and D. B. Tjøstheim. Introducing localgauss, an r package for estimating and visualizing local gaussian correlation. *Journal of Statistical Software*, 56(1):1–18, 2014.

P. Bjerksund, H. Rasmussen, and G. Stensland. Valuation and risk management in the norwegian electricity market. In *Energy, natural resources and environmental economics*, pages 167–185. Springer, 2010.

T. Bjork. *Arbitrage Theory in Continuous Time*. Oxford University Press, 3 edition, 2009. URL `https://EconPapers.repec.org/RePEc:oxp:obooks:9780199574742`.

F. Black. The pricing of commodity contracts. *Journal of financial economics*, 3(1): 167–179, 1976. URL `https://doi.org/10.1016/0304-405X(76)90024-6`.

F. Black and R. W. Jones. Simplifying portfolio insurance. *The Journal of Portfolio Management*, 14(1):48–51, 1987. URL `https://doi.org/10.3905/jpm.1987.409131`.

F. Black and R. Litterman. Global portfolio optimization. *Financial analysts journal*, 48(5):28–43, 1992.

F. Black and M. Scholes. The pricing of options and corporate liabilities. *Journal of political economy*, 81(3):637–654, 1973.

M. D. Bordo. An historical perspective on the crisis of 2007-2008. Technical report, National Bureau of Economic Research, 2008.

E. Bouyé, V. Durrleman, A. Nikeghbali, G. Riboulet, and T. Roncalli. Copulas for finance-a reading guide and some applications. *Available at SSRN 1032533*, 2000.

P. L. Brockett, M. Wang, C. Yang, and H. Zou. Portfolio effects and valuation of weather derivatives. *Financial Review*, 41(1):55–76, 2006.

D. C. Brody, J. Syroka, and M. Zervos. Dynamical pricing of weather derivatives. *Quantitative finance*, 2(3):189, 2002.

R. Caballero, S. Jewson, and A. Brix. Long memory in surface air temperature: detection, modeling, and application to weather derivative valuation. *Climate Research*, 21 (2):127–140, 2002.

R. Campbell, K. Koedijk, and P. Kofman. Increased correlation in bear markets. *Financial Analysts Journal*, 58(1):87–94, 2002.

S. D. Campbell and F. X. Diebold. Weather forecasting for weather derivatives. *Journal of the American Statistical Association*, 100(469):6–16, 2005.

M. Cao and J. Wei. Weather derivatives valuation and market price of weather risk. *Journal of Futures Markets: Futures, Options, and Other Derivative Products*, 24(11): 1065–1089, 2004.

D. Chambers, E. Dimson, and A. Ilmanen. The norway model. *The Journal of Portfolio Management*, 38(2):67–81, 2012.

J.-S. Chen, C.-L. Chang, J.-L. Hou, and Y.-T. Lin. Dynamic proportion portfolio insurance using genetic programming with principal component analysis. *Expert Systems with Applications*, 35(1):273–278, 2008. URL `https://doi.org/10.1016/j.eswa.2007.06.030`.

D. G. Clayton. A model for association in bivariate life tables and its application in epidemiological studies of familial tendency in chronic disease incidence. *Biometrika*, 65(1):141–151, 1978.

CME. Weather products, 2021. `https://www.cmegroup.com/trading/weather/` [Accessed: 2021-11-25].

R. Cont. Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative finance*, 1(2):223, 2001.

M. Davis. Pricing weather derivatives by marginal value. *Quantitative finance*, 1(3): 305–308, 2001.

S. Demarta and A. J. McNeil. The t copula and related copulas. *International statistical review*, 73(1):111–129, 2005.

G. Dionne. Risk management: History, definition, and critique. *Risk management and insurance review*, 16(2):147–166, 2013.

F. Dornier and M. Queruel. Caution to the wind. *Energy & power risk management*, 13 (8):30–32, 2000.

D. Duffie and J. Pan. An overview of value at risk. *Journal of derivatives*, 4(3):7–49, 1997.

P. Embrechts, F. Lindskog, and A. McNeil. Modelling dependence with copulas. *Rapport technique, Département de mathématiques, Institut Fédéral de Technologie de Zurich, Zurich*, 14:1–50, 2001.

EPA. A guide to air quality and your health. *United States Environmental Protection Agency*, 2009.

EPA. Reviewing national ambient air quality standards (naaqs): Scientific and technical information, 2021. URL `https://www.epa.gov/naaqs`.

A. Eydeland and K. Wolyniec. *Energy and power risk management: New developments in modeling, pricing, and hedging*, volume 97. John Wiley & Sons, 2002.

E. F. Fama and K. R. French. Common risk factors in the returns on stocks and bonds. *Journal of financial economics*, 33(1):3–56, 1993.

E. F. Fama and K. R. French. A five-factor asset pricing model. *Journal of financial economics*, 116(1):1–22, 2015.

E. F. Fama and J. D. MacBeth. Risk, return, and equilibrium: Empirical tests. *Journal of political economy*, 81(3):607–636, 1973.

K.-T. Fang, S. Kotz, and K. W. Ng. *Symmetric multivariate and related distributions*. Chapman and Hall/CRC, 2018.

M. J. Frank. On the simultaneous associativity off (x, y) andx+ y- f (x, y). *Aequationes mathematicae*, 19(1):194–226, 1979.

E. W. Frees and P. Wang. Credibility using copulas. *North American Actuarial Journal*, 9(2):31–48, 2005.

E. W. Frees, J. Carriere, and E. Valdez. Annuity valuation with dependent mortality. *Journal of risk and insurance*, pages 229–261, 1996.

F. Galton. I. co-relations and their measurement, chiefly from anthropometric data. *Proceedings of the Royal Society of London*, 45(273-279):135–145, 1889.

R. Garcia and G. Tsafack. Dependence structure and extreme comovements in international equity and bond markets. *Journal of Banking & Finance*, 35(8):1954–1970, 2011.

C. Genest and J. MacKay. The joy of copulas: Bivariate distributions with uniform marginals. *The American Statistician*, 40(4):280–283, 1986.

C. Genest and L.-P. Rivest. Statistical inference procedures for bivariate archimedean copulas. *Journal of the American statistical Association*, 88(423):1034–1043, 1993.

B. Goss, B. Yamey, and L. L. Johnson. *The theory of hedging and speculation in commodity futures*. Springer, 1976.

E. J. Gumbel. Bivariate exponential distributions. *Journal of the American Statistical Association*, 55(292):698–707, 1960.

N. J. Higham. Computing the nearest correlation matrix—a problem from finance. *IMA journal of Numerical Analysis*, 22(3):329–343, 2002.

N. L. Hjort and M. C. Jones. Locally parametric nonparametric density estimation. *The Annals of Statistics*, pages 1619–1647, 1996.

C. Jessen. Constant proportion portfolio insurance: Discrete-time trading and gap risk coverage. *Journal of Derivatives*, 21(3):36–53, 2014. ISSN 1074-1240. doi: 10.3905/jod.2014.21.3.036.

S. Jewson and A. Brix. *Weather derivative valuation: the meteorological, statistical, financial and mathematical foundations*. Cambridge University Press, 2005.

M. G. Kendall. A new measure of rank correlation. *Biometrika*, 30(1/2):81–93, 1938.

D. S. Kirschen and G. Strbac. *Fundamentals of Power System Economics*. John Wiley & Sons, 2018.

P. Kumar, M. Khare, R. Harrison, W. Bloss, A. Lewis, H. Coe, and L. Morawska. New directions: air pollution challenges for developing megacities like delhi. *Atmospheric Environment*, 122:657–661, 2015.

S. Kummer and C. Pauletto. The history of derivatives: A few milestones. In *EFTA Seminar on Regulation of Derivatives Markets*, pages 431–466, 2012.

H.-I. Lee, M.-H. Chiang, and H. Hsu. A new choice of dynamic asset management: the variable proportion portfolio insurance. *Applied Economics*, 40(16):2135–2146, 2008. URL `https://doi.org/10.1080/00036840600949280`.

H. E. Leland. Who should buy portfolio insurance? *The Journal of Finance*, 35(2): 581–594, 1980.

K. G. Lim and Q. Xiao. Computing maximum smoothness forward rate curves. *Statistics and computing*, 12(3):275–279, 2002. URL `https://doi.org/10.1023/A:1020707028156`.

J. Lintner. The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *The Review of Economics and Statistics*, 47(1):13–37, 1965. ISSN 00346535, 15309142. URL `http://www.jstor.org/stable/1924119`.

R. K. Y. Low, J. Alcock, R. Faff, and T. Brailsford. Canonical vine copulas in the context of modern portfolio management: Are they worth it? *Asymmetric Dependence in Finance: Diversification, Correlation and Portfolio Management in Market Downturns*, pages 263–289, 2018.

R. E. Lucas. Asset prices in an exchange economy. *Econometrica*, 46(6):1429–1445, 1978. ISSN 00129682, 14680262. URL http://www.jstor.org/stable/1913837.

Y. Malevergne and D. Sornette. Testing the gaussian copula hypothesis for financial assets dependences. *Quantitative finance*, 3(4):231, 2003.

H. Markowitz. Portfolio selection. *The Journal of Finance*, 7(1):77–91, 1952. ISSN 00221082, 15406261. URL http://www.jstor.org/stable/2975974.

A. J. McNeil, R. Frey, and P. Embrechts. *Quantitative risk management: concepts, techniques and tools-revised edition*. Princeton university press, 2015.

E. Mills. Insurance in a climate of change. *Science*, 309(5737):1040–1044, 2005.

J. Mossin. Equilibrium in a capital asset market. *Econometrica: Journal of the econometric society*, pages 768–783, 1966.

R. B. Nelsen. *An introduction to copulas*. Springer Science & Business Media, 2007.

Q. N. Nguyen, S. Aboura, J. Chevallier, L. Zhang, and B. Zhu. Local gaussian correlations in financial and commodity markets. *European journal of operational research*, 285(1):306–323, 2020.

F. Ollmar. An analysis of derivative prices in the nordic power market. 2003. URL http://hdl.handle.net/11250/164248.

H. Otneim. lg: An R package for Local Gaussian Approximations. *The R Journal*, 2021. doi: 10.32614/RJ-2021-079. URL https://journal.r-project.org/archive/2021/RJ-2021-079/index.html.

A. J. Patton. On the out-of-sample importance of skewness and asymmetric dependence for asset allocation. *Journal of financial econometrics*, 2(1):130–168, 2004.

K. Pearson. Vii. mathematical contributions to the theory of evolution.—iii. regression, heredity, and panmixia. *Philosophical Transactions of the Royal Society of London. Series A, containing papers of a mathematical or physical character*, (187):253–318, 1896.

A. Perold. Constant portfolio insurance. harvard business school. *Unpublished manuscript*, 1986.

E. Platen and J. West. A fair pricing approach to weather derivatives. *Asia-Pacific Financial Markets*, 11(1):23–53, 2004.

R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2021. URL `https://www.R-project.org/`.

W. Schachermayer and J. Teichmann. How close are the option pricing formulas of bachelier and black–merton–scholes? *Mathematical Finance: an international journal of mathematics, statistics and financial economics*, 18(1):155–170, 2008.

F. Schiller, G. Seidler, and M. Wimmer. Temperature models for pricing weather derivatives. *Quantitative Finance*, 12(3):489–500, 2012.

W. F. Sharpe. Capital asset prices: A theory of market equilibrium under conditions of risk. *The journal of finance*, 19(3):425–442, 1964.

M. Sklar. Fonctions de repartition an dimensions et leurs marges. *Publ. inst. statist. univ. Paris*, 8:229–231, 1959.

C. Spearman. The proof and measurement of association between two things. *The American Journal of Psychology*, 15(1):72–101, 1904. ISSN 00029556. URL `http://www.jstor.org/stable/1412159`.

B. Støve and D. Tjøstheim. Asymmetric dependence patterns in financial returns: An empirical investigation using local gaussian correlation. *Essays in Nonlinear Time Series Econometrics*, 307, 2014.

B. Støve, D. Tjøstheim, and K. O. Hufthammer. Using local gaussian correlation in a nonlinear re-examination of financial contagion. *Journal of Empirical Finance*, 25: 62–82, 2014.

E. J. Swan. *Building the global market: A 4000 year history of derivatives*. Kluwer Law International London, 2000.

D. Tasche. Expected shortfall and beyond. *Journal of Banking & Finance*, 26(7):1519–1533, 2002.

S. Thind. As temperatures tumble in north america, weather derivatives warm up. *Institutional Investor, January*, 23, 2014.

H. Till. Why haven't weather derivatives been more successful as futures contracts? a case study. *A Case Study (December 1, 2014)*, 2014.

D. Tjøstheim and K. O. Hufthammer. Local gaussian correlation: A new measure of dependence. *Journal of Econometrics*, 172(1):33–48, 2013.

D. Tjøstheim, H. Otneim, and B. Støve. Statistical dependence: Beynd pearson's ρ. *Statistical Science, to appear*, 2021.

D. Tjostheim, H. Otneim, and B. Stove. *Statistical Modeling Using Local Gaussian Approximation*. Academic Press, 2021.

J. L. Treynor. Market value, time, and risk. *Time, and Risk (August 8, 1961)*, 1961.

H. Wickham. Reshaping data with the reshape package. *Journal of Statistical Software*, 21(12):1–20, 2007. URL `http://www.jstatsoft.org/v21/i12/`.

H. Wickham. testthat: Get started with testing. *The R Journal*, 3:5–10, 2011. URL `https://journal.r-project.org/archive/2011-1/RJournal_2011-1_Wickham.pdf`.

H. Wickham. *ggplot2: Elegant Graphics for Data Analysis*. Springer-Verlag New York, 2016. ISBN 978-3-319-24277-4. URL `https://ggplot2.tidyverse.org`.

WRMA. *Weather Risk Derivative Survey*. Weather Risk Management Association, 2011.

D. M. Zimmer. The role of copulas in the housing crisis. *Review of Economics and Statistics*, 94(2):607–620, 2012.

Chapter 7

Papers

Paper A

7.1 Portfolio Allocation under Asymmetric Dependence in Asset Returns using Local Gaussian Correlations

Anders D. Sleire, Bård Støve, Håkon Otneim, Geir Drage Berentsen, Dag Tjøstheim, Sverre Hauso Haugen, (2021), Finance Research Letters, 102475.

Finance Research Letters 46 (2022) 102475

Contents lists available at ScienceDirect

Finance Research Letters

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

Portfolio allocation under asymmetric dependence in asset returns using local Gaussian correlations

Anders D. Sleire [b], Bård Støve [b,*], Håkon Otneim [a], Geir Drage Berentsen [a], Dag Tjøstheim [b], Sverre Hauso Haugen [b]

[a] Department of Business and Management Science, Norwegian School of Economics, Norway
[b] Department of Mathematics, University of Bergen, Norway

ARTICLE INFO

Keywords:
Asymmetric dependence
Local Gaussian correlation
Mean–variance
Portfolio allocation

ABSTRACT

It is well known that there are asymmetric dependence structures between financial returns. This paper describes a portfolio selection method rooted in the classical mean–variance framework that incorporates such asymmetric dependency structures using a nonparametric measure of local dependence, the local Gaussian correlation (LGC). It is shown that the portfolio optimization process for financial returns with asymmetric dependence structures is straightforward using local covariance matrices. The new method is shown to outperform the equally weighted ("1/N") portfolio and the classical Markowitz portfolio when applied to historical data on six assets.

1. Introduction

Modern portfolio theory aims to allocate assets by maximizing the expected return while minimizing risk. Markowitz (1952) provides the foundation for the mean–variance (MV) approach under the crucial assumption that the asset returns follow a joint-Gaussian distribution. The idea is simple; highly correlated assets should be avoided to obtain a diverse portfolio. Several empirical studies, however, document asymmetries in the distribution of financial returns. In particular, one often observes stronger dependence between assets when the market is going down. This phenomenon is known as *asymmetric dependence structures*, see e.g. Silvapulle and Granger (2001), Campbell et al. (2002), Okimoto (2008), Ang and Chen (2002), Hong et al. (2007), Chollete et al. (2009), Aas et al. (2009), Garcia and Tsafack (2011).

Asymmetric dependence may lead to less effective diversification of mean–variance balanced portfolios. Several studies seek to overcome this shortcoming by modeling the dependence structure using copula theory and then applying this modeling into the portfolio allocation problem, see e.g. Patton (2004), Hatherley and Alcock (2007), Low et al. (2013), Kakouris and Rustem (2014), Bekiros et al. (2015) and Han et al. (2017). These procedures are quite complicated, and a non-technical asset manager might be overwhelmed by such choices. Moreover, there is no guarantee that portfolio allocations based on complex models will improve performance compared with simpler methods, see e.g., DeMiguel et al. (2009) and Low et al. (2016), who show that outperforming the naive $1/N$ portfolio remains an elusive task.

Without making assumptions about the nature of the underlying probability model, we present a simple adjustment to the MV approach by replacing the correlation matrix of the assets with a *local* correlation matrix. This approach is based on the local Gaussian correlation (LGC), see Tjøstheim and Hufthammer (2013), and has been applied successfully to analyze dependence structures between asset returns, see e.g., Støve and Tjøstheim (2014), Støve et al. (2014), Bampinas and Panagiotidis (2017) and Nguyen

* Corresponding author.
E-mail address: Bard.Stove@math.uib.no (B. Støve).

https://doi.org/10.1016/j.frl.2021.102475
Received 3 July 2021; Received in revised form 15 September 2021; Accepted 19 September 2021
Available online 25 September 2021

et al. (2020), but has yet to be utilized in the portfolio allocation setting. The LGC provides a market-dependent adjustment to the correlation matrix that takes the current state of the market into account, and the main goal of this paper is to extend the classical MV framework by using the theory of the local Gaussian correlation, hence taking into account any asymmetric dependence structures between returns, and providing a simple alternative to the copula-based approaches. The organization of the paper is as follows. Section 2 presents the classical mean–variance portfolio approach and the extension using the local Gaussian correlation. In Section 3, we present a data set consisting of six asset returns, and in Section 4, we investigate the performance of portfolios constructed by our new approach and other methods. Finally, Section 5 offers some conclusions. In the supplementary material, we briefly review the local Gaussian correlation and provide some additional empirical study results.

2. Portfolio allocation using local Gaussian correlation

Denote the returns on N risky assets at time $t = 1, \dots, n$ by $\boldsymbol{R}_t \in \mathbb{R}^N$. Let $f_t(\boldsymbol{r}_t)$ denote the probability density function of $\boldsymbol{R}_t$, and let $\boldsymbol{\mu}_t \in \mathbb{R}^N$ and $\boldsymbol{\Sigma}_t \in \mathbb{R}^N \times \mathbb{R}^N$ denote its expectation vector and covariance matrix, respectively. Finally, let $\boldsymbol{w}_t = (w_{1,t}, \dots, w_{N,t})^T \in \mathbb{R}^N$ be the vector of portfolio weights at time t, to be determined by the portfolio selection strategy. We adopt the *full investment* constraint: $w_{1,t} + \cdots + w_{N,t} = 1$, for $t = 1, \dots, n$. Moreover, in the empirical example in Section 4, we investigate the performance of our proposed procedure with and without the *long-only* constraint ($0 \leq w_{i,t} \leq 1, i = 1, \dots, N, t = 1, \dots n$). We do not include a risk-free asset in our treatment of the portfolio allocation problem, but this does not impact our main findings.

The general portfolio optimization problem requires the investor to select weights $\boldsymbol{w}_t$ that maximizes an expected utility function at each time t. We consider the classical mean–variance as well as the minimum-variance utility functions, given respectively as

$$U_1 = \boldsymbol{w}_t^T \boldsymbol{\mu}_t - \frac{\gamma}{2} \boldsymbol{w}_t^T \boldsymbol{\Sigma}_t \boldsymbol{w}_t \qquad \text{and} \qquad U_2 = -\boldsymbol{w}_t^T \boldsymbol{\Sigma}_t \boldsymbol{w}_t, \tag{2.1}$$

where γ represents the investor's degree of risk aversion. Maximizing U_1 with respect to the portfolio weights provides a trade-off between expected volatility and expected returns for a given level of risk aversion (for simplicity fixed at $\gamma = 1$ throughout the paper). Maximizing U_2 results in the minimum variance portfolio.

We take as our point of departure the portfolio allocation approach as described by DeMiguel et al. (2009), Tu and Zhou (2011), and Low et al. (2016) when estimating the expected return vector $\boldsymbol{\mu}_t$ and covariance matrix $\boldsymbol{\Sigma}_t$ for monthly data. The approach is given as follows;

1. Select a sampling window of M trading months.
2. In each month $t > M$, estimate the expected return vector $\boldsymbol{\mu}_t$ and the covariance matrix $\boldsymbol{\Sigma}_t$ by their empirical counterparts, using the M preceding monthly returns.
3. Rebalance the portfolio on the first trading day of each month by solving the relevant optimization problem, i.e. optimizing one of the utility functions in (2.1).

The above algorithm implicitly assumes that the covariance matrix $\boldsymbol{\Sigma}_t$ completely describes the dependence structure among the assets under consideration. This property is not true in general unless the returns are jointly normally distributed, which is a strong assumption that is rarely satisfied in practice. Indeed, as mentioned in Section 1, there have been many attempts to replace the normality assumption in portfolio selection with more sophisticated distributions that better fit the return density f_t. However, this also results in a more complicated optimization routine than under the classical Markowitz framework indicated in the three steps listed above.

We propose to describe asymmetric dependence by making adjustments directly to the covariance matrix $\boldsymbol{\Sigma}_t$, which allows us to use the classical Markowitz formulation. To this end, we employ the *local Gaussian correlation*. The idea originated in Tjøstheim and Hufthammer (2013), who in turn based themselves on the local parameter concept of Hjort and Jones (1996). The latter authors approximate an unknown density function $f(\boldsymbol{r})$ by fitting a parametric family $f(\boldsymbol{r}, \theta)$ *locally* to $f(\boldsymbol{r})$, where $\theta \in \Theta$ is an unknown parameter in a parameter space Θ. This means that instead of constructing a single estimate $\widehat{\theta}$ of θ, they rather estimate a parameter *function*, $\widehat{\theta}(\boldsymbol{r})$, meaning that different members of the parametric family $\{f(\boldsymbol{r}, \theta), \theta \in \Theta\}$ approximate $f(\boldsymbol{r})$ in different parts of the domain of $f(\boldsymbol{r})$. Here, $\boldsymbol{r}$ represents a generic location in the domain of $\boldsymbol{R}_t$.

Hjort and Jones (1996) estimate $\theta(\boldsymbol{r})$ using a nonparametric local likelihood procedure, and Tjøstheim and Hufthammer (2013) consider the special case where $\{f(\boldsymbol{r}, \theta)\}$ is the multivariate Gaussian distribution, that is, $\theta = (\boldsymbol{\mu}, \boldsymbol{\Sigma})$. Under this specification, it is natural to interpret the local covariance matrix $\boldsymbol{\Sigma}(\boldsymbol{r})$ as a measure of local dependence, which in particular gives a natural description of the asymmetric dependence relationships so often observed in financial returns.

Consider, for example, the observed returns on two of the assets in our data set displayed in the left panel of Fig. 1. The classical Gaussian assumption results in a single estimated covariance matrix $\widehat{\boldsymbol{\Sigma}}$; while the local likelihood estimate $\widehat{\boldsymbol{\Sigma}}(\boldsymbol{r})$ is a function of $\boldsymbol{r}$. In the right-hand panel of Fig. 1 we see the corresponding local Gaussian correlation, which clearly indicates that these returns are most strongly dependent in the lower left part of the distribution.

In order to incorporate the asymmetry observed in the right-hand panel of Fig. 1, we propose to replace step 2 in the above procedure with the following:

2.' In each month $t > M$, estimate the expected return vector $\boldsymbol{\mu}_t$ by its empirical counterpart, and the local Gaussian covariance matrix $\boldsymbol{\Sigma}_t(\boldsymbol{r})$, using the M preceding monthly returns.

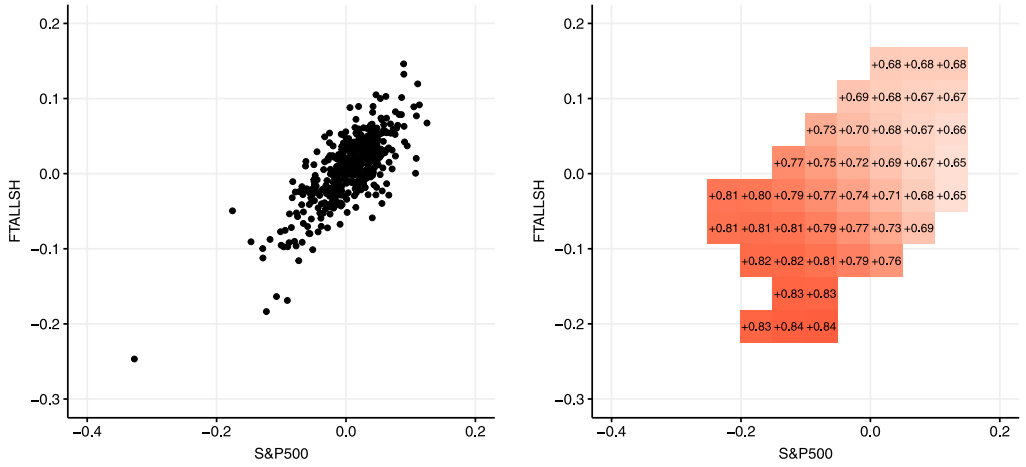

Fig. 1. Observations on two of the assets considered in the empirical analysis and the corresponding estimated local Gaussian correlation.

Table 1
Overview of the data series and abbreviations.

Name	Description
FTALLSH	FTSE Actuaries All Share Index
S&P500	Standard and Poor's 500 Index
BMUK10Y	UK Benchmark 10 Year DS government bond Index
BMUS10Y	US Benchmark 10 Year DS government bond Index
EWCI	Thomson Reuters Equal Weight Commodity Index
GSGCSPT	Standard and Poor's GSCI Gold Index

There are several approaches for choosing the evaluation point $\boldsymbol{r}$. A risk-averse investor can guard against large losses by selecting an evaluation point representing the asset returns during crisis periods. In this way, the corresponding estimated local covariance matrix reflects the (historical) dependence structure during crisis periods, provided that the window length M is chosen sufficiently large. However, the selection of the evaluation point can also be dynamic, i.e. $\boldsymbol{r} = \boldsymbol{r}_t$. For instance, the evaluation point may correspond to a subjective opinion of where the investor thinks the market is heading in the following trading month. The selection may also be based on more advanced statistical predictions.

In the empirical analysis in Section 4, we opt for a simple data-driven selection of evaluation points by computing the average of the last three months of observed returns. More specifically, the evaluation point at time t is defined for all pairs of assets i, j as

$$\boldsymbol{r}_t = \left(\frac{1}{3}\sum_{k=1}^{3} R^i_{t-k}, \frac{1}{3}\sum_{k=1}^{3} R^j_{t-k} \right). \tag{2.2}$$

This is a simple way of letting the covariance matrix dynamically adapt to the dependence structure of the market under the naïve assumption that the dependence structure between asset returns in month t is similar to the dependence structure of asset returns in the neighborhood of $\boldsymbol{r}_t$. As the empirical analysis in Section 4 will demonstrate, this simple selection of evaluation points performs well in practice.

3. Data

Our data set consists of monthly closing prices on six US dollar-denominated indices sourced from Refinitiv (past Thompson Reuters) Datastream. We calculate the returns as 100 times the difference in the log of the price indices. The sample period extends from February 1980 to August 2018, yielding 463 monthly return observations of the following assets: FTSE Actuaries All Share Index, (FTALLSH), Standard and Poor's 500 Index (S&P500), UK Benchmark 10 Year DS government bond Index (BMUK10Y), US Benchmark 10 Year DS government bond Index (BMUS10Y), Thomson Reuters Equal Weight Commodity Index (EWCI), and Standard and Poor's GSCI Gold Index (GSGCSPT) (see Table 1).

From the descriptive statistics in Table 2, we note that all of the returns are skewed and show relatively high kurtosis. Normality is rejected with the Jarque–Bera test on the 1% level for all series. This suggests that the multivariate normal distribution with a global covariance matrix does not provide a sufficient description of the dependence structure, particularly in the tails of the distribution.

Table 2
Correlations and descriptive statistics for the asset returns studied in the empirical analysis.

	FTALLSH	S&P500	BMUK10Y	BMUS10Y	EWCI	GSGCSPT
Global correlation matrix						
FTALLSH	1					
S&P500	0.760	1				
BMUK10Y	0.184	0.017	1			
BMUS10Y	−0.067	−0.029	0.489	1		
EWCI	0.246	0.288	−0.094	−0.185	1	
GSGCSPT	0.038	0.031	0.080	0.077	0.483	1
Local correlation matrix, bear market (lower 5% percentiles)						
FTALLSH	1					
S&P500	0.843	1				
BMUK10Y	0.174	−0.017	1			
BMUS10Y	0.020	0.034	0.635	1		
EWCI	0.161	0.185	−0.140	−0.224	1	
GSGCSPT	−0.135	−0.131	0.204	0.215	0.480	1
Descriptive statistics						
Observations	463	463	463	463	463	463
Mean	0.628	0.704	0.769	0.583	0.079	0.177
Std. Dev.	4.588	4.406	2.376	2.417	3.511	5.211
Variance	21.050	19.413	5.643	5.839	12.326	27.159
Skewness	−1.300	−0.968	−0.128	0.453	−0.592	0.026
Kurtosis	6.288	3.665	1.325	1.960	3.775	3.036
Jarque–Bera	903.903	335.969	36.135	91.622	306.377	180.971
Sharpe ratio	0.137	0.160	0.324	0.241	0.023	0.034
Max. drawdown	49.887	59.811	15.764	12.035	48.397	73.680
Min	−32.711	−24.677	−7.824	−7.600	−20.050	−21.887
1 Quartile	−1.474	−1.694	−0.585	−0.922	−1.794	−2.668
Median	1.176	1.242	0.843	0.497	0.151	−0.161
3 Quartile	3.559	3.265	2.151	1.853	1.998	2.899
Max	12.523	14.612	8.851	12.660	13.384	26.336

The two top panels in Table 2 show the global and local correlation matrices over the entire sampling period. The latter is constructed for a *bear market* scenario using the lower 5% percentiles for the evaluation point selection in the pairwise calculation approach described in the supplementary material. The strongest positive and negative correlation is observed between the stock indices FTALLSH and S&P500 ($\hat{\rho} = 0.76$), and between EWCI and BMUS10Y ($\hat{\rho} = -0.185$), respectively. The corresponding LGCs in the bear market scenario are $\hat{\rho} = 0.843$ and $\hat{\rho} = -0.224$, respectively, indicating the ability of the LGC to capture asymmetric dependence structures, see also Tjøstheim and Hufthammer (2013).

4. Empirical results

Our analysis[1] compares the portfolio selection strategies listed in Table 3 by evaluating their performance using both terminal wealth as well as a range of risk-adjusted performance measures. Following Low et al. (2016), we use the naïve $1/N$ weighted portfolio strategy as a benchmark model in the analysis. This strategy distributes weights equally across the portfolio at the start of the sampling period and is left unadjusted for the rest of the investment horizon. Tu and Zhou (2011) find that longer sampling windows result in improved portfolio strategy performance; hence we use both $M = 120$ and $M = 240$ month sampling windows. We report the $M = 240$ in the following sections. The corresponding results using $M = 120$ months are given in the supplementary material.

Inspired by Low et al. (2013), we proceed to evaluate the portfolio rebalancing, terminal wealth as well as the risk-adjusted performance for each of the strategies. A descriptive analysis of out-of-sample results is available in the supplementary material.

4.1. Portfolio rebalancing and terminal wealth

Table 4 provides a summary of the portfolio rebalancing analysis and the terminal wealth reached by each of the strategies. The average standard deviation within target portfolio weights across the entire out-of-sample time period is calculated as follows:

$$\bar{\sigma}_k = \frac{\sum_{t=1}^{n-M} \sigma_{t,k}}{n-M},$$

where

$$\sigma_{t,k} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(\hat{w}_{i,t,k} - \bar{w}_{\cdot,t,k})^2},$$

[1] Reproduce results or perform new studies with: https://gitlab.com/sleire/lgportf

Table 3
Portfolio strategies applied in the empirical analysis. All strategies allowing short sales have a lower limit on portfolio weights equal to −0.5.

Strategy	Description
Benchmark strategy	
EW	$1/N$ (equal weight portfolio) without rebalancing
Global approach	
MVS	Mean–variance portfolio with short sales
MVSC	Mean–variance portfolio with short sales constraint
MIN	Minimum variance portfolio
MINC	Minimum variance portfolio with short sales constraint
Local approach	
MVS-L	Mean–variance portfolio with short sales using local covariance matrices
MVSC-L	Mean–variance portfolio with short sales constraint using local covariance matrices
MIN-L	Minimum variance portfolio using local covariance matrices
MINC-L	Minimum variance portfolio with short sales constraint using local covariance matrices

Table 4
Portfolio rebalancing analysis and terminal wealth based on an initial investment of $1 for the different strategies (cf. Table 3) considered. Window size $M = 240$ months.

	$\bar{\sigma}_k$	Max. adj.	Min. adj.	Avg.turnover	Wealth	Wealth incl.tcost
Benchmark strategy						
EW	0	0	0	0	2.231	2.231
Global approach						
MVS	16.663	9.228	−6.741	6.325	2.605	2.551
MVSC	15.813	9.228	−7.426	5.622	2.571	2.524
MIN	16.249	5.036	−4.434	3.090	2.729	2.701
MINC	15.952	5.039	−4.428	2.793	2.742	2.717
Local approach						
MVS-L	15.660	73.184	−89.316	16.967	2.855	2.698
MVSC-L	14.253	18.096	−26.245	12.105	2.745	2.637
MIN-L	17.014	53.636	−79.931	21.617	2.953	2.749
MINC-L	15.243	19.646	−26.028	13.762	2.910	2.780

The Max.adj. and Min.adj. is the maximum values for positive and negative weight adjustments, respectively. Avg.turnover is defined in Eq. (4.1). Basis points of 15 per transaction are imposed as costs.

and where $\hat{w}_{i,t,k}$ is the portfolio weight for asset i at time t using portfolio strategy k, and $\bar{w}_{\cdot,t,k}$ is the average weight across the N assets in portfolio k at time t. The maximum values for positive and negative weight adjustments are the largest positive and negative weight changes on the asset level. Following DeMiguel et al. (2009), we also report the average turnover, which is calculated as

$$\text{Average turnover} = \frac{1}{n-M}\sum_{t=1}^{n-M}\sum_{i=1}^{N}(|\hat{w}_{i,t+1,k} - \hat{w}_{i,t,k}|), \tag{4.1}$$

We compute terminal wealth with and without a transaction cost of 15 basis points; such cost is comparable to prior studies, see e.g Low et al. (2016).

The variability of portfolio weights reported in Table 4 shows no systematic differences between the local and global approaches. Looking at the maximum and minimum adjustments of portfolio weights, however, there are clear differences. The local Gaussian strategies require adjustments of larger magnitude in both directions. This is particularly the case for the unconstrained models allowing short sales. Viewed across all strategies, we see larger and more frequent adjustments in the local portfolio strategies.

We see that increased trading volume translates into lower terminal wealth when transaction costs are included in the analysis. However, all local Gaussian strategies achieve higher terminal wealth than their traditional counterparts, also when transaction costs are included. The top-ranked strategy exclusive costs is MIN-L. When costs are included, the long-only portfolio MINC-L achieves the best result.

Fig. 2 shows wealth accumulation and drawdowns for the hypothetical investment of $1 in each of the nine strategies included in the analysis. As seen in the upper part of the figure, the local Gaussian MIN-L produces the largest final wealth when disregarding trade costs. It remains top-ranked during most months in the sample and suffers from smaller drawdowns in volatile periods such as the 2008 Financial Crisis. During this period, the EW strategy loses out substantially, which partially explains the overall poor performance of this strategy relative to the rest. The other local strategies also performs better than the corresponding global ones during this period. This is as expected, as we typically observe higher local dependence between asset returns during crisis periods (see e.g. Støve and Tjøstheim (2014)). When transaction costs are included, the strategy MIN-L still performs well but is surpassed by the constrained MINC-L, which has a lower turnover.

In Fig. 3, we see an illustration of the difference between the traditional Markowitz minimum variance, short sale constrained portfolio (MINC), and the proposed local counterpart (MINC-L). In the top panel, we see the estimated global and local variances of one of the assets in our example (EWCI) and the corresponding weight as a function of time. The local estimates are naturally

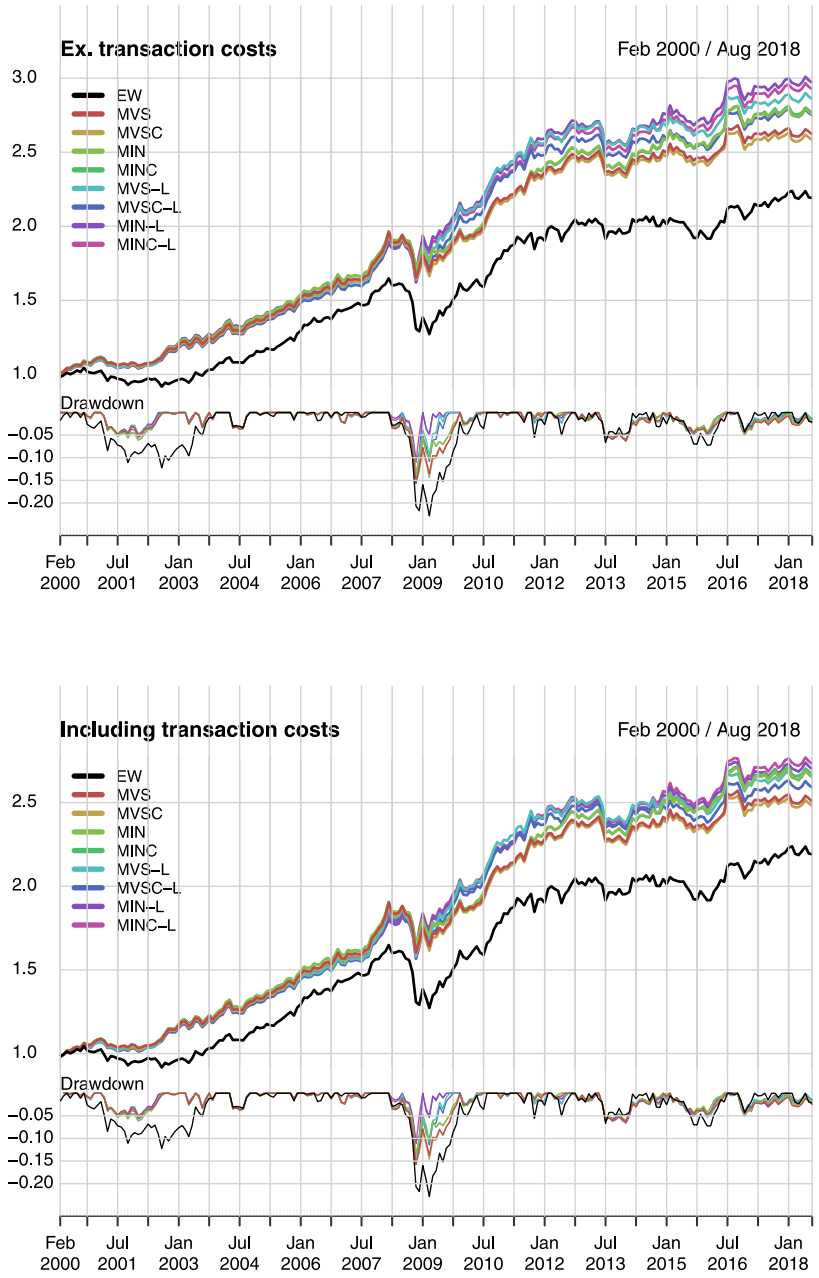

Fig. 2. Wealth accumulation of the different strategies based on an initial investment of $ 1, using a rolling window of size $M = 240$ monthly observations, top plot excluding transaction costs, bottom plot including transaction cost of 15 basis points per transaction.

more volatile due to their nonparametric estimation. However, they are also more sensitive to the state of the market, which is most easily visible during the financial crisis of 2008. The local variance of the asset increases sharply in this period, which is immediately reflected in the portfolio weight that quickly decreases to zero. We see similar effects for other instances of increased

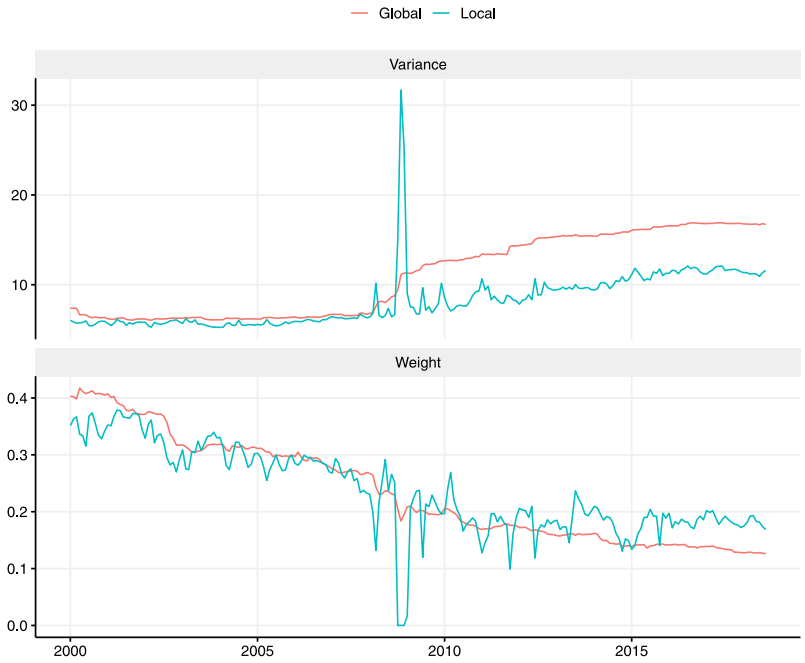

Fig. 3. Top panel: The estimated variance for one of the assets (EWCI) in our sample for the minimum variance and short sales constrained portfolio with traditional (global) Markowitz optimization (MINC), as well as the corresponding local version (MINC-L). In the bottom panel, we see the corresponding weight for this asset under the two strategies. See the online appendix for a corresponding plot covering all the assets in this example. The sample window is $M = 240$ months.

local variance. Note also that under the traditional (global) estimate, the estimated global variance reflects the financial crisis in the entire remaining sample period, leading to smaller investments in this asset than for the locally estimated portfolio.

4.2. Evaluation of risk-adjusted performance

Table 5 reports the out-of-sample performance of the different strategies considered using the following risk-adjusted metrics; The Sharpe ratio (Sharpe, 1966), the two modified Sharpe ratios VaR Sharpe, and ES Sharpe, where the Value at Risk and Expected Shortfall are used as risk measures Gregoriou and Gueyie (2003) and Favre and Galeano (2002). We also consider the Certainty Equivalent (CEQ), the Sortino ratio (Sortino and Price, 1994) and finally, the Omega ratio (Keating and Shadwick, 2002). All metrics produce high values for the best-performing strategies. Furthermore, we have performed the z-test of Ledoit and Wolf (2008), which is applied to the Sharpe Ratios to indicate the statistical differences for all MV optimizations against the $1/N$ benchmark.

Results excluding and including transaction costs are reported in Panel A and B, respectively. In both cases, the local portfolio strategies systematically outperform their traditional counterparts. Furthermore, we note that the local minimum variance portfolio has the highest performance across all metrics. The Sharpe and the annualized Sharp ratios prefer the long-only version (MINC-L), while the Var Sharpe, ES Sharpe, Sortino, and Omega ratios prefer the unconstrained version (MIN-L). The CEQ prefers the constrained version when costs are excluded and the unconstrained version when costs are included.

5. Concluding remarks

The results in this paper suggest that challenges related to return asymmetries may be handled in a familiar and well-established framework for portfolio management by replacing the global covariance matrix with a local version. Improved performance and simplicity are some of the appeals with the local Gaussian approach to portfolio management, even when considering higher transaction costs due to increased rebalancing requirements. There are, however, matters to keep in mind when implementing the approach. The selection of evaluation points for calculating the pairwise local correlations will affect the local Gaussian covariance matrix. We have evaluated alternative approaches to the moving evaluation point selection without observing substantial changes in results and conclusions. Nevertheless, there is a variety of options and possibilities for this choice. A more thorough analysis of these effects is left for future studies.

Table 5
Out-of-sample performance for the different portfolio strategies (cf. Table 3) considered. The sample window is $M = 240$ months.

	Sharpe	VaR Sharpe	ES Sharpe	Ann. Sharpe	CEQ	Sortino	Omega
Panel A: Ex. transaction costs							
Benchmark strategy							
EW	0.174	0.110	0.057	0.577	0.353	0.263	1.593
Global approach							
MVS	0.267	0.179	0.099	0.919	0.433	0.427	2.025
MVSC	0.264	0.175	0.097	0.905	0.426	0.417	2.009
MIN	0.289*	0.198	0.116	0.999	0.455	0.474	2.120
MINC	0.290*	0.199	0.117	1.000	0.457	0.474	2.120
Local approach							
MVS-L	0.276*	0.203	0.115	0.951	0.472	0.472	2.126
MVSC-L	0.270*	0.193	0.107	0.929	0.454	0.453	2.075
MIN-L	0.301*	**0.317**	**0.317**	1.041	**0.489**	**0.567**	**2.303**
MINC-L	**0.309***	0.240	0.157	**1.070**	0.482	0.554	2.268
Panel B: Incl. transaction costs							
Benchmark strategy							
EW	0.179	0.113	0.058	0.593	0.362	0.270	1.613
Global approach							
MVS	0.261	0.174	0.097	0.895	0.423	0.415	1.989
MVSC	0.258	0.171	0.095	0.885	0.418	0.407	1.978
MIN	0.285*	0.194	0.114	0.982	0.449	0.465	2.095
MINC	0.285*	0.195	0.115	0.985	0.451	0.466	2.097
Local approach							
MVS-L	0.262	0.188	0.106	0.900	0.448	0.440	2.041
MVSC-L	0.260	0.183	0.101	0.890	0.438	0.430	2.014
MIN-L	0.281	**0.279**	**0.279**	0.969	**0.456**	0.518	2.174
MINC-L	**0.295***	0.225	0.148	**1.019**	0.461	**0.522**	**2.182**

The maximum values for the risk-adjusted performance metrics are in **bold**. 15 basis points per transaction is imposed as costs in Panel B. '*' indicates that the Sharpe ratio is statistically different on the 5 percent level from the benchmark (EW) strategy using the z-test of Ledoit and Wolf (2008).

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work has been partly supported by the Finance Market Fund (Norway). We thank two reviewers and the Editor for suggestions leading to an improved paper.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.frl.2021.102475.

References

Aas, K., Czado, C., Frigessi, A., Bakken, H., 2009. Pair-copula constructions of multiple dependence. Insurance Math. Econom. 44, 182–198.
Ang, A., Chen, J., 2002. Asymmetric correlations of equity portfolios. J. Financ. Econ. 63 (3), 443–494.
Bampinas, G., Panagiotidis, T., 2017. Oil and stock markets before and after financial crises: A local Gaussian correlation approach. J. Futures Mark. 37 (12), 1179–1204.
Bekiros, S., Hernandez, J.A., Hammoudeh, S., Nguyen, D.K., 2015. Multivariate dependence risk and portfolio optimization: An application to mining stock portfolios. Resour. Policy 46, 1–11.
Campbell, R., Koedijk, K., Kofman, P., 2002. Increased correlation in bear markets. Financ. Anal. J. 58 (1), 87–94.
Chollete, L., Heinen, A., Valdesogo, A., 2009. Modeling international financial returns with a multivariate regime switching copula. J. Financ. Econom. 7, 437–480.
DeMiguel, V., Garlappi, L., Uppal, R., 2009. Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy? Rev. Financ. Stud. 22 (5), 1915–1953.
Favre, L., Galeano, J.-A., 2002. Mean-modified value-at-risk optimization with hedge funds. J. Altern. Invest. 5 (2), 21–25.
Garcia, R., Tsafack, G., 2011. Dependence structure and extreme comovements in international equity and bond markets. J. Bank. Financ. 35, 1954–1970.
Gregoriou, G., Gueyie, J.-P., 2003. Risk-adjusted performance of funds of hedge funds using a modified sharpe ratio. J. Wealth Manag. 6 (3), 77–83.
Han, Y., Li, P., Xia, Y., 2017. Dynamic robust portfolio selection with copulas. Finance Res. Lett. 21, 190–200.
Hatherley, A., Alcock, J., 2007. Portfolio construction incorporating asymmetric dependence structures: a user's guide. Account. Finance 47 (3), 447–472.
Hjort, N.L., Jones, M.C., 1996. Locally parametric nonparametric density estimation. Ann. Statist. (ISSN: 0090-5364) 24 (4), 1619–1647.
Hong, Y., Tu, J., Zhou, G., 2007. Asymmetries in stock returns: Statistical tests and economic evaluation. Rev. Financ. Stud. 20 (5), 1547–1581.

Kakouris, I., Rustem, B., 2014. Robust portfolio optimization with copulas. European J. Oper. Res. 235 (1), 28–37.
Keating, C., Shadwick, W.F., 2002. A universal performance measure. J. Perform. Meas. 6 (3), 59–84.
Ledoit, O., Wolf, M., 2008. Robust performance hypothesis testing with the sharpe ratio. J. Empir. Financ. 15 (5), 850–859.
Low, R.K.Y., Alcock, J., Faff, R., Brailsford, T., 2013. Canonical vine copulas in the context of modern portfolio management: Are they worth it? J. Bank. Financ. 37 (8), 3085–3099.
Low, R.K.Y., Faff, R., Aas, K., 2016. Enhancing mean–variance portfolio selection by modeling distributional asymmetries. J. Econ. Bus. 85, 49–72.
Markowitz, H., 1952. Portfolio selection. J. Finance 7 (1), 77–91.
Nguyen, Q.N., Aboura, S., Chevallier, J., Lyuyuan, Z., Zhu, B., 2020. Local Gaussian correlations in financial and commodity markets. European J. Oper. Res..
Okimoto, T., 2008. New evidence of asymmetric dependence structures in international equity markets. J. Financ. Quant. Anal. 43 (3), 787–816.
Patton, A.J., 2004. On the out-of-sample importance of skewness and asymmetric dependence for asset allocation. J. Financ. Econom. 2 (1), 130–168.
Sharpe, W.F., 1966. Mutual fund performance. J. Bus. 39 (1), 119–138.
Silvapulle, P., Granger, C.W.J., 2001. Large returns, conditional correlation and portfolio diversification: a value-at-risk approach. Quant. Finance 1 (5), 542–551.
Sortino, F.A., Price, L.N., 1994. Performance measurement in a downside risk framework. J. Invest. 3 (3), 59–64.
Støve, B., Tjøstheim, D., 2014. Measuring asymmetries in financial returns: An empirical investigation using local Gaussian correlation. In: N. Haldrup, M.M., Saikkonen, P. (Eds.), Essays in Nonlinear Time Series Econometrics. Oxford University Press, Oxford, pp. 307–329.
Støve, B., Tjøstheim, D., Hufthammer, K.O., 2014. Using local Gaussian correlation in a nonlinear re-examination of financial contagion. J. Empir. Financ. 25, 785–801.
Tjøstheim, D., Hufthammer, K.O., 2013. Local Gaussian correlation: A new measure of dependence. J. Econometrics 172, 33–48.
Tu, J., Zhou, G., 2011. Markowitz meets talmud: A combination of sophisticated and naive diversification strategies. J. Financ. Econ. 99 (1), 204–215.

Supplementary material for Portfolio Allocation under Asymmetric Dependence in Asset Returns using Local Gaussian Correlations

Anders D. Sleire[2], Bård Støve[2]*, Håkon Otneim[1], Geir Drage Berentsen[1], Dag Tjøstheim[2], Sverre Hauso Haugen[2],

[1]Department of Business and Management Science, Norwegian School of Economics
[2] Department of Mathematics, University of Bergen, Norway

September 14, 2021

1 Local Gaussian correlation

This paper relies on the recent developed dependence measure, the local Gaussian correlation (LGC), introduced by Tjøstheim and Hufthammer [2013]. This is a local characterization of dependence, and this idea has also been extended to several different situations, as test of independence, see Berentsen and Tjøstheim [2014], Lacal and Tjøstheim [2017] and Lacal and Tjøstheim [2019], density and conditional density estimation, see Otneim and Tjøstheim [2017] and Otneim and Tjøstheim [2018], a local Gaussian partial correlation, Otneim and Tjøstheim [2021], and local Gaussian spectral estimation, see Jordanger and Tjøstheim [2020]. Finally, the relationship between the local Gaussian correlation and different copulas has been studied in Berentsen et al. [2014]. For completeness, we briefly present the local Gaussian correlation in a standard way, and we note that this section closely follows the presentation of the LGC in Tjøstheim et al. [2021].

Let $\boldsymbol{R} = (R_1, R_2)$ represent the return on two risky assets with density $f(\boldsymbol{r}) = f(r_1, r_2)$. For simplicity we drop the time index in the following. We approximate f locally in each point $\boldsymbol{r} = (r_1, r_2)$ by a Gaussian bivariate density, $\psi_{r_1,r_2}(\boldsymbol{v})$, where $\boldsymbol{v} = (v_1, v_2)$ are running variables. Let $\boldsymbol{\mu}(\boldsymbol{r}) = (\mu_1(\boldsymbol{r}), \mu_2(\boldsymbol{r}))$ be the mean vector in the normal distribution having density $\psi_{\boldsymbol{r}}$, $\boldsymbol{\sigma}(\boldsymbol{r}) = (\sigma_1(\boldsymbol{r}), \sigma_2(\boldsymbol{r}))$ is the vector of standard deviations, and $\rho(\boldsymbol{r})$ is the correlation coefficient

*Corresponding author. Tel. +47 55 58 28 86. E-mail: Bard.Stove@math.uib.no.

in the normal distribution $\psi_{\boldsymbol{r}}$. The approximating density is then given as

$$\psi_{\boldsymbol{r}} = \psi(\boldsymbol{v}, \mu_1(\boldsymbol{r}), \mu_2(\boldsymbol{r}), \sigma_1^2(\boldsymbol{r}), \sigma_2^2(\boldsymbol{r}), \rho(\boldsymbol{r})) = \frac{1}{2\pi\sigma_1(\boldsymbol{r})\sigma_2(\boldsymbol{r})\sqrt{1-\rho^2(\boldsymbol{r})}}$$
$$\times \exp\Big[-\frac{1}{2}\frac{1}{1-\rho^2(\boldsymbol{r})}\Big(\frac{(v_1-\mu_1(\boldsymbol{r}))^2}{\sigma_1^2(\boldsymbol{r})} - 2\rho(\boldsymbol{r})\frac{(v_1-\mu_1(\boldsymbol{r}))(v_2-\mu_2(\boldsymbol{r}))}{\sigma_1(\boldsymbol{r})\sigma_2(\boldsymbol{r})}$$
$$+ \frac{(v_2-\mu_2(\boldsymbol{r}))^2}{\sigma_2^2(\boldsymbol{r})}\Big)\Big]. \tag{1.1}$$

Moving to another point $\boldsymbol{r}' = (r_1', r_2')$ gives another approximating normal distribution $\psi_{\boldsymbol{r}'}$ depending on a new set of parameters $(\mu_1(\boldsymbol{r}'), \mu_2(\boldsymbol{r}'), \sigma_1(\boldsymbol{r}'), \sigma_2(\boldsymbol{r}'), \rho(\boldsymbol{r}'))$. One exception to this is the case where f itself is Gaussian with parameters $(\mu_1, \mu_2, \sigma_1, \sigma_2, \rho)$, in which case $(\mu_1(\boldsymbol{r}), \mu_2(\boldsymbol{r}), \sigma_1(\boldsymbol{r}), \sigma_2(\boldsymbol{r}), \rho(\boldsymbol{r})) \equiv (\mu_1, \mu_2, \sigma_1, \sigma_2, \rho)$.

The population parameter vector, $\boldsymbol{\theta}(\boldsymbol{r}) \stackrel{\text{def}}{=} (\mu_1(\boldsymbol{r}), \mu_2(\boldsymbol{r}), \sigma_1(\boldsymbol{r}), \sigma_2(\boldsymbol{r}), \rho(\boldsymbol{r}))$, are obtained by minimizing the local penalty function measuring the difference between f and $\psi_{\boldsymbol{r}}$. It is defined by

$$q = \int K_{\boldsymbol{b}}(\boldsymbol{v}-\boldsymbol{r})[\psi(\boldsymbol{v}, \boldsymbol{\theta}(\boldsymbol{r})) - \ln\{\psi(\boldsymbol{v}, \boldsymbol{\theta}(\boldsymbol{r}))\}f(\boldsymbol{v})]\mathrm{d}\boldsymbol{v} \tag{1.2}$$

where $K_{\boldsymbol{b}}(\boldsymbol{v}-\boldsymbol{x}) = (b_1 b_2)^{-1} K_1(b_1^{-1}(v_1-r_1))K_2(b_2^{-1}(v_2-r_2))$ is a product kernel with bandwidths $\boldsymbol{b} = (b_1, b_2)$. As is seen in Hjort and Jones (1996, pp 1623-1624), the expression in (1.2) can be interpreted as a locally weighted Kullback-Leibler distance from f to $\psi(\cdot, \boldsymbol{\theta}(\boldsymbol{r}))$. Hence, the minimizer $\boldsymbol{\theta}_{\boldsymbol{b}}(\boldsymbol{r})$ (also depending on K) should satisfy

$$\int K_{\boldsymbol{b}}(\boldsymbol{v}-\boldsymbol{r})\frac{\partial}{\partial\theta_j}[\ln\{\psi(\boldsymbol{v}, \boldsymbol{\theta}(\boldsymbol{r}))\}f(\boldsymbol{v}) - \psi(\boldsymbol{v}, \boldsymbol{\theta}(\boldsymbol{r}))]\mathrm{d}\boldsymbol{v} = 0, \quad j = 1, \ldots, 5. \tag{1.3}$$

In the first step we define the population value $\boldsymbol{\theta}_{\boldsymbol{b}}(\boldsymbol{r})$ as the minimizer of (1.2), assuming that there is a unique solution to (1.3). The definition of $\boldsymbol{\theta}_{\boldsymbol{b}}(\boldsymbol{r})$ and the assumption of uniqueness are essentially identical to those used in Hjort and Jones [1996] for more general parametric families of densities.

In the next step we let $\boldsymbol{b} \to \boldsymbol{0}$ and consider the limiting value $\boldsymbol{\theta}(\boldsymbol{r}) = \lim_{\boldsymbol{b}\to\boldsymbol{0}} \boldsymbol{\theta}_{\boldsymbol{b}}(\boldsymbol{r})$. This is in fact considered indirectly by Hjort and Jones [1996] and more directly in Tjøstheim and Hufthammer [2013], both using Taylor expansion arguments. In the following we assume that a limiting value $\boldsymbol{\theta}(\boldsymbol{r})$ independent of $\boldsymbol{b}$ and K exists.

When estimating $\boldsymbol{\theta}(\boldsymbol{r})$ and $\boldsymbol{\theta}_{\boldsymbol{b}}(\boldsymbol{r})$ we have to use a neighborhood with a finite bandwidth, which is in analogy to nonparametric density estimation. The estimate $\widehat{\boldsymbol{\theta}}(\boldsymbol{r}) = \widehat{\boldsymbol{\theta}}_{\boldsymbol{b}}(\boldsymbol{r})$ is then obtained from maximizing a local likelihood. Given observations $\boldsymbol{R}_1, \ldots, \boldsymbol{R}_n$ the local log

likelihood is determined by

$$L(\boldsymbol{R}_1, \ldots, \boldsymbol{R}_n, \boldsymbol{\theta}(\boldsymbol{r})) = n^{-1} \sum_i K_{\boldsymbol{b}}(\boldsymbol{R}_i - \boldsymbol{r}) \log \psi(\boldsymbol{R}_i, \boldsymbol{\theta}(\boldsymbol{r})) \\ - \int K_b(\boldsymbol{v} - \boldsymbol{r}) \psi(\boldsymbol{v}, \boldsymbol{\theta}(\boldsymbol{r})) \mathrm{d}\boldsymbol{v}. \tag{1.4}$$

When $\boldsymbol{b} \to \infty$, the last term has 1 as its limiting value, and the likelihood reduces to the ordinary global likelihood. This last term is essential, as it implies that $\psi(\boldsymbol{r}, \boldsymbol{\theta}_{\boldsymbol{b}}(\boldsymbol{r}))$ is not allowed to stray far away from $f(\boldsymbol{r})$ as $\boldsymbol{b} \to \boldsymbol{0}$. Indeed, using the notation

$$u_j(\cdot, \boldsymbol{\theta}) \stackrel{\text{def}}{=} \frac{\partial}{\partial \theta_j} \log \psi(\cdot, \boldsymbol{\theta}), \tag{1.5}$$

by the law of large numbers, or by the ergodic theorem in the time series case, assuming $\mathbb{E}(K_{\boldsymbol{b}}(\boldsymbol{R}_i - \boldsymbol{r}) \log \psi(\boldsymbol{R}_i, \boldsymbol{\theta}_{\boldsymbol{b}}(\boldsymbol{r}))) < \infty$, we have almost surely

$$\begin{aligned} \frac{\partial L}{\partial \theta_j} &= n^{-1} \sum_i K_{\boldsymbol{b}}(\boldsymbol{R}_i - \boldsymbol{r}) u_j(\boldsymbol{R}_i, \boldsymbol{\theta}_{\boldsymbol{b}}(\boldsymbol{r})) \\ &\quad - \int K_{\boldsymbol{b}}(\boldsymbol{v} - \boldsymbol{r}) u_j(\boldsymbol{v}, \boldsymbol{\theta}_{\boldsymbol{b}}(\boldsymbol{r})) \psi(\boldsymbol{v}, \boldsymbol{\theta}_{\boldsymbol{b}}(\boldsymbol{r})) \mathrm{d}\boldsymbol{v} \\ &\to \int K_{\boldsymbol{b}}(\boldsymbol{v} - \boldsymbol{r}) u_j(\boldsymbol{v}, \boldsymbol{\theta}_{\boldsymbol{b}}(\boldsymbol{r})) [f(\boldsymbol{v}) - \psi(\boldsymbol{v}, \boldsymbol{\theta}_{\boldsymbol{b}}(\boldsymbol{r}))] \mathrm{d}\boldsymbol{v}. \end{aligned} \tag{1.6}$$

Setting the expression in the first line of (1.6) equal to zero yields the local maximum likelihood estimate $\widehat{\boldsymbol{\theta}}_{\boldsymbol{b}}(\boldsymbol{r})$ $(= \widehat{\boldsymbol{\theta}}(\boldsymbol{r}))$ of the population value $\boldsymbol{\theta}_{\boldsymbol{b}}(\boldsymbol{r})$ (and $\boldsymbol{\theta}(\boldsymbol{r})$ which satisfies (1.3)).

An asymptotic theory has been developed in Tjøstheim and Hufthammer [2013] for $\widehat{\boldsymbol{\theta}}_{\boldsymbol{b}}(\boldsymbol{r})$ for the case that $\boldsymbol{b}$ is fixed and for $\widehat{\boldsymbol{\theta}}(\boldsymbol{r})$ in the case that $\boldsymbol{b} \to \boldsymbol{0}$. The first case is much easier to treat than the second one. In fact for the first case the theory of Hjort and Jones [1996] can be used almost directly, although it is extended to the ergodic time series case in Tjøstheim and Hufthammer [2013]. In the case that $\boldsymbol{b} \to \boldsymbol{0}$, this leads to a slow convergence rate of $(n(b_1 b_2)^3)^{-1/2}$, which is the same convergence rate as for the the estimated dependence function treated in Jones [1996].

We have thus far concentrated on the bivariate case, in which we estimate a single local Gaussian correlation based on a bivariate sample. In principle, it is straightforward to extend to the case of more than two variables. Assume that we observe a multivariate sample $\boldsymbol{R}_i = \{R_{1i}, \ldots, R_{pi}\}$, $i = 1, \ldots, n$ with dimension $p > 2$. We can then estimate the $p \times p$ local correlation matrix $\boldsymbol{\rho}(\boldsymbol{r}) = \{\rho_{k\ell}(\boldsymbol{r})\}$, $1 \le k < \ell \le p$, $\boldsymbol{r} = (r_1, \ldots, r_p)$, as well as the p local means and local variances $\boldsymbol{\mu}(\boldsymbol{r}) = \{\mu_1(\boldsymbol{r}), \ldots, \mu_p(\boldsymbol{r})\}$ and $\boldsymbol{\sigma}(\boldsymbol{r}) = \{\sigma_1(\boldsymbol{r}), \ldots, \sigma_p(\boldsymbol{r})\}$ by maximizing the local likelihood function (1.4). The precision of such estimates, however, deteriorates quickly as the dimension p grows, due to the curse of dimensionality.

However, a simplifying technique that reduces the complexity of this estimation problem

is to estimate each local correlation $\rho_{k\ell}(\boldsymbol{z})$ as a bivariate problem by only considering the corresponding pair of observation vectors $\{R_{ik}, R_{i\ell}\}$, $i = 1, \ldots, n$. Thus, we reduce the p-variate problems of estimating the local parameters depending on all coordinates, to a series of bivariate problems of estimating pairwise local correlations depending on their respective pairs of coordinates. In this way, we obtain a simplification that is analogous to an additive approximation in nonparametric regression. In the event that the resulting local covariance matrix is not positive definite, it is adjusted with the method described in Higham [2002]. his technique is applied in the empirical analysis that follows. For more details regarding this pairwise modeling approach, see Otneim and Tjøstheim [2017]. In particular, they show that the convergence speed is improved to $(nb^2)^{-1/2}$.

As already mentioned, the local estimates depend on the smoothing device - the bandwidth vector $\boldsymbol{b}$ and a specific choice of the kernel function, K. In the empirical analysis, we use the Gaussian kernel, and the bandwidth selector used is the plug-in selector suggested in Støve et al. [2014] – the global standard deviation of the observations times a constant equal to 1.1.

Finally, we note that the local Gaussian correlation has been used in several studies examining the dependence structure between asset returns, and in testing for financial contagion, see e.g. Støve and Tjøstheim [2014], Støve et al. [2014], Bampinas and Panagiotidis [2017] and Nguyen et al. [2020].

2 Descriptive statistics portfolio strategies

Descriptive statistics of the portfolio strategies out-of-sample returns are shown in Table 1. We report the mean, standard deviation, skewness, kurtosis, minimum value, maximum value, and the maximum portfolio drawdown, which is the maximum observed loss from a peak to a trough of the portfolio before a new peak is attained, for window size $M = 120$ (top) and $M = 240$ (bottom).

The mean return tends to increase with the different local Gaussian approaches, and all portfolios achieve moderately higher average returns. In the $M = 120$ case, MVS-L reaches the highest mean, followed by MVSC-L, MIN-L, MINC-L, MVS, MVSC, MIN, MVSC and EW. For the $M = 240$ window size, the average return ranking is MIN-L, MINC-L, MVS-L, MVSC-L/MINC, MIN, MVS, MVSC and EW. These findings indicate that the local Gaussian approach may capture asymmetries and outperform the corresponding benchmark models.

The lowest standard deviation for $M = 120$ is achieved by MINC-L. However, this is an exception, as all other strategies have slightly higher values when the local Gaussian method is applied. For $M = 240$, all local Gaussian portfolios have moderately higher standard deviations, except for MINC-L. As noted in Low et al. [2013], this can be due to a larger upside variation, which is desirable for investors. We will follow their approach and include downside risk measures when evaluating performance below.

Table 1: Descriptive statistics of the different portfolio strategies examined

	Mean	Std.dev.	Skewness	Kurtosis	Min	Max	Max. drawdown
Window size M = 120							
EW	0.423	1.999	-0.714	4.052	-11.916	7.342	22.857
MVS	0.455	1.492	-0.209	1.645	-5.451	5.778	9.479
MVSC	0.444	1.486	-0.256	1.644	-5.451	5.688	9.479
MIN	0.435	1.426	-0.190	1.211	-5.108	5.382	9.404
MINC	0.427	1.430	-0.195	1.181	-5.108	5.382	9.404
MVS-L	0.491	1.539	-0.214	1.693	-5.701	6.177	8.959
MVSC-L	0.484	1.494	-0.156	1.592	-5.020	6.079	8.881
MIN-L	0.462	1.459	-0.218	1.182	-5.194	5.429	8.577
MINC-L	0.460	1.401	-0.080	1.222	-4.620	5.743	8.065
Window size M = 240							
EW	0.376	2.158	-0.750	4.356	-11.916	7.342	22.857
MVS	0.447	1.671	-0.672	3.012	-8.294	5.061	15.677
MVSC	0.440	1.670	-0.703	3.047	-8.294	5.122	15.677
MIN	0.468	1.619	-0.564	2.316	-7.493	5.053	13.530
MINC	0.470	1.624	-0.558	2.272	-7.476	5.053	13.524
MVS-L	0.488	1.765	-0.243	3.461	-8.522	7.516	15.354
MVSC-L	0.470	1.737	-0.348	3.434	-8.522	7.182	15.364
MIN-L	0.503	1.673	0.805	7.357	-6.023	11.198	10.458
MINC-L	0.495	1.604	0.012	1.981	-5.926	7.142	11.036

All strategy returns exhibit slight negative skewness, except for MIN-L and MINC-L for the $M = 240$ window, with values of 0.805 and 0.012, respectively. Disregarding the EW strategies, the largest negative skew of -0.703 can be found in MVSC, for $M = 240$. Hence, the strategies are all moderately skewed or approximately symmetric.

The MIN-L for $M = 240$ holds the largest kurtosis value in the analysis. When examining the minimum and maximum returns for this strategy, we observe larger values for both. The MIN-L also achieves the lowest drawdown for $M = 240$. The smallest maximum drawdown for $M = 120$ is produced by MINC-L. Overall, the local gaussian strategies all have lower maximum drawdowns when compared to their benchmarks in both windows.

3 Portfolio rebalancing and terminal wealth, $M = 120$

In this section and the following, we present a corresponding performance analysis of the portfolio selection strategies as found in the main article and an estimation window of $M = 120$ months. We see that the smaller estimation window naturally leads to more sampling variation in the estimates and thus in the performance metrics of the various portfolio selection strategies, but the conclusions from the main article remain essentially unchanged.

Table 2 provides a summary of the portfolio rebalancing analysis and the terminal wealth

Table 2: Portfolio rebalancing and terminal wealth (based on an initial investment of $1) for the different strategies considered

	$\bar{\sigma}(\hat{w}_{k,M})$	Max. adj.	Min. adj.	Avg.turnover	Wealth	Wealth incl.tcost
Window size M = 120						
EW	0	0	0	0	4.052	4.052
MVS	20.671	18.759	-14.112	8.402	4.680	4.483
MVSC	17.944	18.737	-17.349	7.117	4.506	4.346
MIN	19.332	6.966	-8.769	4.135	4.376	4.285
MINC	18.536	6.966	-8.769	3.533	4.254	4.178
MVS-L	20.760	122.681	-132.874	32.084	5.283	4.484
MVSC-L	16.521	29.529	-30.509	17.045	5.166	4.736
MIN-L	21.174	117.820	-147.244	38.222	4.775	3.927
MINC-L	17.710	45.461	-46.261	18.301	4.747	4.323
Window size M = 240						
EW	0	0	0	0	2.231	2.231
MVS	16.663	9.228	-6.741	6.325	2.605	2.551
MVSC	15.813	9.228	-7.426	5.622	2.571	2.524
MIN	16.249	5.036	-4.434	3.090	2.729	2.701
MINC	15.952	5.039	-4.428	2.793	2.742	2.717
MVS-L	15.660	73.184	-89.316	16.967	2.855	2.698
MVSC-L	14.253	18.096	-26.245	12.105	2.745	2.637
MIN-L	17.014	53.636	-79.931	21.617	2.953	2.749
MINC-L	15.243	19.646	-26.028	13.762	2.910	2.780

reached by each of the strategies. We present the results for $M = 240$ as well for easy comparison.

Figure 4 shows wealth accumulation and drawdowns for the hypothetical investment of $1 in each of the nine strategies included in the analysis when a sampling window of $M = 120$ is used. As seen in the upper part of the figure, the local Gaussian MVS-L produces the largest final wealth when disregarding trade costs. It remains top-ranked during most months in the sample and suffers from smaller drawdowns in volatile periods such as the 2008 financial crisis. When transaction costs (15 basis points per transaction, as in the main paper) are considered, the strategy still performs well but is pushed down from the top position by the constrained MVSC-L, which has a lower turnover.

4 Comparison between Local and Global portfolio strategies

This section provides a comparison between the local and global strategies. As the difference between these strategies is the choice of covariance matrix, we start by inspecting the differences between local and global variances and correlations.

Figure 1 displays the local variances of the asset returns, i.e. the diagonal of the matrix

$\widehat{\boldsymbol{\Sigma}}(\boldsymbol{r})$, together with their global counterparts. For BMUK10Y and BMUS10Y (top panel), the local estimates agree well with the global estimates. However, for the four other assets (the two lower panels), there are two interesting differences. Firstly, all four local estimates capture the increased volatility during the 2008 financial crisis. Secondly, the local variances are typically smaller than the global ones in other periods, particularly after the 2008 financial crisis. This is expected since the local variance estimates will quickly reflect high volatility once estimated at extreme evaluation points $\boldsymbol{r}_t$. Moreover, in less extreme evaluation points, events such as the 2008 financial crisis will be smoothed away in the estimation of $\boldsymbol{\Sigma}_t(\boldsymbol{r}_t)$, regardless of window size. In contrast, the global estimates suffer from an averaging effect that may result in variances that are neither representative for the high volatility periods (too low) or calm periods (too high).

Figure 2 displays the local and global correlations between the asset returns. Though not to the same degree, Figure 2 displays similar characteristics as Figure 1; The local correlations between many assets increases significantly during the 2008 financial crisis and for these assets the global correlations are typically are larger than the local ones in the period that follows.

Figure 3 displays the weights of the global and local MINC strategy (MINC and MINC-L) and mainly reflects the results displayed in 1;, i.e., increasing variances translates to higher risk and thus lower weights. For example, the assets BMUS10Y,EWCI, GSGCSPT, and S&P500 are assigned higher local weights after the 2008 financial crisis compared to the global weights, which, apart from BMUS10Y, is in line with our observations in Figure 1. It is difficult to evaluate the impact of correlations on the weights directly. However, it is reasonable to believe that better diversifications are obtained since $\widehat{\boldsymbol{\Sigma}}(\boldsymbol{r}_t)$ picks up the increased dependence between assets during the 2008 financial crisis and perhaps reflects a more realistic dependence structure estimates (weaker) later on.

Figure 1: Local and global variances of asset returns.

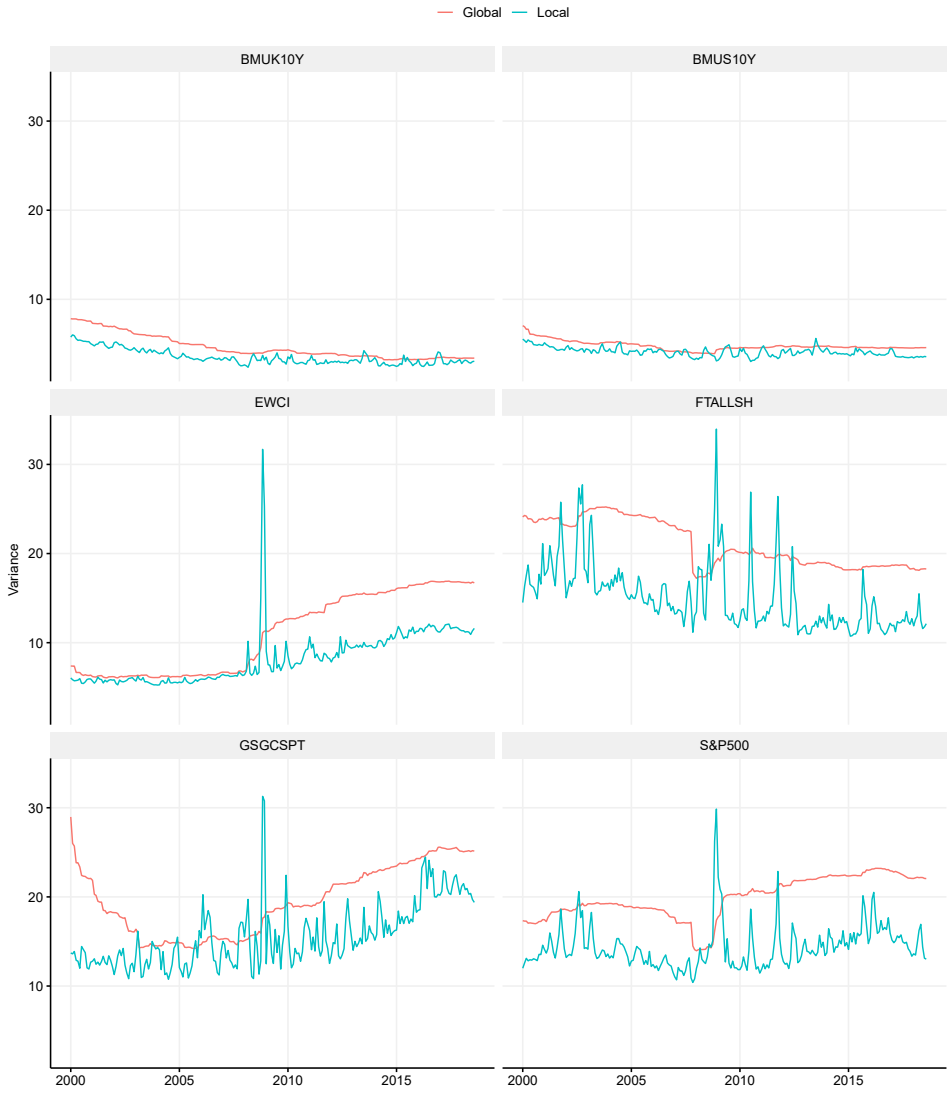

Figure 2: Local and global correlations between asset returns.

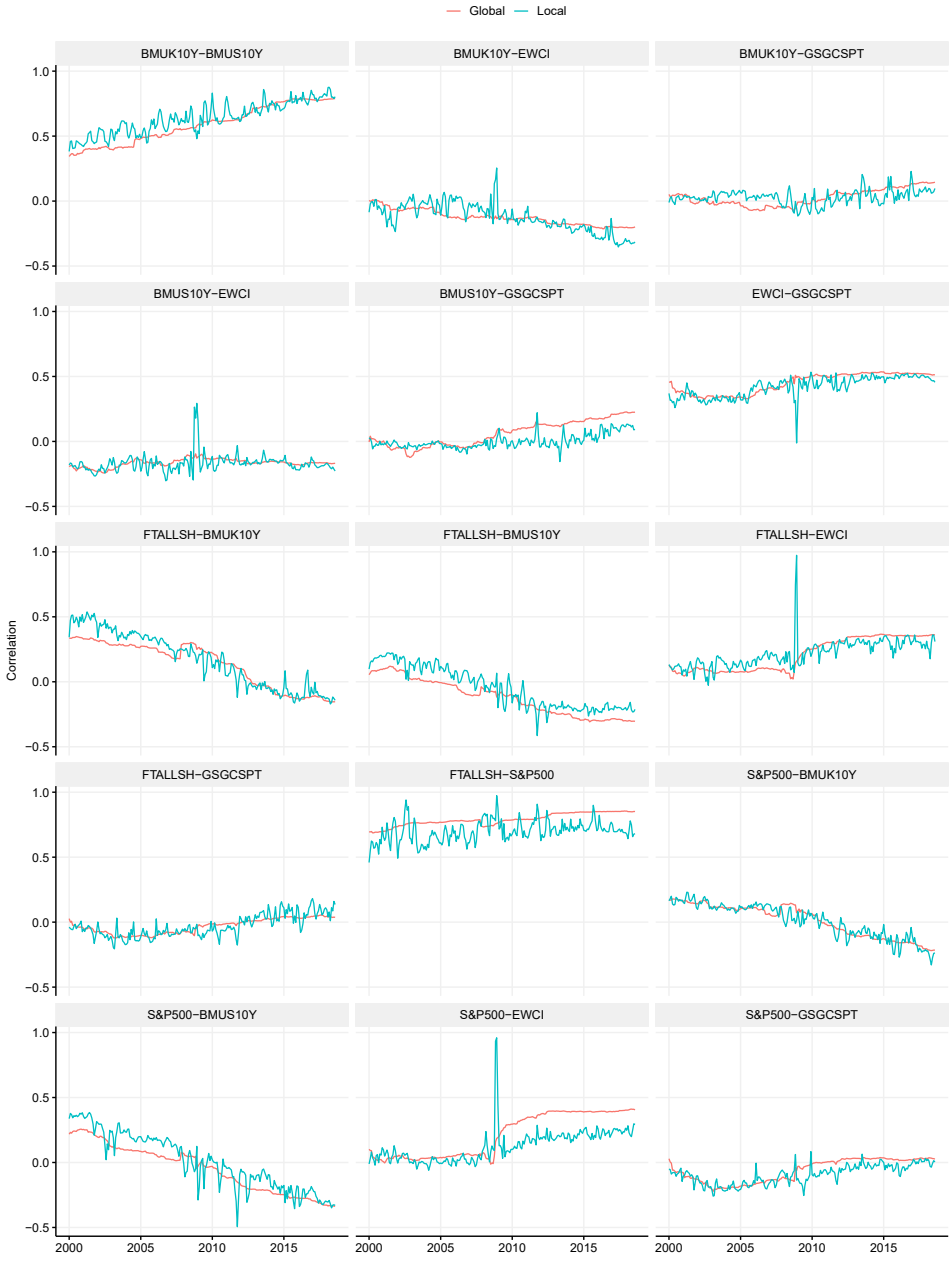

Figure 3: Local and global portfolio weights for the MINC strategy.

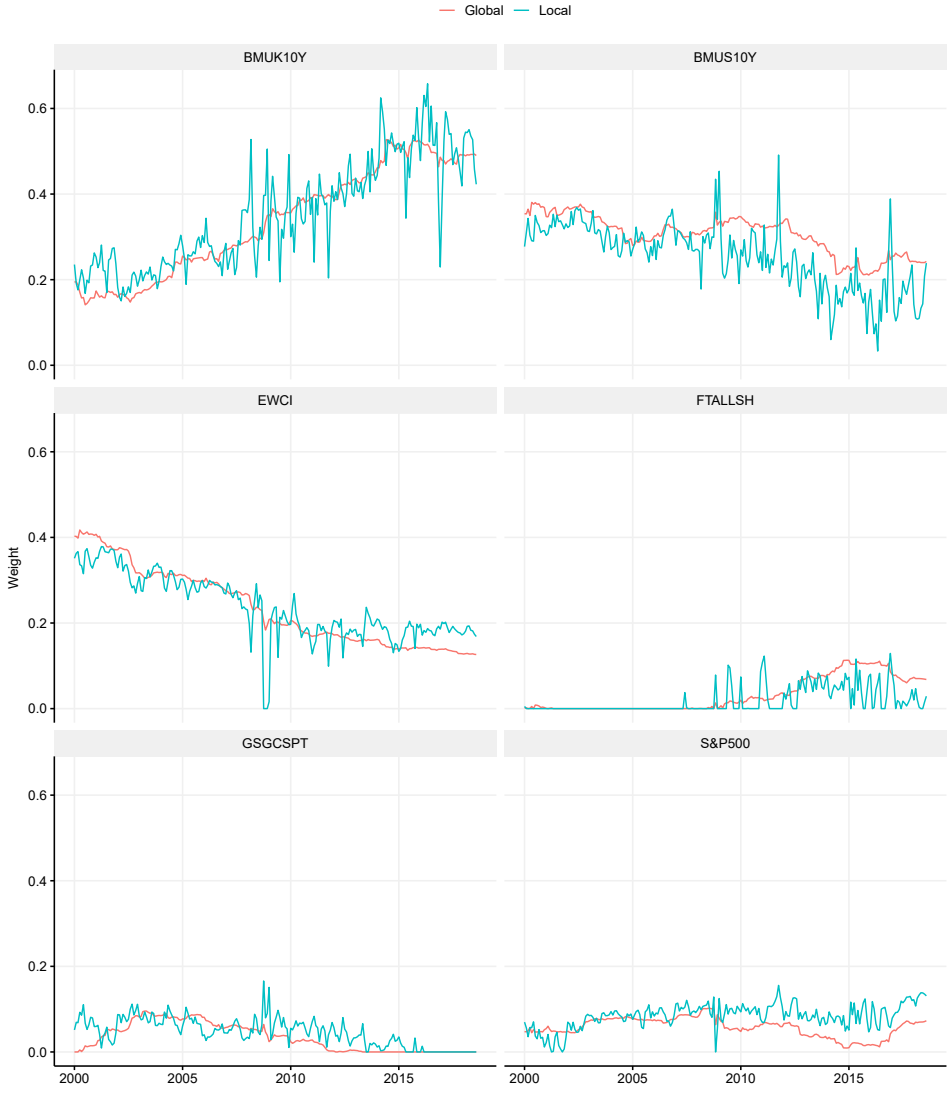

5 Evaluation of risk-adjusted performance, $M = 120$

Table 3 reports out-of-sample performance by evaluating portfolio strategy returns using the same metrics as in the main article. Again, we include the corresponding results for $M = 240$ here as well for easy comparison between the two strategies. When using the smaller estimation windows, all but one of the performance measures prefer the long-only version of the local minimum variance portfolio, the only exception being the CEQ, which prefers the unconstrained local mean-variance portfolio, both with and without transaction costs. Furthermore, we have performed the z-test of Ledoit and Wolf [2008] which is applied to the Sharpe Ratios to indicate the statistical differences for all MV optimizations against the $1/N$ benchmark.

For windows $M = 120$ and $M = 240$, the remaining top rankings are: MINC-L and MIN-L (VaR Sharpe), MINC-L and MIN-L (ES Sharpe ratio), MVSC-L and MINC-L (CEQ), MINC-L and MIN-L (Sortino ratio), MINC-L and MIN-L (Omega ratio). These results again confirms that the local Gaussian approach seems to improve performance.

Figure 4: Wealth accumulation for the different strategies based on an initial investment of \$1, using a rolling window of size $M = 120$ monthly observations, top plot excluding transaction costs, bottom plot including transaction cost of 15 basis point

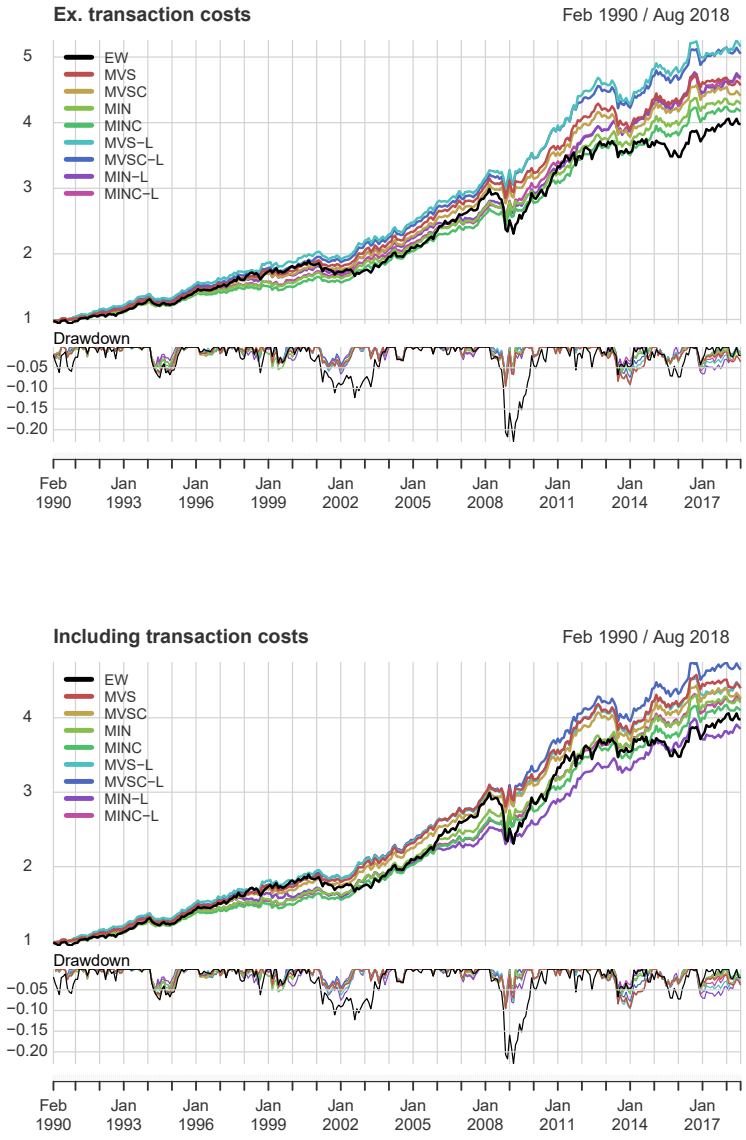

Table 3: Portfolio strategies out-of-sample performance

	Sharpe	VaR Sharpe	ES Sharpe	Ann. Sharpe	CEQ	Sortino	Omega
Panel A: Ex. transaction costs							
Window size M = 120							
EW	0.212	0.137	0.071	0.715	0.403	0.327	1.737
MVS	0.305	0.224	0.143	1.057	0.444	0.525	2.221
MVSC	0.299	0.216	0.138	1.035	0.433	0.508	2.191
MIN	0.305	0.223	0.148	1.057	0.425	0.528	2.199
MINC	0.299	0.217	0.145	1.034	0.417	0.514	2.169
MVS-L	0.319*	0.237	0.150	1.108	**0.479**	0.557	2.298
MVSC-L	0.324*	0.244	0.157	1.126	0.473	0.569	2.315
MIN-L	0.317*	0.232	0.154	1.099	0.451	0.555	2.264
MINC-L	**0.328***	**0.250**	**0.168**	**1.141**	0.450	**0.587**	**2.340**
Window size M = 240							
EW	0.174	0.110	0.057	0.577	0.353	0.263	1.593
MVS	0.267	0.179	0.099	0.919	0.433	0.427	2.025
MVSC	0.264	0.175	0.097	0.905	0.426	0.417	2.009
MIN	0.289*	0.198	0.116	0.999	0.455	0.474	2.120
MINC	0.290*	0.199	0.117	1.000	0.457	0.474	2.120
MVS-L	0.276*	0.203	0.115	0.951	0.472	0.472	2.126
MVSC-L	0.270*	0.193	0.107	0.929	0.454	0.453	2.075
MIN-L	0.301*	**0.317**	**0.317**	1.041	**0.489**	**0.567**	**2.303**
MINC-L	**0.309***	0.240	0.157	**1.070**	0.482	0.554	2.268
Panel B: Incl. transaction costs							
Window size M = 120							
EW	0.215	0.140	0.072	0.727	0.410	0.332	1.753
MVS	0.303	0.222	0.141	1.051	0.440	0.521	2.213
MVSC	0.298	0.215	0.136	1.031	0.430	0.506	2.188
MIN	0.307	0.225	0.149	1.064	0.426	0.532	2.213
MINC	0.301	0.219	0.145	1.041	0.419	0.517	2.183
MVS-L	0.294	0.212	0.132	1.015	0.440	0.500	2.152
MVSC-L	**0.313***	0.233	0.150	**1.087**	**0.456**	0.545	2.256
MIN-L	0.282*	0.199	0.130	0.971	0.401	0.476	2.065
MINC-L	**0.313**	**0.236**	**0.158**	**1.087**	0.429	**0.553**	**2.256**
Window size M = 240							
EW	0.179	0.113	0.058	0.593	0.362	0.270	1.613
MVS	0.261	0.174	0.097	0.895	0.423	0.415	1.989
MVSC	0.258	0.171	0.095	0.885	0.418	0.407	1.978
MIN	0.285*	0.194	0.114	0.982	0.449	0.465	2.095
MINC	0.285*	0.195	0.115	0.985	0.451	0.466	2.097
MVS-L	0.262	0.188	0.106	0.900	0.448	0.440	2.041
MVSC-L	0.260	0.183	0.101	0.890	0.438	0.430	2.014
MIN-L	0.281	**0.279**	**0.279**	0.969	0.456	0.518	2.174
MINC-L	**0.295***	0.225	0.148	**1.019**	**0.461**	**0.522**	**2.182**

References

G. Bampinas and T. Panagiotidis. Oil and stock markets before and after financial crises: A local gaussian correlation approach. *Journal of Futures Markets*, 37(12):1179–1204, 2017.

G. D. Berentsen and D. Tjøstheim. Recognizing and visualizing departures from independence in bivariate data using local Gaussian correlation. *Statistics and Computing*, 24(5):785–801, 2014.

G.D. Berentsen, B. Støve, D. Tjøstheim, and T. Nordbø. Recognizing and visualizing copulas: an approach using local gaussian approximation. *Insurance: Mathematics and Economics*, 57:90–103, 2014.

N.J. Higham. Computing the nearest correlation matrix—a problem from finance. *IMA journal of Numerical Analysis*, 22(3):329–343, 2002.

N.L. Hjort and M.C. Jones. Locally parametric nonparametric density estimation. *Annals of Statistics*, 24(4):1619–1647, 1996.

M. Chris Jones. The local dependence function. *Biometrika*, 83(4):899–904, 1996.

L.A. Jordanger and D. Tjøstheim. Nonlinear spectral analysis: A local gaussian approach. *Journal of the American Statistical Association*, pages 1–55, 2020.

V. Lacal and D. Tjøstheim. Local gaussian autocorrelation and tests of serial dependence. *Journal of Time Series Analysis*, 38(1):51–71, 2017.

V. Lacal and D. Tjøstheim. Estimating and testing nonlinear local dependence between two time series. *Journal of Business and Economic Statistics*, 37(4):648–660, 2019.

Oliver Ledoit and Michael Wolf. Robust performance hypothesis testing with the sharpe ratio. *Journal of Empirical Finance*, 15(5):850–859, 2008.

R. K. Low, J. Alcock, R. Faff, and T. Brailsford. Canonical vine copulas in the context of modern portfolio management: Are they worth it? *Journal of Banking & Finance*, 37(8): 3085–3099, 2013.

Q.N. Nguyen, S. Aboura, J. Chevallier, Z. Lyuyuan, and B. Zhu. Local gaussian correlations in financial and commodity markets. *European Journal of Operational Research*, 2020.

H. Otneim and D. Tjøstheim. The locally Gaussian density estimator for multivariate data. *Statistics and Computing*, 27(6):1595–1616, 2017.

H. Otneim and D. Tjøstheim. Conditional density estimation using the local gaussian correlation. *Statistics and Computing*, 28(2):303–321, 2018.

H. Otneim and D. Tjøstheim. The locally gaussian partial correlation. *Journal of Business & Economic Statistics*, pages 1–13, 2021.

B. Støve and D. Tjøstheim. Measuring asymmetries in financial returns: An empirical investigation using local gaussian correlation. In M. Meitz N. Haldrup and P. Saikkonen, editors, *Essays in Nonlinear Time Series Econometrics*, pages 307–329. Oxford University Press, Oxford, 2014.

B. Støve, D. Tjøstheim, and K.O. Hufthammer. Using local gaussian correlation in a nonlinear re-examination of financial contagion. *Journal of Empirical Finance*, 25:785–801, 2014.

D. Tjøstheim and K.O. Hufthammer. Local Gaussian correlation: A new measure of dependence. *Journal of Econometrics*, 172:33–48, 2013.

D. Tjøstheim, H. Otneim, and B. Støve. Statistical dependence: Beynd pearson's ρ. *Statistical Science, to appear*, 2021.

Paper B

7.2 etrm: Energy Trading and Risk Management in R

Anders D. Sleire, (2022), The R-Journal, 14/1, p. 320-341.

etrm: Energy Trading and Risk Management in R

by Anders D. Sleire

Abstract This paper introduces **etrm**, an R package with tools for trading and financial risk management in energy markets. Contracts for electric power and natural gas differ from most other commodities due to the fact that physical delivery takes place over a time interval, and not at a specific point in time. There is typically strong seasonality, limited storage and transmission capacity and strong correlation between price and required volume. Such characteristics need to be taken into account when pricing contracts and managing financial risk related to energy procurement. Tools for these task are usually bundled into proprietary Energy Trading Risk Management (ETRM) systems delivered by specialized IT vendors. The **etrm** package offers a transparent solution for building a forward price curve for energy commodities which is consistent with methods widely used in the industry. The user's fundamental market view may be combined with contract price quotes to form a forward curve that replicate current market prices, as described in Ollmar (2003) and Benth et al. (2007b). **etrm** also provides implementations of five portfolio insurance trading strategies for energy price risk management. The forward market curve and the energy price hedging strategies are core elements in an ETRM system, which to the best of the author's knowledge has not been previously available in the R ecosystem.

1 Introduction

The purpose of this paper is to introduce the R package **etrm** and its tools for energy trading and financial risk management. Substantial fluctuations in energy prices represent a significant risk for market players, in particular for large consumers, producers and utility companies, see Benini et al. (2002). The price dynamics is complex due to strong weather dependency and physical constraints related to storage, distribution, and the introduction of new technology. See for example Nicolosi (2010) for an analysis of renewable energy production and the negative prices following extreme events in the German power market. Derivatives securities, such as futures contracts, are often used to hedge against the commodity price risk. Specialized Energy Trading and Risk Management (ETRM) systems provide the necessary tools to handle key activities such as position management, valuation and risk reporting. Several proprietary alternatives exist. The annual *Energy Risk's Software Survey* in Farrington (2020) gives an overview of major providers along with rankings based on industry polls. There, ETRM systems are divided into the operational categories *derivatives software*, *physical trading and operations software* and *front- and middle-office functionality*. Historically, many system providers within this domain have bundled modules into large monolithic architectures serving a variety of purposes, including accounting and regulatory reporting. During the last decade, a general trend within system development has moved towards splitting software into smaller stand-alone components. The **etrm** package solely focus on financial trading, and may be viewed as a module for front- and middle-office functionality for energy derivatives. The package currently offers transparent tools for two main ETRM activities 1) construction of forward market curves and 2) implementation of trading strategies for price risk management.

After the liberalization of electricity and gas markets started in the 1990s, a rich research literature emerged. Topics studied include pricing and hedging in the forward market and modelling of spot price processes, see for example Bessembinder and Lemmon (2002), Janczura et al. (2013) and Benth et al. (2007a). Alternative methods for pricing options in power markets can be found in Burger et al. (2004) and Benth and Schmeck (2014). Textbooks such as Eydeland and Wolyniec (2002), Benth et al. (2008) and Kirschen and Strbac (2018) may be used to gain a more detailed overview of market structure, available instruments, methods for risk management and the related markets for fuel, freight and weather products. In this paper, we will cover some of the theory regarding forward curve modelling from Ollmar (2003) and Benth et al. (2007b). The theoretical framework for price risk management is gathered from the *portfolio insurance* literature, see Leland (1980), Perold and Sharpe (1988), Leland and Rubinstein (1976). This will be presented in further detail below.

We would like to note that there are some tools available outside the domain of proprietary ETRM software providers. Two examples are the MathWorks case studies Sundar (2021) and Deoras (2021), focusing on risk assessment for gas-fired power plants and electricity load and price forecasting using MATLAB. These topics are however somewhat ad-hoc, and the supplied code cannot be easily incorporated into a generic ETRM system for general use. Similarly within the R ecosystem, there are related tools in the Rmetrics suite of packages, such as **fOptions** and **fPortfolio**, but they are not

directly applicable. Due to the unique properties of energy markets, standard methods for generating forward curves in interest rate markets cannot be used either. To fill this gap and provide practitioners and researchers with tools dedicated to energy price risk management, we have created **etrm**. The package is available on CRAN, and may be installed and loaded into the R environment by running the following commands:

```
if(!require(etrm)==TRUE) {install.packages("etrm")}
library(etrm)
```

The rest of the article is organized as follows. First, we give a brief introduction to energy market forward curves and the maximum smoothness forward curve (MSFC) model. We describe the **etrm** implementation and provide examples of use. Second, a short treatment of energy price risk management with futures contracts is provided, followed by a presentation of five portfolio insurance models. The **etrm** implementation is described and illustrated with practical examples for energy portfolios with both short and long market exposure. The third part provide an overview of the **etrm** package structure, available functions and included data sets. Finally, the last section summarizes the paper and provide some suggestions for future work.

2 Energy market forward price curves

The standardised forwards for electricity and gas are contracts for flow delivery. The underlying commodity is not received at a fixed point in time, but over a time interval. In mature markets, participants can trade a variety of products, both over-the-counter (OTC) and on exchanges such as Nasdaq Commodities, European Energy Exchange and the Intercontinental Exchange. Liquidity is often best in the so called front-products, and there is normally higher activity in contracts for next week, month, quarter and year compared to similar products further ahead in time. Shorter period contracts may not even be available on a longer horizon, and seasonal price variation is thus not directly observable in prices far ahead in time. Transacted volume and prices also inhibit quite pronounced seasonality, during the year, week and within a specific day. For this reason, forward contracts are divided into categories based on a *load pattern*. In the *base load* contracts, volume is delivered at a constant rate during the contract period, while *peak load* products are linked to high volume hours, such as Monday to Friday from 8 am to 8 pm. Other, more exotic load patterns exist, but they are less common. Further details can be found in Eydeland and Wolyniec (2002) and Kirschen and Strbac (2018).

The aim of the energy market forward price curve calculation is to create a compact representation of the forward market, at a given point in time. The curve must be able to price the quoted instruments correctly, while accounting for typical energy market characteristics such as seasonality and (possibly overlapping) contracts for flow delivery. The curve is an essential decision making tool with many uses, such as pricing non-standard supply agreements, investment decisions and risk management.

The topic of forward curve fitting has been studied for decades in interest rate markets, see for example McCulloch (1971) and Anderson et al. (1996). These techniques cannot be applied directly to commodities with flow delivery and strong seasonality in prices. There are several alternative approaches to calculating a forward price curve for energy commodities. In Fleten and Lemming (2003), market data is combined with forecasts generated by a bottom-up model constrained by the bid/ask spread in order to meet the no-arbitrage condition. Borak and Weron (2008) propose a semiparametric factor model for the forward curve dynamics in electricity markets, while Hildmann et al. (2012) develop a calculation method by combining parametric estimation and prediction of futures prices under constraints.

In **etrm**, we have opted for a method that combines a seasonal function with the maximum smoothness-approach from interest rate markets, see Adams and Van Deventer (1994). Hence, we follow the approach in Ollmar (2003) and Benth et al. (2007b). *Base load* contracts are used to calculate a curve with daily granularity. This method has several benefits. It produces a continuous curve which has a closed form solution, it is fast to calculate, flexible and used by many practitioners in the industry. The next section provide a brief overview of the method, followed by a description of the **etrm** implementation with examples.

Maximum smoothness forward curve model

Consider a market at time t with m forward contracts available for trading. Let the list

$$S_t = \{(\tau_1^s, \tau_1^e), (\tau_2^s, \tau_2^e), ..., (\tau_m^s, \tau_m^e)\}$$

contain the start and end dates for each of these contracts. The time distance between τ_i^s and τ_i^e for a contract i in $1, .., m$ cover standardized periods such as *week*, *month*, *quarter* and *year*. Some of these settlement intervals might overlap, and in order to handle this we create a new list of dates $\{t_0, t_1, ..., t_n\}$ to identify each separate sub period, see Figure 1. The new list is made by sorting all dates in S_t in ascending order and removing duplicates.

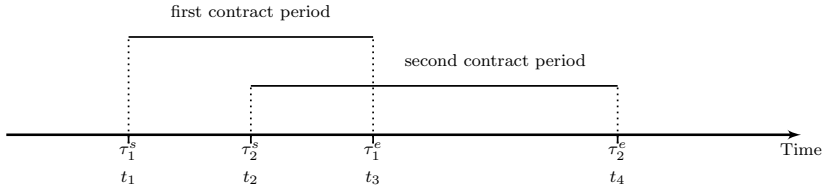

Figure 1: Illustration of two overlapping contracts with start (τ^s) and end (τ^e) dates. Due to the overlap, the total delivery period is split into sub intervals identified with $\{t_1, t_2, t_3, t_4\}$.

The forward price at time t for one unit of energy delivered at a constant rate $(\tau^e - \tau^s)^{-1}$ over the time interval (τ^s, τ^e) is denoted by $F(t, \tau^s, \tau^e)$, where $t \leq \tau^s < \tau^e$. A forward contract for a flow delivery may be thought of as the average of hypothetical single-delivery contracts. At time t, each of these would have a unique price $f(t, u)$ for the delivery at u with an infinitesimal delivery period. This leads to $F(t, \tau^s, \tau^e)$ being the weighted average

$$F(t, \tau_s, \tau_e) = \int_{\tau^s}^{\tau^e} w(u, \tau_s, \tau_e) f(t, u)\, \mathrm{d}u \tag{1}$$

where $w(u, \tau_s, \tau_e) = \frac{\hat{w}(u)}{\int_{\tau_s}^{\tau_e} \hat{w}(v) dv}$ is a weight function accounting for the rate of interest r and the time value of money. If the contract in question is settled at the end of the delivery period (forward contract), the weight function is given by $w(u, \tau_s, \tau_e) = 1/(\tau_e - \tau_s)$. If the contract is settled continuously over the delivery period (futures contract), $w(u, \tau_s, \tau_e) = \frac{re^{-ru}}{e^{-r\tau_s} - e^{-r\tau_e}}$. In the following we construct a forward market price curve for the entire horizon using a simplified notation $f(u)$ for the function describing the forward curve at time t. In order to model the strong seasonality in energy markets, the forward curve function is decomposed into two elements:

$$f(u) = \Lambda(u) + \epsilon(u) \qquad u \in [t_0, t_n] \tag{2}$$

Following Ollmar (2003) we calculate $f(u)$ by combining a prior function $\Lambda(u)$ which contain our subjective views on the future prices with an adjustment function $\epsilon(u)$ to ensure match with the observed closing prices for the m contracts. The prior could be generated with a simple sinusoidal function or from a fundamental model more capable of describing the seasonality and calendar effects observed in energy markets. Should the prior be excluded, the seasonal price patterns will not be visible in the far end of the curve, where only yearly or seasonal contracts are available. Smoothing is calculated on the adjustment function, we aim to minimize the total curvature of $\Lambda(u)$ while preserving the information from the prior. Smoothness is defined as the integral of the second-order derivative of the function, and the smoothest possible curve over $[t_0, t_n]$ is achieved by minimising

$$\int_{t_0}^{t_n} [\epsilon''(u)]^2\, \mathrm{d}u$$

under five constraints presented below. Lim and Xiao (2002) show the smoothest possible curve is found when the n sub periods are modelled by fourth-degree polynomials. We write $\epsilon(u)$ as a spline

$$\epsilon(u) = \begin{cases} a_1 u^4 + b_1 u^3 + c_1 u^2 + d_1 u + e_1, & u \in [t_0, t_1], \\ a_2 u^4 + b_2 u^3 + c_2 u^2 + d_2 u + e_2, & u \in [t_1, t_2], \\ \quad . & \\ \quad . & \\ a_n u^4 + b_n u^3 + c_n u^2 + d_n u + e_n, & u \in [t_{n-1}, t_n]. \end{cases}$$

In order to construct the forward curve function we need to identify the parameters of $\epsilon(u)$

$$\mathbf{x}^\intercal = [a_1\ b_1\ c_1\ d_1\ e_1 \quad a_2\ b_2\ c_2\ d_2\ e_2 \ldots\ a_n\ b_n\ c_n\ d_n\ e_n]$$

by solving the quadratic optimisation problem

$$\min_x \int_{\tau^s}^{\tau^e} [\epsilon''(u;x)]^2\,\mathrm{d}u \tag{3}$$

subject to

$$\begin{aligned}
(a_{j+1}-a_j)u_j^4 + (b_{j+1}-b_j)u_j^3 + (c_{j+1}-c_j)u_j^2 + (d_{j+1}-d_j)u_j + e_{j+1} - e_j &= 0 \quad \text{(a)}\\
4(a_{j+1}-a_j)u_j^3 + 3(b_{j+1}-b_j)u_j^2 + 2(c_{j+1}-c_j)u_j + (d_{j+1}-d_j) &= 0 \quad \text{(b)}\\
12(a_{j+1}-a_j)u_j^2 + 6(b_{j+1}-b_j)u_j + 2(c_{j+1}-c_j) &= 0 \quad \text{(c)}\\
\epsilon'(u_n;x) &= 0 \quad \text{(d)}\\
\frac{1}{\tau_i^e-\tau_i^s}\int_{\tau_i^s}^{\tau_i^e} (\Lambda(u)+\epsilon(u))\,\mathrm{d}u &= F_i^c \quad \text{(e)}
\end{aligned}$$

for spline knot $j = 1, ..., n-1$ and contract $i = 1, ..., m$. The constraint in (a) ensures the adjustment function is continuous in the knots, while (b) and (c) imposes this restriction also for the first and second order differentials. The (d) constraint require the adjustment function to be horizontal at time T, and finally (e) also require the average value of the forward price function $f(u)$ over the delivery period for contract i to match the quoted closing price F_i^c. Here, we could take the interest rate effect from r into account and set the present value of the average of the forward price function equal to present value of the forward contract. Instead of doing that, we will follow Benth et al. (2008) and Ollmar (2003) and assume $r = 0$ such that the weight function in 1 is approximated with $w(u, \tau^s, \tau^e) \approx \frac{1}{\tau^e-\tau^s}$. Like Benth et al. (2008) we will argue that both the prior and the smoothing will outweigh a marginal interest rate effect. This minimisation problem can be expressed as

$$\min_{\mathbf{x}} \quad \mathbf{x}^\intercal \mathbf{H} \mathbf{x},$$

where $\mathbf{x}$ is a $(5n \times 1)$ vector and

$$\mathbf{H} = \begin{bmatrix} h_1 & & 0 \\ & \ddots & \\ 0 & & h_n \end{bmatrix}, \mathbf{h}_j = \begin{bmatrix} \frac{144}{5}\Delta_j^5 & 18\Delta_j^4 & 8\Delta_j^3 & 0 & 0 \\ 18\Delta_j^4 & 12\Delta_j^3 & 6\Delta_j^2 & 0 & 0 \\ 8\Delta_j^3 & 6\Delta_j^2 & 4\Delta_j & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

The block diagonal matrix $\mathbf{H}$ has dimensions $(5n \times 5n)$ and $\Delta_j = t_{j+1}^l - t_j^l$. As the constraints in (3) are all linear in x and may be expressed on the form $\boldsymbol{Ax} = \boldsymbol{B}$, the problem may be rephrased as an unconstrained minimisation problem via the Lagrange multiplier method:

$$\min_{\mathbf{x},\lambda} \quad \boldsymbol{x}^\intercal \boldsymbol{H}\boldsymbol{x} + \boldsymbol{\lambda}^\intercal(\boldsymbol{A}\boldsymbol{x} - \boldsymbol{B})$$

where $\mathbf{A}$ is a $(3n + m - 2 \times 5n)$ matrix and $\mathbf{B}$ is a $(3n + m - 2 \times 1)$ vector. The spline parameter vector and Lagrange multipliers are identified by solving

$$\begin{bmatrix} \mathbf{2H} & \mathbf{A}^\intercal \\ \mathbf{A} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{x} \\ \boldsymbol{\lambda} \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \mathbf{B} \end{bmatrix} \tag{4}$$

where the left matrix has dimension of $(8n + m - 2) \times (8n + m - 2)$ and both vectors are of dimension $(8n + m - 2)$.

The "MSFC" class with examples

The forward curve calculation in **etrm** is implemented in the S4 `"MSFC"` class. By supplying all required arguments to the constructor function `msfc()`, the user may create an object that contains the calculation results, input arguments and further calculation details. In addition to the arguments from the list of contracts, the user may also provide a prior to the calculation. By default the prior is set to zero, but the user can input any vector expressing a belief regarding the market to be combined with the observed prices. An overview of the `msfc()` arguments can be found in Table 1.

Argument	Description	Default value
`tdate`	Trading date	none
`include`	Logical vector for contract selection	none
`contract`	Character vector with contract names	none
`sdate`	Date vector with contract start dates	none
`edate`	Date vector with contract end dates	none
`f`	Numeric vector with contract prices	none
`prior`	Numeric prior curve vector	0

Table 1: Arguments for the `msfc()` constructor function for forward curve calculation.

The `"MSFC"` class properties and available methods are summarised in Figure 2, and a brief description is provided in the following. The `"Name"` slot describe the type of forward curve model used by storing the character value "MSFC", while `"TradeDate"` keeps the trade date used in the calculation. A data frame containing details for selected contracts along with the calculated forward price based on the curve can be found in `"BenchSheet"`. A count of the $(n-1)$ number of polynomials used in the spline is available as a scalar value in `"Polynomials"`, and the prior curve vector in `"PriorFunc"`. The main calculation result is stored in a data frame which contains daily values for all selected contracts along with the calculated forward curve. The data frame span the date range from `"TradeDate"` to the end date of the contract furthest ahead in time, and can be found in `"Results"`. The interested user may also extract additional information regarding the spline itself. Coefficients for all polynomials can be found in the `"SplineCoef"` list and the knotpoints separating them in the numeric vector `"KnotPoints"`. Further details regarding the calculation of the daily forward curve values are available in the `"CalcDat"` data frame. This table is essentially an extended version of `"Results"`, where numeric time vectors and the spline coefficients have been added.

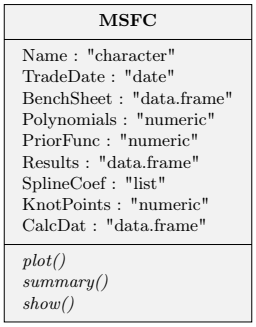

Figure 2: Attributes and methods of the `"MSFC"` class.

The `"MSFC"` class has the generic methods `plot()`, `summary()` and `show()`. The `plot()` method may be used to create a chart of the calculated curve and underlying contracts from the `"Results"` data frame. All plot methods in **etrm** are based on **ggplot2**, see Wickham (2011a). The `summary()` method returns a list with three elements; a description string, a sample of the prior vector, and the bench sheet. Finally, the `show()` method returns the `"Results"` data frame.

We proceed with a practical example using two of the embedded **etrm** data sets to represent information available to a European power market participant. All market-related inputs to the `msfc()` constructor (trade date and contract properties) are required arguments. These are collected from a synthetic data set for the trading date 2013-05-13, and can be found in `"powfutures130513"` presented in Table 2. Contracts covering long time spans are excluded with `include = FALSE` if futures of shorter duration are available for the same time interval in order to preserve the seasonality available in market prices. We use a seasonal prior with high energy prices during the winter season, followed by

a drop toward the lower summer levels. It also take into account some well known calendar effects, such as weekends. The prior is simple, and merely used for illustrative purposes. It can be found in the data set "powpriors130513" included in the package. A calculation excluding the prior is also added for comparison.

Include	Contract	Start	End	Closing
TRUE	W21-13	2013-05-20	2013-05-26	33.65
TRUE	W22-13	2013-05-27	2013-06-02	35.77
TRUE	W23-13	2013-06-03	2013-06-09	36.58
TRUE	W24-13	2013-06-10	2013-06-16	35.93
TRUE	W25-13	2013-06-17	2013-06-23	33.14
TRUE	W26-13	2013-06-24	2013-06-30	34.16
FALSE	MJUN-13	2013-06-01	2013-06-30	35.35
TRUE	MJUL-13	2013-07-01	2013-07-31	33.14
TRUE	MAUG-13	2013-08-01	2013-08-31	35.72
TRUE	MSEP-13	2013-09-01	2013-09-30	38.41
TRUE	MOCT-13	2013-10-01	2013-10-31	38.81
TRUE	MNOV-13	2013-11-01	2013-11-30	40.94
FALSE	Q3-13	2013-07-01	2013-09-30	35.72
TRUE	Q4-13	2013-10-01	2013-12-31	40.53
TRUE	Q1-14	2014-01-01	2014-03-31	42.40
TRUE	Q2-14	2014-04-01	2014-06-30	33.39
TRUE	Q3-14	2014-07-01	2014-09-30	31.78
TRUE	Q4-14	2014-10-01	2014-12-31	38.25
TRUE	Q1-15	2015-01-01	2015-03-31	40.73
TRUE	Q2-15	2015-04-01	2015-06-30	32.64
TRUE	Q3-15	2015-07-01	2015-09-30	30.87
TRUE	Q4-15	2015-10-01	2015-12-31	37.22
FALSE	CAL-14	2014-01-01	2014-12-31	36.43
FALSE	CAL-15	2015-01-01	2015-12-31	35.12
TRUE	CAL-16	2016-01-01	2016-12-31	34.10
FALSE	CAL-17	2017-01-01	2017-12-31	35.22
FALSE	CAL-18	2018-01-01	2018-12-31	36.36

Table 2: Closing prices for futures contracts used in the forward curve calculation for 2013-05-13. Contracts are selected for the calculations with the `include` vector. Prices for these contracts can be found as horizontal lines in Figure 3

As shown in Figure 3, the shorter contracts close in time to the trading date clearly reflect a seasonal pattern. This is typical in power markets, where weather and calendar effects have strong influence on transacted volume and price formation. On a longer horizon however, this information is not observable in market prices, as the quoted contracts cover longer time spans. This is where price data may be supplemented with prior knowledge in order to create a representation of the market consistent with both the underlying fundamentals and the listed contracts. The following code will create the "MSFC" objects and plot calculation results:

```
library(etrm)
library(gridExtra)
data(powfutures130513)
data(powpriors130513)

# instance of MSFC class with prior
fwd.fut.wpri <- msfc(tdate = as.Date("2013-05-13"),
                     include = powfutures130513$Include,
                     contract = powfutures130513$Contract,
                     sdate = powfutures130513$Start,
                     edate = powfutures130513$End,
                     f = powfutures130513$Closing,
                     prior = powpriors130513$mod.prior)

# instance of MSFC class without prior
fwd.fut.npri <- msfc(tdate = as.Date("2013-05-13"),
                     include = powfutures130513$Include,
                     contract = powfutures130513$Contract,
```

```
                    sdate = powfutures130513$Start,
                    edate = powfutures130513$End,
                    f = powfutures130513$Closing,
                    prior = 0)


# the generic plot() method
pw <- plot(fwd.fut.wpri, ylab = "EUR/MWh", legend = "")
pn <- plot(fwd.fut.npri, ylab = "EUR/MWh", legend = "")

# combine plots
gridExtra::grid.arrange(pw, pn)
```

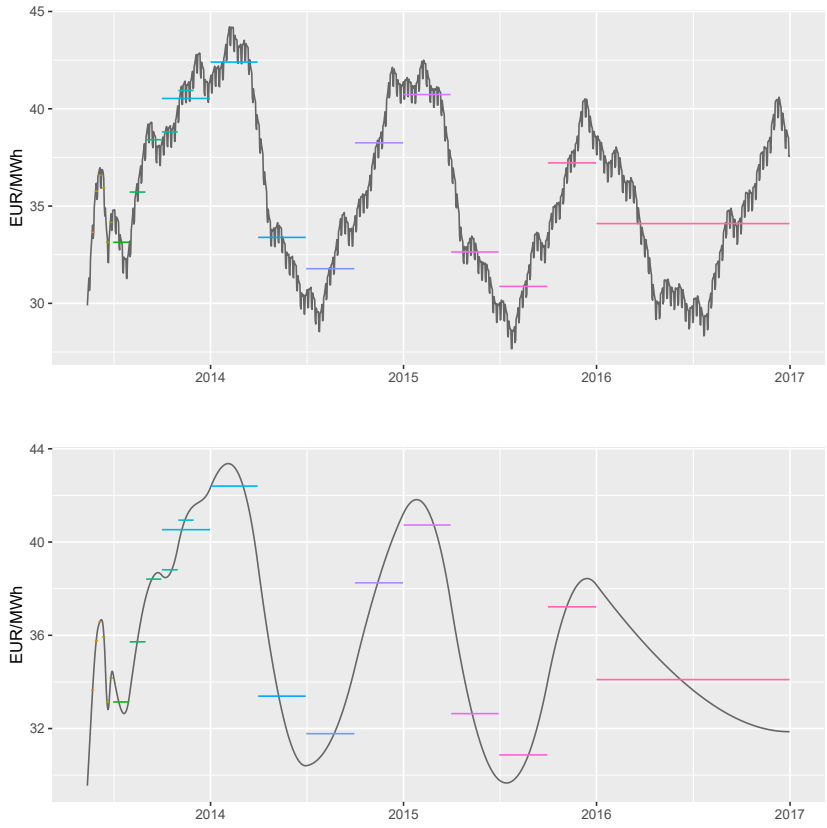

Figure 3: Two alternative forward curve calculations based on the same contract selection. In the top panel, a prior function is included in the calculation. This curve shows price variation on the weekly level, with lower prices during weekends. The prior also ensures that yearly seasonality is visible in the far end of the curve. The bottom plot is based solely on market prices, which does not reflect seasonality on such long horizon.

The computed prices may be verified via the `summary()` method, which also return a sample of the prior and information regarding the spline calculation:

```
> summary(fwd.fut.wpri)
$Description
[1] "MSFC of length 1329 built with 41 polynomials at trade date 2013-05-13"

$PriorFunc
[1] 30.10842 30.16396 30.19572 30.16144 29.06268 28.93272

$BenchSheet
   Include Contract       From         To Price  Comp
1     TRUE   W21-13 2013-05-20 2013-05-26 33.65 33.65
2     TRUE   W22-13 2013-05-27 2013-06-02 35.77 35.77
3     TRUE   W23-13 2013-06-03 2013-06-09 36.58 36.58
4     TRUE   W24-13 2013-06-10 2013-06-16 35.93 35.93
5     TRUE   W25-13 2013-06-17 2013-06-23 33.14 33.14
6     TRUE   W26-13 2013-06-24 2013-06-30 34.16 34.16
8     TRUE  MJUL-13 2013-07-01 2013-07-31 33.14 33.14
9     TRUE  MAUG-13 2013-08-01 2013-08-31 35.72 35.72
10    TRUE  MSEP-13 2013-09-01 2013-09-30 38.41 38.41
11    TRUE  MOCT-13 2013-10-01 2013-10-31 38.81 38.81
12    TRUE  MNOV-13 2013-11-01 2013-11-30 40.94 40.94
14    TRUE    Q4-13 2013-10-01 2013-12-31 40.53 40.53
15    TRUE    Q1-14 2014-01-01 2014-03-31 42.40 42.40
16    TRUE    Q2-14 2014-04-01 2014-06-30 33.39 33.39
17    TRUE    Q3-14 2014-07-01 2014-09-30 31.78 31.78
18    TRUE    Q4-14 2014-10-01 2014-12-31 38.25 38.25
19    TRUE    Q1-15 2015-01-01 2015-03-31 40.73 40.73
20    TRUE    Q2-15 2015-04-01 2015-06-30 32.64 32.64
21    TRUE    Q3-15 2015-07-01 2015-09-30 30.87 30.87
22    TRUE    Q4-15 2015-10-01 2015-12-31 37.22 37.22
25    TRUE   CAL-16 2016-01-01 2016-12-31 34.10 34.10
```

The forward curve values can be extracted along with daily prices for the contracts used in the calculation with the `show()` method:

```
> head(show(fwd.fut.wpri), 20)[1:5]
         Date     MSFC W21-13 W22-13 W23-13
1  2013-05-13 29.89373     NA     NA     NA
2  2013-05-14 30.40235     NA     NA     NA
3  2013-05-15 30.88704     NA     NA     NA
4  2013-05-16 31.30634     NA     NA     NA
5  2013-05-17 30.66200     NA     NA     NA
6  2013-05-18 30.98687     NA     NA     NA
7  2013-05-19 32.33591     NA     NA     NA
8  2013-05-20 32.74655  33.65     NA     NA
9  2013-05-21 33.19772  33.65     NA     NA
10 2013-05-22 33.63844  33.65     NA     NA
11 2013-05-23 34.02161  33.65     NA     NA
12 2013-05-24 33.34168  33.65     NA     NA
13 2013-05-25 33.62327  33.65     NA     NA
14 2013-05-26 34.91272  33.65     NA     NA
15 2013-05-27 35.24208     NA  35.77     NA
16 2013-05-28 35.59669     NA  35.77     NA
17 2013-05-29 35.92499     NA  35.77     NA
18 2013-05-30 36.17633     NA  35.77     NA
19 2013-05-31 35.34194     NA  35.77     NA
```

```
20 2013-06-01 35.44437     NA  35.77     NA
```

We have excluded columns from the data frame for the sake of presentation. Further details regarding the calculation such as spline coefficients and knot points can be found in the slots:

```
> slotNames(fwd.fut.wpri)
[1] "Name"        "TradeDate"   "BenchSheet"
[4] "Polynomials" "PriorFunc"   "Results"
[7] "SplineCoef"  "KnotPoints"  "CalcDat"
```

See for example the numeric vector with knot points, measured in years from the trading date:

```
> fwd.fut.npri@KnotPoints
 [1] 0.00000000 0.01917808 0.03561644 0.03835616 0.05479452 0.05753425 0.07397260
 [8] 0.07671233 0.09315068 0.09589041 0.11232877 0.11506849 0.13150685 0.13424658
[15] 0.21643836 0.21917808 0.30136986 0.30410959 0.38356164 0.38630137 0.46849315
[22] 0.47123288 0.55068493 0.63561644 0.63835616 0.88219178 0.88493151 1.13150685
[29] 1.13424658 1.38356164 1.38630137 1.63561644 1.63835616 1.88219178 1.88493151
[36] 2.13150685 2.13424658 2.38356164 2.38630137 2.63561644 2.63835616 3.63835616
```

The coefficients for the first polynomial in the adjustment function spline can be found with

```
> fwd.fut.npri@SplineCoef[[1]]
[1] -355585.14451   10911.10580     -78.47028     151.90713      29.54903
```

The most elaborate presentation of the curve calculation is available in `fwd.fut.npri@CalcDat`. This slot contains a data frame with all calculation details for each of the daily values returned by `msfc()`. It is not included here due to space requirements.

3 Energy price risk management

Energy market participants may be exposed to number of risk factors such as volume, profile and basis risk, counter party defaults, foreign exchange and market liquidity, to name a few. See for example Eydeland and Wolyniec (2002) and Kirschen and Strbac (2018) for a comprehensive treatment of the topic. The main focus in **etrm** is on the market price risk of the energy commodity. Consider the price risk associated with the constant base load volume q to be delivered over a future time interval (τ^s, τ^e). The risk can be mitigated by taking positions in the futures market during a trading period (t_0, T), which ends before the actual delivery of the energy takes place at $T < \tau^s$.

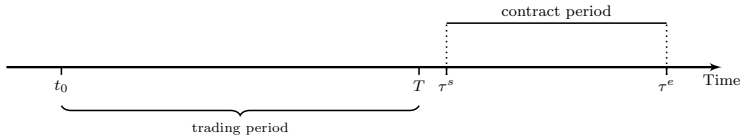

Figure 4: Trading and settlement periods for energy forward contract.

The price risk may be reduced by constructing a portfolio, consisting of the physical energy market exposure and derivatives contracts. The *portfolio price* per energy unit p_t is calculated as the weighted average of the value of the transacted volumes and the open volume evaluated mark-to-market

$$p_t = \frac{1}{q}\left(f_0 h_0 q + \sum_{i=1}^{t} f_i(h_i - h_{i-1})q + f_t(1-h_t)q\right) = f_0 h_0 + \sum_{i=1}^{t} f_i(h_i - h_{i-1}) + f_t(1-h_t) \qquad (5)$$

where $h_t \in (0,1)$ is the hedge rate and f_t the futures price at time t. In the simplest possible scenario, the risk can be managed by locking the entire volume in the forward market. This removes the price risk and the portfolio owner knows up front what to pay or receive when delivery of the energy takes place. On the downside, one might regret locking if the market develops in a favourable way.

Portfolio insurance strategies

In the portfolio insurance approach, dynamic hedging strategies that allow buying and selling the hedging instrument are used to protect the portfolio, while seeking to benefit from advantageous market developments. Historically, the theory of portfolio insurance has focused on protection against downside risk in financial investment portfolios, see Leland and Rubinstein (1976) , Perold and Sharpe (1988) and Leland (1980). Here, we apply the same principles to manage commodity price risk in the forward market. A consumer following a dynamic hedging program may control price risk by locking a share of future volume in the futures market, and increase (decrease) the share if the price increase (falls). A seller can implement similar strategies to maximise value of the energy portfolio. The size of the initial hedged share and how it is adjusted affects both the protective properties of the hedging scheme as well as its ability to exploit opportunities in the market. Trading activity needs

to be carefully managed and harmonized with overall objectives and risk preferences. A variety of portfolio insurance strategies offer different approaches to this task. The allocation strategies presented below all aim to control p_t and prevent breach of a pre-specified cap (or floor) price, p^*, under the hedge rate restriction $h_t \in (0,1)$.

The **Constant proportion portfolio insurance (CPPI)** strategy was introduced by Perold (1986) and Black and Jones (1987) for management of investment portfolios with capital guarantees. When applied to an energy portfolio, it aims to insure the portfolio by protecting a target price, a cap (floor) value for the portfolio price. Prior to start of hedging, p^* is set equal to the highest (lowest) acceptable portfolio price. This target price must be set higher than the first day's market price f_0 to implement cap protection, or lower for a floor protection model. The difference between the target price and the current portfolio price is termed the cushion. The key idea of CPPI is that the proportion of the volume exposed to the market should be calculated as a constant multiple m of the cushion. The multiple is given by $m = \mu^{-1}$, where $\mu > 0$ is a risk factor set to handle the maximum daily price change to be handled by the hedging model, which again affects strategy gearing. The unexposed proportion, the hedge rate, is

$$h_t = \begin{cases} 0 & \text{if} \quad c_t > \mu \\ 1 - mc_t & \text{if} \quad \mu \geq c_t \geq 0 \\ 1 & \text{if} \quad c_t < 0 \end{cases} \tag{6}$$

where $c_t = p^* - p_{t-1}$ for a short hedger, and $c_t = p_{t-1} - p^*$ for the long hedger.

In the **Dynamic proportion portfolio insurance (DPPI)** strategy, a decision rule similar to CPPI is applied, but the multiple m_t is allowed to vary. Changing market conditions may require re-evaluation of the risk factor in the multiple. Methods such as Value-at-Risk or Expexted Shortfall, or even simple heuristics can be used for this purpose. For a further treatment of DPPI type strategies, the reader is referred to Lee et al. (2008) and Chen et al. (2008). In **etrm** we also allow adjustments in the target price p_t^*, to catch opportunities to lower the capped value, or to increase the floor value. The hedge rate is determined similarly to CPPI, with

$$p_t^* = \begin{cases} min(\lambda p_{t-1}, p_{t-1}^*) & \text{short hedger} \\ \\ max(\lambda p_{t-1}, p_{t-1}^*) & \text{long hedger} \end{cases} \tag{7}$$

where $\lambda = \frac{p_o^*}{p_o}$ for a short hedger, and $\lambda = \frac{p_o}{p_o^*}$ for the long hedger.

Option based portfolio insurance (OBPI) was first introduced in Leland and Rubinstein (1976) as a means of providing insurance for investment portfolios. By combining an investment in a risky asset with a put option on the asset, the portfolio value is prevented from falling below the option strike price, K. A similar approach can be taken for the energy portfolio. As we are using futures contracts to manage the energy price risk, the Black-76 formula introduced in Black (1976) is used to approximate the contingent claim premium. The price at time t of the European call and put options with exercise date T and strike price K, on a futures contract with delivery start $\tau^s \geq T$ is given by

$$C(f_t, t, K, \sigma, r) = e^{-r(T-t)}[f_t N(d_1) - KN(d_2)] \tag{8}$$

$$P(f_t, t, K, \sigma, r) = e^{-r(T-t)}[KN(-d_2) - f_t N(-d_1)] \tag{9}$$

where f_t is the futures price and N is the cumulative distribution function of $N(0,1)$, where

$$d_1 = \frac{\ln(f_t/K) + (\sigma^2/2)(T-t)}{\sigma\sqrt{(T-t)}} \tag{10}$$

$$d_2 = d_1 - \sigma\sqrt{(T-t)} \tag{11}$$

and r is the risk free rate of interest, σ the volatility of the underlying futures price and $(T-t)$ is the time to exercise date. The sensitivity in the option premiums with respect to changes in the underlying futures price is given by the call and put option deltas:

$$\frac{\partial C}{\partial f} = e^{-r(T-t)}N(d_1) \tag{12}$$

$$\frac{\partial P}{\partial f} = e^{-r(T-t)}N(-d_1) \tag{13}$$

These are used to synthesise the option, by setting portfolio hedge rate $h_t \in (0,1)$ with the call (buyer) and put (seller) option deltas. By implementing this *delta hedging* scheme, a cap (floor) for the portfolio price is set at the option strike price K, adjusted for the option premium/ replication costs. For a more

detailed presentation of the underlying theory, the reader is referred to Bjork (2009).

Step hedge portfolio insurance (SHPI) is a simple and mechanical benchmark strategy that builds hedging positions gradually by transacting identical volumes each day through the trading period (t_0, T), reaching a full hedge prior to the start of the settlement period. The hedge rate for a buyer at time t is given by

$$h_t = \begin{cases} \frac{t}{T-t_0+1} & \text{if} \quad p_t < p^* \\ 1 & \text{if} \quad p_t \geq p^* \end{cases} \tag{14}$$

The hedges for a seller is entered mechanically in a similar manner as long as $p_t > p^*$. In the event $p_t \leq p^*$ a full hedge $h_t = 1$ is implemented. By distributing the transacted volumes evenly across the trading period while monitoring the target, the strategy portfolio price will either be locked in at the target price, or end up equal to the average forward market price over (t_0, T).

Finally, the **Stop loss portfolio insurance (SLPI)** is another simple benchmark, where no hedge positions are entered unless the target level is reached. For a buyer, this may be expressed as

$$h_t = \begin{cases} 0 & \text{if} \quad p_t < p^* \\ 1 & \text{if} \quad p_t \geq p^* \end{cases} \tag{15}$$

For a seller, the logic is reversed with $h_t = 0$ for $p_t > p^*$ and $h_t = 1$ for $p_t \leq p^*$. In the event that the target level is reached, the portfolio is kept fully hedged until start of settlement. If this does not occur, the portfolio follows the forward market, leaving an option to lock in the price at contract expiration.

The strategies presented above all have strengths and weaknesses. CPPI, SHPI and SLPI are simple and intuitive, but can be vulnerable to so-called *lock in*, the inability to improve portfolio price once the target level has been reached. This is also the case for DPPI. The OBPI does not suffer from this trait, but it relies on more assumptions regarding model parameters. In some scenarios, it will also generate more trading activity (costs), for example if the market fluctuate around the option strike price, K. Finally, as the strategies must be implemented in discrete time, they will all be exposed to gap risk.

The strategy classes with examples

The portfolio insurance strategies in **etrm** are implemented as S4 classes. Since they share many characteristics, they inherit most of their properties from a parent class, `"GenericStrat"`. In fact, the implementation of the simple benchmark strategies SLPI and SHPI do not require any additional properties to be added to the parent. The remaining strategy classes have some additional model specific features, in accordance with the descriptions in previous section. This modular design offers flexibility, and new strategies for price risk management can easily be added to the package.

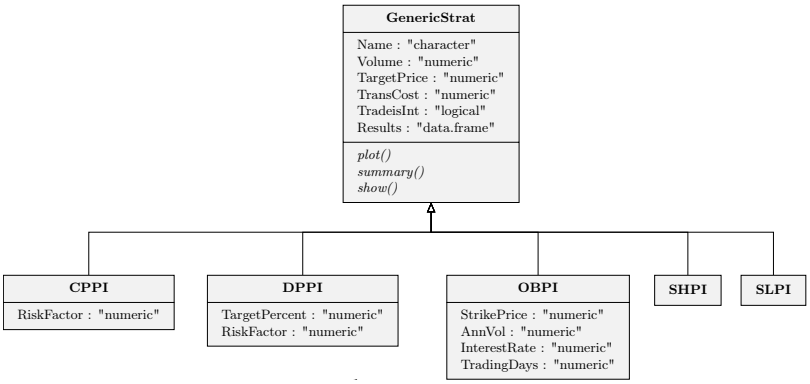

Figure 5: Attributes and methods for portfolio insurance strategy classes in the **etrm** package.

Figure 5 provide an overview of the class hierarchy, and a brief description is given in the following. In `"GenericStrat"`, the `"Name"` property is used to store a strategy identifier. Allowed character values are "CPPI", "DPPI", "OBPI", "SHPI" and "SLPI". The volume to be managed and the corresponding price cap (floor) can be found in `"Volume"` and `"TargetPrice"`, respectively. If a transaction cost has been included in the calculation of the portfolio price, this is to be found in `"TransCost"`. One may also set a restriction on transactions by requiring that the smallest volume available for trading is

equal to 1 unit. This lot size limitation is stored as TRUE/FALSE in `"TradeisInt"`. The main output from a strategy calculation can be found in `"Results"`. This data frame keeps daily values for market prices, transactions, exposed volume, open volume, hedge rate, target price and portfolio price.

The generic methods `plot()`, `summary()` and `show()` are implemented in `"GenericStrat"` and inherited by the strategy classes. The `plot()` method returns a chart based on `"Results"`, with daily values for portfolio, market and target prices and portfolio hedge rate. The `summary()` method returns a list with five elements; a description string, portfolio volume, target price, calculated churn rate and a data frame with summary statistics for the trading period. Finally, the `show()` method returns the `"Results"` data frame.

Argument	Description	Default value
`q`	Numeric volume	none
`tdate`	Date vector with trading days	none
`f`	Numeric price vector	none
`tcost`	Numeric transaction cost	0
`int`	Logical lot size integer restriction	TRUE

Table 3: Arguments shared by the portfolio insurance strategy functions.

The strategy constructor functions `cppi()`, `dppi()`, `obpi()`, `shpi()` and `slpi()` share five of the arguments, see Table 3. Each strategy require some additional arguments to implement the models presented in previous section. All of these inputs are of `"numeric"` data type. They are summarised in Table 4.

Function	Argument	Description	Default value
`cppi()`	`tper`	Target price factor	none
	`rper`	Risk factor percentage	none
`dppi()`	`tper`	Target price factor	none
	`rper`	Risk factor percentage	none
`obpi()`	`k`	Option strike price	$k = f_0$
	`vol`	Annualized volatility	none
	`r`	Interest rate	0
	`tdays`	Trading days per year	250
	`daysleft`	Days left to expiry	none
`shpi()`	`tper`	Target price factor	none
	`daysleft`	Days left to expiry	none
`slpi()`	`tper`	Target price factor	none

Table 4: Model specific arguments for the portfolio insurance strategy functions.

To illustrate further, we proceed with an example. Consider a European consumer of electricity procuring 30 MW to be delivered in 2006. The CAL-06 baseload power future from the synthetic **etrm** `"powcal"` data set is used as hedging instrument. Trading is started 500 days prior to the contract expiry, approximately a horizon of 2 years. For the `"OBPI"` strategy presented below, the target price is calculated as an expected price cap given by the option premium-adjusted strike price selected for the delta hedging scheme within a standard Black-76 option pricing framework. The default `obpi()` strike price is set at-the-money, in this case at 26.82 EUR/MWh. The expected target price illustrated with the horizontal dotted line in Figure 6 is calculated to be 29.84 EUR/MWh. The `"OBPI"` delta hedging scheme dictate an initial hedge rate of 57 percent, and the consumer enters a 17 MW position in CAL-06 on the first day of trading.

As time progresses and the market price changes, the `obpi()` function adjust the required hedge rate in order to replicate the call option on the CAL-06 contract. Hedge rate is gradually built up as the market increase from the second quarter of 2004, followed by a reduction after the sharp price drop starting late in the same year. Eventually, the volume is fully hedged due to the strong upwards price trend in 2005. The CAL-06 contract closes at 37.81 EUR/MWh on the expiry date, while the consumer has a hedge of 30 MW and a portfolio price of 29.29 EUR/MWh. The calculated price of the option to be synthesized (and the delta hedges) will depend on the Black-76 model parameters. In this example the risk free rate of interest is set to $r = 0$ and annualized volatility σ is assumed to be 20 percent. The following code will implement the strategy and create the plot in Figure 6:

```
# data frame with final 500 trading days for CAL-06 contract
dat06 <- tail(na.omit(powcal[, c(1,2)]), 500)

# instance of the OBPI class
cal06_obpi_b <- obpi(q = 30,
                     tdate = dat06$Date,
                     f = dat06$`CAL-06`,
                     k = dat06$`CAL-06`[1],
                     vol = 0.2,
                     r = 0,
                     tdays = 250,
                     daysleft = 500,
                     tcost = 0,
                     int = TRUE)

# the generic plot() method
plot(cal06_obpi_b, title = "", legend = "right", ylab.1 = "EUR/MWh")
```

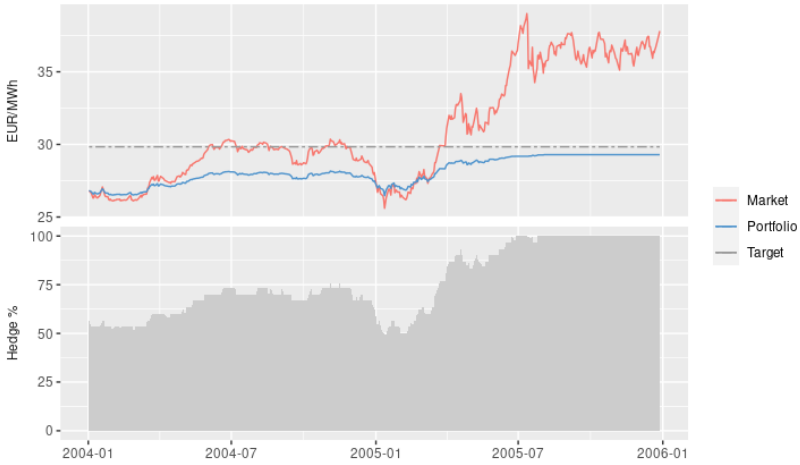

Figure 6: Option based portfolio insurance (OBPI) strategy for buyer CAL-06. Daily observations for prices (top panel) and hedge rate (bottom panel). As the market price continue to rise, the hedge rate is increased and the portfolio price is locked below the target price level.

An aggregated view of the trading activity over the 2 year period and final results can be retrieved by running the `summary()` method on the object created above:

```
> summary(cal06_obpi_b)
$Description
[1] "Hedging strategy of type OBPI and length 500"

$Volume
[1] 30

$Target
[1] 29.83626

$ChurnRate
[1] 4.333333
```

```
$Stats
      Market Trade Exposed Position     Hedge   Target Portfolio
First  26.82    17      13       17 0.5666667 29.83626  26.82000
Max    39.01    17      17       30 1.0000000 29.83626  29.29433
Min    25.60    -3       0       13 0.4333333 29.83626  26.46833
Last   37.81     0       0       30 1.0000000 29.83626  29.29433
```

We note from the `"ChurnRate"` that the underlying 30 MW volume had to be traded 4.33 times in order to synthesize the call option and achieve the results summarised in `"Stats"`. By also considering the trading costs in the calculations, the user can get valuable inputs when considering alternatives, such as simply buying the option in the market. However, such contract may not always be available.

Finally, the `show()` method provide details regarding daily values for market price, transactions, exposed volume, futures contract position, hedge rate, the target price and the calculated portfolio price:

```
> head(show(cal06_obpi_b))
        Date Market Trade Exposed Position     Hedge   Target Portfolio
1 2004-01-02  26.82    17      13       17 0.5666667 29.83626  26.82000
2 2004-01-05  26.63    -1      14       16 0.5333333 29.83626  26.73767
3 2004-01-07  26.31     0      14       16 0.5333333 29.83626  26.58833
4 2004-01-08  26.31     0      14       16 0.5333333 29.83626  26.58833
5 2004-01-09  26.54     0      14       16 0.5333333 29.83626  26.69567
6 2004-01-12  26.32     0      14       16 0.5333333 29.83626  26.59300
```

For the sake of comparison, the OBPI strategy for CAL-06 from a sellers point of view can be implemented with similar assumptions by setting the volume to $q = -30$. Using the default at-the-money strike price, the seller calculates an expected target floor to protect at 23.80 EUR/MWh and an initial hedge rate of 43 percent. As the market starts to rise, the hedge is reduced. The seller increases the hedge in late 2004 to dampen the effect from the market drop, and finally exits the forward market positions as the price increases during 2005. The portfolio price follows the market upwards with a premium for the put option replication, as expected for an insurance scheme. The CAL-06 contract closes at 37.81 EUR/MWh, and the seller has a portfolio price of 35.34 EUR/MWh, which may be locked in on the final trading day.

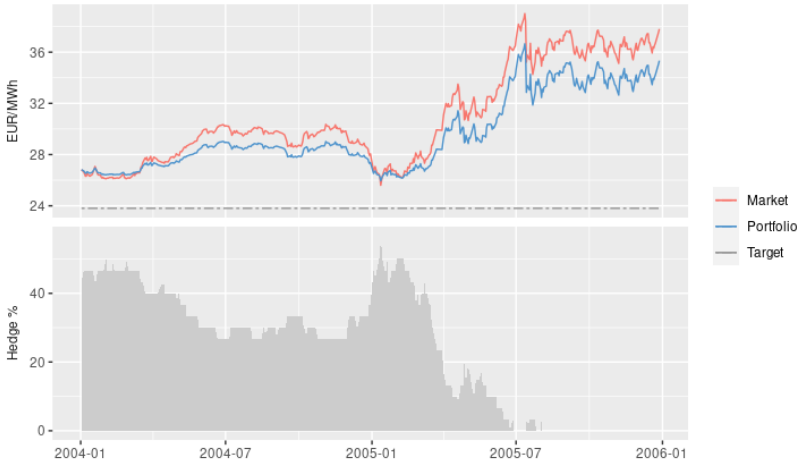

Figure 7: Option based portfolio insurance (OBPI) strategy for seller CAL-06. Daily observations for prices (top panel) and hedge rate (bottom panel). The hedge rate is lowered in the upwards trending market, and the portfolio price continue to increase.

```
# instance of the OBPI class
cal06_obpi_s <- obpi(q = - 30,
                     tdate = dat06$Date,
                     f = dat06$`CAL-06`,
                     k = dat06$`CAL-06`[1],
                     vol = 0.2,
                     r = 0,
                     tdays = 250,
                     daysleft = 500,
                     tcost = 0,
                     int = TRUE)

# the generic plot() method
plot(cal06_obpi_s, title = "", legend = "right", ylab.1 = "EUR/MWh")

> summary(cal06_obpi_s)
$Description
[1] "Hedging strategy of type OBPI and length 500"

$Volume
[1] -30

$Target
[1] 23.80374

$ChurnRate
[1] 4.2

$Stats
      Market Trade Exposed Position     Hedge   Target Portfolio
First  26.82   -13     -17      -13 0.4333333 23.80374  26.82000
Max    39.01     2     -13        0 0.5666667 23.80374  36.64867
Min    25.60   -13     -30      -17 0.0000000 23.80374  25.95167
Last   37.81     0     -30        0 0.0000000 23.80374  35.33567
```

In the examples above, we have implicitly assumed that both the consumer and the seller have a flat volume corresponding to 30 MW over the entire year which can be covered by a base load contract such as the CAL-06. In practice, this is typically not the case. Industrial energy consumers will have consumption profiles determined by the activity level in their production facilities, and often face seasonal shifts due to variation in demand, or holidays. Weather also play a large role, both for consumers and producers such as hydroelectric plants. In order to hedge the predicted volume more precisely, some of the other contract types presented in Table 2 will need to be included in the portfolio. Market players will "roll forward" and start trading contracts covering shorter periods such as quarters, months and weeks, as they become available. The mandate for the energy portfolio will typically be broken down into smaller time intervals with expected volume and required hedge levels. All strategies presented here may be used to make decisions for several years and their sub periods, and the market value of a specific volume prognosis and corresponding futures positions can be evaluated using the forward curve discussed in previous sections.

In order to maintain focus on the strategies themselves, we continue with the baseload example with 30 MW. In Figure 8 we plot results for the remaining four strategies for the consumer hedging with CAL-06. The benchmark strategies `"SHPI"`, and `"SLPI"` follow simple, mechanical patterns. The `"SHPI"` builds a full hedge gradually over the trading period, ending at either the average forward market price for the period, or the target price. This approach will always ensure a full hedge at expiry, without intervention. The `"SLPI"` does not take any positions unless the target is reached, ending either at the target level, or leaving an option to close at the contract expiry price. As the CAL-06 increase significantly during 2005, both end up at the target level.

The `"CPPI"`, and `"DPPI"` strategies are more dynamic and adjust hedge rate according to market developments and the model parameters. As the `"DPPI"` implements a dynamic risk factor, μ_t, the strategies are geared differently. In this example, the `"DPPI"` successfully adjusts the target price downward on one of the first trading days, and achieves a lower portfolio price on last trading day.

A similar overview from a seller's perspective is provided in Figure 9. As the market trends upwards during the hedging period, none of the strategies end up at the initial target price. The `"SHPI"` builds the hedge positions in a step-wise manner, ending up with a portfolio price equal to

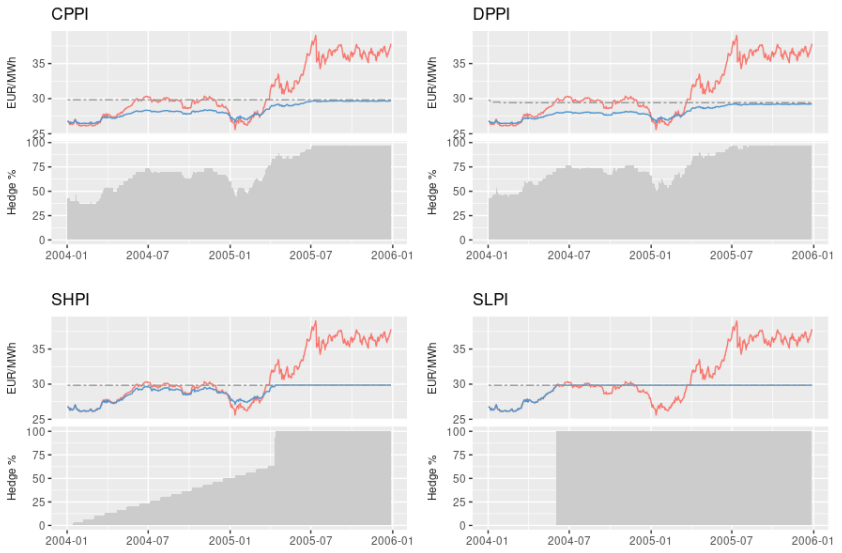

Figure 8: Achieved results for the strategies CPPI, DPPI, SHPI and SLPI for buyer CAL-06. Daily observations for prices (top panels) and hedge rate (bottom panels).

the average futures market price for the period. The `"SLPI"` does not enter any positions, leaving an option to lock in market price at expiry. Finally, we can also here see some differences between `"CPPI"`, and `"DPPI"`. This is due to the dissimilar gearing of the portfolios, but also because of the rather frequent adjustments of the target price by `"DPPI"`. In order to protect the higher targets, hedge rate must be increased and `"DPPI"` falls behind `"CPPI"` and ends up at a lower portfolio price for the seller.

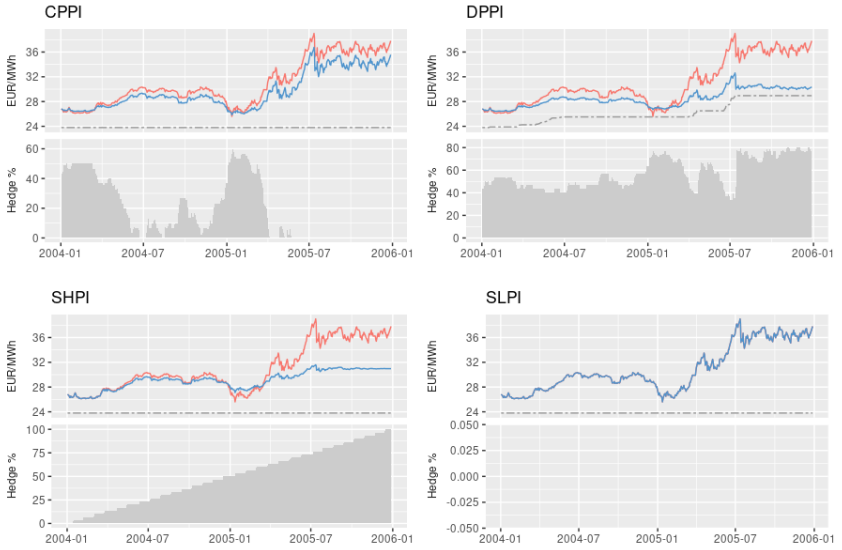

Figure 9: Achieved results for the strategies CPPI, DPPI, SHPI and SLPI for seller CAL-06. Daily observations for prices (top panels) and hedge rate (bottom panels).

etrm can also be used in conjunction with other R packages to evaluate risks related to energy procurement. Metrics such as *Value-at-Risk* and *Expected Shortfall* can for example be calculated using the **PerformanceAnalytics** package. We will proceed with a simple, illustrative example. Consider the OBPI portfolios `"cal06_obpi_b"` and `"cal06_obpi_s"` in the code example above. If we need to calculate risk measures at a specific point in time, say at day 350 in the trading period, we can execute the following code:

```
library(PerformanceAnalytics)

# CAL-06 returns prior to t=350
ret_06 <- head(diff(log(show(cal06_obpi_b)$Market)), 349)

# portfolio status at t=350
pdat <- rbind(
  Buyer =show(cal06_obpi_b)[350,],
  Seller =show(cal06_obpi_s)[350,]
)

# add risk measures to pdat
pdat <-cbind(pdat,
    VaR = abs(rep(VaR(ret_06, p=.95, method="historical"), 2)*pdat$Market*pdat$Exposed*8760),
    ES = abs(rep(ES(ret_06, p=.95, method="historical"), 2)*pdat$Market*pdat$Exposed*8760))
```

The calculation above evaluate market risk related to the unhedged volume (exposed MW $\times$8760 hours in the year 2006) at current market prices under the (simplistic) assumption of symmetry in the returns distribution. The portfolio status, including risk metrics is

```
> pdat[c(-1, -3)]
       Market Exposed Position Hedge   Target Portfolio VaR         ES
Buyer   32.54       3       27   0.9    29.84  28.98     10671.55    18237.06
Seller  32.54     -27       -3   0.1    23.81  30.38    96043.94    164133.52
```

4 Overview of the etrm package

Package **etrm** offers an open source implementation of core functionalities of an ETRM system:

- Construction of forward curves
- Strategies for price risk management

Functions included in the package are listed in Table 5.

Function	Description
`msfc()`	Maximum Smoothness Forward Curve
`cppi()`	Constant Proportion Portfolio Insurance
`dppi()`	Dynamic Proportion Portfolio Insurance
`obpi()`	Option Based Portfolio Insurance
`shpi()`	Step Hedge Portfolio Insurance
`slpi()`	Stop Loss Portfolio Insurance

Table 5: Overview of **etrm** package functions

All functions act as constructors for their corresponding S4 classes, as described in further detail in previous sections. The classes all have generic methods `plot()`, `summary()` and `show()`. Unit tests covering all functions in **etrm** have been implemented using the **testthat** framework introduced in Wickham (2011b).

Three synthetic data sets are included in the package, see Table 6. The `"powfutures130513"` and `"powpriors130513"` data may be used to create forward curves with the `msfc()` function for the trading date 2013-05-13. The portfolio insurance strategies may be tested on the `"powcal"` data set, which contains historical prices for 11 base load power futures.

Data set	Description
`powfutures130513`	Synthetic data for a set of electricity base load futures quoted at 2013-05-13. Closing prices for contracts with weekly, monthly, quarterly and yearly settlement periods
`powpriors130513`	Two simple priors for forward market price curve Daily values for calculation to be used with powfutures130513
`powcal`	Synthetic data set with daily closing prices for 11 electricity base load futures with yearly settlement periods for 2006-2016

Table 6: Overview of **etrm** package data sets

5 Summary and suggestions for future work

This paper introduces **etrm**, an R package for energy market risk management. The package contains tools previously not available in the R ecosystem, such as the `msfc()` function for building a forward curve for energy commodities with flow delivery contracts and strong seasonality. The forward curve is a key decision making tool with many uses, such as pricing non-standard supply agreements, investment decisions and risk management. **etrm** also provides implementations of portfolio insurance strategies for handling price risk, suitable for both long and short hedgers. The functions can be used for back testing strategies on historical futures price data, risk and strategy evaluations, and as decision support tools for trade execution.

The **etrm** package may be developed further by incorporating new elements. First, the forward curve calculation may be done on an hourly level. The bid-ask spread can be used as price constraint for the optimization, as an extension of the current solution based on closing prices. Competing forward curve calculation methods can also be added to the package, and new asset allocation strategies for price risk management could be included.

A further extension of **etrm** functionality can be to implement a `"PORTFOLIO"` class, consisting of a daily volume prognosis covering the full management horizon, supplemented with authorized volumes per (sub)period and hedging strategy objects implemented in accordance with these authorizations. The portfolio object could contain multiple strategy objects for contracts such as "year", "quarter", "month" and "week", depending on the shape of the volume prognosis. This construction can be priced using the forward curve, and portfolio wide risk measures could be calculated via Monte Carlo simulations on the curve.

6 Acknowledgments

We thank the editor and two anonymous reviewers for their constructive comments, which helped us to improve the manuscript. This work has been supported by The Norwegian Research Council via the industrial Ph.D. scheme. The author also wish to thank Montel AS and former colleagues in Kinect Energy Markets AS for interesting discussions and helpful comments.

Bibliography

K. J. Adams and D. R. Van Deventer. Fitting yield curves and forward rate curves with maximum smoothness. *Journal of Fixed Income*, pages 52–62, 1994. URL https://doi.org/10.3905/jfi.1994.408102. [p321]

N. Anderson, F. Breedon, M. Deacon, A. Derry, and G. Murphy. *Estimating and Interpreting the Yield Curve*, volume 6. John Wiley & Sons Incorporated, 1996. [p321]

M. Benini, M. Marracci, P. Pelacchi, and A. Venturini. Day-ahead market price volatility analysis in deregulated electricity markets. In *IEEE Power Engineering Society Summer Meeting*, volume 3, pages 1354–1359 vol.3, 2002. URL https://doi.org/10.1109/PESS.2002.1043596. [p320]

F. E. Benth and M. Schmeck. Pricing and hedging options in energy markets using black-76. *Journal of Energy Markets*, 7(2), 2014. URL https://doi.org/10.21314/JEM.2014.114. [p320]

F. E. Benth, J. Kallsen, and T. Meyer-Brandis. A non-gaussian ornstein–uhlenbeck process for electricity spot price modeling and derivatives pricing. *Applied Mathematical Finance*, 14(2):153–169, 2007a. URL https://doi.org/10.1080/13504860600725031. [p320]

F. E. Benth, S. Koekkebakker, and F. Ollmar. Extracting and applying smooth forward curves from average-based commodity contracts with seasonal variation. *The Journal of Derivatives*, 15(1):52–66, 2007b. URL https://doi.org/10.3905/jod.2007.694791. [p320, 321]

F. E. Benth, J. S. Benth, and S. Koekebakker. *Stochastic Modelling of Electricity and Related Markets*, volume 11. World Scientific, 2008. URL https://doi.org/10.1142/6811. [p320, 323]

H. Bessembinder and M. L. Lemmon. Equilibrium pricing and optimal hedging in electricity forward markets. *The Journal of Finance*, 57(3):1347–1382, 2002. URL https://doi.org/10.1111/1540-6261.00463. [p320]

T. Bjork. *Arbitrage Theory in Continuous Time*. Oxford University Press, 3 edition, 2009. URL https://EconPapers.repec.org/RePEc:oxp:obooks:9780199574742. [p331]

F. Black. The pricing of commodity contracts. *Journal of Financial Economics*, 3(1):167–179, 1976. URL https://doi.org/10.1016/0304-405X(76)90024-6. [p330]

F. Black and R. W. Jones. Simplifying portfolio insurance. *The Journal of Portfolio Management*, 14(1): 48–51, 1987. URL https://doi.org/10.3905/jpm.1987.409131. [p330]

S. Borak and R. Weron. *A Semiparametric Factor Model for Electricity Forward Curve Dynamics*. Humboldt-Universität zu Berlin, Wirtschaftswissenschaftliche Fakultät, 2008. URL http://dx.doi.org/10.18452/4142. [p321]

M. Burger, B. Klar, A. Müller, and G. Schindlmayr. A spot market model for pricing derivatives in electricity markets. *Quantitative Finance*, 4:109–122, 2004. URL https://doi.org/10.1088/1469-7688/4/1/010. [p320]

J.-S. Chen, C.-L. Chang, J.-L. Hou, and Y.-T. Lin. Dynamic proportion portfolio insurance using genetic programming with principal component analysis. *Expert Systems with Applications*, 35(1):273–278, 2008. URL https://doi.org/10.1016/j.eswa.2007.06.030. [p330]

A. Deoras. Electricity load and price forecasting webinar case study, 2021. https://www.mathworks.com/matlabcentral/fileexchange/28684-electricity-load-and-price-forecasting-webinar-case-study [Accessed: 2021-04-06]. [p320]

A. Eydeland and K. Wolyniec. *Energy and Power Risk Management: New Developments in Modeling, Pricing, and Hedging*, volume 97. John Wiley & Sons, 2002. [p320, 321, 329]

S. Farrington. Energy risk software rankings: A different world, 2020. https://www.risk.net/commodities/energy/7511636/energy-risk-software-rankings-a-different-world [Accessed: 2021-04-05]. [p320]

S.-E. Fleten and J. Lemming. Constructing forward price curves in electricity markets. *Energy Economics*, 25(5):409–424, 2003. URL https://doi.org/10.1016/S0140-9883(03)00039-2. [p321]

M. Hildmann, E. Kaffe, Y. He, and G. Andersson. Combined estimation and prediction of the hourly price forward curve. In *2012 IEEE Power and Energy Society General Meeting*, pages 1–8, 2012. URL 10.1109/PESGM.2012.6345333. [p321]

J. Janczura, S. Trück, R. Weron, and R. C. Wolff. Identifying spikes and seasonal components in electricity spot price data: A guide to robust modeling. *Energy Economics*, 38:96–110, 2013. [p320]

D. S. Kirschen and G. Strbac. *Fundamentals of Power System Economics*. John Wiley & Sons, 2018. [p320, 321, 329]

H.-I. Lee, M.-H. Chiang, and H. Hsu. A new choice of dynamic asset management: the variable proportion portfolio insurance. *Applied Economics*, 40(16):2135–2146, 2008. URL https://doi.org/10.1080/00036840600949280. [p330]

H. Leland and M. Rubinstein. The evolution of portfolio insurance. *Dynamic Hedging*, 01 1976. URL https://www.researchgate.net/publication/265430746_The_Evolution_of_Portfolio_Insurance. [p320, 329, 330]

H. E. Leland. Who should buy portfolio insurance? *The Journal of Finance*, 35(2):581–594, 1980. URL http://www.jstor.org/stable/2327419. [p320, 329]

K. G. Lim and Q. Xiao. Computing maximum smoothness forward rate curves. *Statistics and Computing*, 12(3):275–279, 2002. URL https://doi.org/10.1023/A:1020707028156. [p322]

J. H. McCulloch. Measuring the term structure of interest rates. *The Journal of Business*, 44(1):19–31, 1971. URL http://www.jstor.org/stable/2351832. [p321]

M. Nicolosi. Wind power integration and power system flexibility–an empirical analysis of extreme events in germany under the new negative price regime. *Energy Policy*, 38(11):7257–7268, 2010. ISSN 0301-4215. URL https://www.sciencedirect.com/science/article/pii/S0301421510005860. Energy Efficiency Policies and Strategies with regular papers. [p320]

F. Ollmar. An analysis of derivative prices in the Nordic power market. 2003. URL http://hdl.handle.net/11250/164248. [p320, 321, 322, 323]

A. Perold. Constant proportion portfolio insurance. Harvard business school. *Unpublished manuscript*, 1986. [p330]

A. F. Perold and W. F. Sharpe. Dynamic strategies for asset allocation. *Financial Analysts Journal*, 44(1): 16–27, 1988. URL http://www.jstor.org/stable/4479087. [p320, 329]

S. Sundar. Energy trading & risk management with Matlab webinar case study, 2021. https://www.mathworks.com/matlabcentral/fileexchange/28056-energy-trading-risk-management-with-matlab-webinar-case-study [Accessed: 2021-04-06]. [p320]

H. Wickham. ggplot2. *WIREs Computational Statistics*, 3(2):180–185, 2011a. URL https://onlinelibrary.wiley.com/doi/abs/10.1002/wics.147. [p324]

H. Wickham. testthat: Get Started with Testing. *The R Journal*, 3(1):5–10, 2011b. URL https://doi.org/10.32614/RJ-2011-002. [p337]

Anders D. Sleire
Department of Mathematics
University of Bergen
P.O.Box 7803
N-5020 Bergen
Norway
Anders.Sleire@uib.no

Paper C

7.3 Modelling and Pricing Air Pollution Derivatives

Anders D. Sleire, working paper.

Modelling and Pricing Air Pollution Derivatives

Anders D. Sleire
Department of Mathematics, University of Bergen, Norway

March 3, 2023

Abstract

Severe air pollution in densely populated urban areas represent a significant health threat to millions of people. The alarming cases are often found in some of the world's most populous cities, where incidents with extreme pollution levels directly affect a significant proportion of the global population. In recent years, there has been increased focus on the financial impact of air pollution, including consequences for business and industry. The pollution level is typically assessed by measuring concentrations of the so-called criteria air pollutants; ground-level ozone, particulate matter, carbon monoxide, sulfur dioxide, and nitrogen dioxide. Results are reported to the public on a standardized scale, such as the widely-adopted Air Quality Index introduced by the US Environmental Protection Agency (US EPA AQI). Building upon weather derivatives theory, we design contracts whose payoff depend on publicly available AQI data. We develop stochastic models for the pollution dynamics for three major cities in China, and show how these can be used to price derivatives contracts written on AQI-based indices. Results are compared to valuations from alternative pricing strategies, and a practical use case is presented and discussed.

Keywords: Air pollution derivatives; weather derivatives; option pricing; stochastic modelling; financial risk management; air quality index.

JEL Classification: G13, G32, Q53.

1 Introduction

Rapid development and urbanization in fast-growing developing economies of the world have led to increased levels of air pollution, which is becoming a growing concern. The soaring demand for energy, transportation, and industrial activities associated with growth have resulted in increased emissions and higher levels of air pollution in many of the world's most populous cities

[Chan and Yao, 2008], [Mage et al., 1996]. This has become a significant concern for public health and the environment. In many Asian cities, air pollution levels are among the highest in the world, and the health effects of pollution, such as respiratory and cardiovascular diseases, are becoming increasingly prevalent [Manisalidis et al., 2020], [Gurjar et al., 2010]. In regions heavily impacted, the consequences are far-reaching and can include reduced life expectancy, decreased worker productivity, and negative impacts on industries such as agriculture, manufacturing, construction and tourism [Marlier et al., 2016]. Governments aim to play an active role in mitigating the risks through the implementation of regulations and programs for sustainable urban development. Other initiatives include creating incentives for clean energy production and investment in public transportation infrastructure. Many countries have adopted national air quality standards, which are used to regulate and monitor the levels of pollutants in the air. These standards typically specify limits for a range of pollutants, including ozone, particulate matter, carbon monoxide, nitrogen dioxide, and sulfur dioxide [Kuklinska et al., 2015], [Marlier et al., 2016]. The implementation of these standards is supported by air quality monitoring programs, which are designed to measure the levels of pollutants in the air and provide information on their impacts. This information is used by governments to assess the effectiveness of air quality management strategies and to identify areas where additional measures may be needed.

Megacities such as Beijing and Delhi are facing significant air pollution challenges. The emissions generated by their rapidly growing populations and economies have contributed to multiple episodes of extreme air pollution, which have required immediate action from the authorities. While the generally poor air quality pose a long-term threat to public health, these incidents represent an acute risk. As a result, governments have been forced to take drastic measures, such as implementing lockdowns, to reduce emissions and improve air quality [An et al., 2007], [Kumar et al., 2015]. Pollution incidents and lockdowns can have a significant impact on the economy, hindering business operations and causing economic disruption in regions that play an important role in the global supply chain. Both the extraordinary circumstances during a pollution incident, and the imposed shut down of activities may represent a significant financial risk for business owners.

While insurance agreements can to some extent be used to reduce this exposure, they may not cover all relevant outcomes and be costly and time-consuming to implement. In an insurance settlement process, the policyholder would have to file a claim and provide documentation of loss, which needs to be evaluated by the provider of the insurance. As an alternative, we suggest using standardized financial derivatives contracts as tools for managing financial pollution risk. Similar instruments have already been used for decades in the market for weather derivatives Thind [2014]. The traded contracts are settled against standardized indices that are built on weather data reported by trusted third parties, typically meteorological organizations, gov-

ernment agencies, and private weather companies. Depending on the outcome of the measured weather event, settlement is conducted without delay, according to contract specifications. This provides increased transparency and flexibility, as there is no need for further documentation, and the derivative may be freely bought and sold in the market.

The remainder of the article is organized as follows. In Section 2 we give an introduction to weather markets, and consider how weather derivative-like contracts may be suited to manage pollution risk. Thereafter, an overview of previous studies is presented. In Section 3 we give an introduction to the US Environmental Protection Agency AQI [1] standard, and show how publicly available AQI data can be used to define indices for the derivatives contracts. We continue with a discussion of alternative approaches to pricing, and present a stochastic model for the AQI dynamics. Section 4 contains an empirical analysis of pollution data from three large cities in China, and stochastic models for the reported AQI series. In section 5 we compare results from alternative pricing methods and explore some practical applications. Main conclusions and suggestions for further work are presented in Section 6.

2 Motivation and review of literature

Weather derivatives are financial instruments that allow market participants to hedge against the impacts of adverse weather conditions. They are similar to other financial derivatives, in that their value is derived from an underlying settlement reference, which in this case is weather. The Chicago Mercantile Exchange (CME) offers trade in futures and options based on several weather indices, including temperature, snow fall and precipitation, and the listed contracts are written for locations mainly in North America, Australia, Japan, and Europe. A significant proportion of the global weather market is managed OTC, as this provides more flexibility [WRMA, 2011]. By trading bilaterally, the parties are able to specify the geographical location and tailor contract terms to meet their risk management needs. The first recorded weather trade was executed OTC. This was an energy supply contract with a temperature clause for volume risk management between Aquila Energy and Consolidated Edison Thind [2014]. The first pure weather trade was based on temperature and took place in Milwaukee for the winter of 1997-1998 between Enron and Koch Industries. This trade marked the beginning of the modern weather derivatives market and demonstrated the potential for financial instruments to manage weather-related risks.

Temperature-based contracts are the most popular and actively traded weather derivatives currently available [WRMA, 2011]. The most common instrument types are based on indices

[1] `https://www.epa.gov/outdoor-air-quality-data/how-aqi-calculated`

that measure the cumulative average temperature (CAT), the cumulative heating-degree days (HDD), and the cumulative cooling-degree days (CDD). For a given measurement period $[\tau_1, \tau_2]$, the indices are calculated with

$$CAT(\tau_1, \tau_2) = \sum_{t=\tau_1}^{\tau_2} T(t) \tag{2.1}$$

$$HDD(\tau_1, \tau_2) = \sum_{t=\tau_1}^{\tau_2} max\{b - T(t), 0\} \tag{2.2}$$

$$CDD(\tau_1, \tau_2) = \sum_{t=\tau_1}^{\tau_2} max\{T(t) - b, 0\} \tag{2.3}$$

where $T(t)$ is the average of the daily maximum and minimum temperature for day t, and b is a base temperature level, typically set at $18^oC/\ 65^oF$. By accumulating the days above the base level, the CDD index quantifies the need for cooling during the measurements period. Similarly, the HDD reflects the need for heating. The CAT index reflects the aggregated temperature during a period of time by accumulating the average temperatures, without considering any threshold. Derivatives, such as futures, swaps and options, are written on these indices, and the contracts can cover periods such as weeks, months or seasons. There is a rich research literature on weather markets and pricing methods for traded instruments [Jewson and Brix, 2005]. See [Schiller et al., 2012] and [Benth and Saltyte Benth, 2012] for an overview of approaches to temperature derivative pricing. Commonly used techniques include historical pricing (burn analysis), temperature index modeling daily temperature modeling [Alaton et al., 2002], [Benth and Saltyte Benth, 2005] , [Benth and Saltyte Benth, 2007], [Benth and Saltyte Benth, 2011] and utility based approaches [Cao and Wei, 2004], [Davis, 2001], [Platen and West, 2004] and [Brockett et al., 2006]. Alternative approaches to pricing are further described in Section 3.3.

These instruments are used in a variety of sectors that are impacted by weather, including energy production, agriculture and tourism [WRMA, 2011]. By taking a position in a CDD futures contract, an energy company can swap a fixed level of the index against the floating CDD to offset financial loss due to reduced electricity demand. Farmers can hedge against crop failure by procuring contracts that generate a payoff in the event of extreme temperature and drought. For companies with significant weather exposure, these instruments can serve as effective tools for volume risk management. See for example [Pérez-González and Yun, 2013] for an evaluation of the relationship between active weather risk management practices and firm value for weather-sensitive companies.

Episodes with extreme air pollution have led to lockdowns in several Asian megacities, causing a range of negative consequences for businesses, individuals, and the wider economy. Both the magnitude of the pollution incident itself, and the related shutdown of factories, schools,

airports, and restrictions on traffic may contribute to a significant reduction in economic activity [An et al., 2007], [Kumar et al., 2015]. As these challenges occur in fast-growing cities of considerable size, they may affect large populations and number of industries active in the international trade. The degree of air pollution during an incident can be evaluated by assessing the official Air Quality Index (AQI) for the location. Using the AQI as settlement reference for weather-type derivatives is a fairly new idea. Early works include [Li and Zhu, 2016] and [Sleire, 2018] [2] . Both of these independently developed studies design option contracts for financial pollution risk management based on air quality data from the city of Beijing, but the underlying AQI used are solely based on $PM_{2.5}$ measurements. Although particulate matter often is the dominant pollutant, a broader index including all of the air quality index constituents would be preferable. The quality and trustworthiness of this data is of vital importance, both for public health and the relevance of the suggested pollution derivatives, which depend on the AQI. Pollution levels are typically reported by local authorities or environmental protection agencies, and published studies have documented likely misreporting of official air quality data from China, see [Stoerk, 2016]. Further details regarding the index constituents and calculation methodology for the internationally established US EPA AQI standard are presented in Section 3.1. In [Xue et al., 2019], multiple pollutants are considered when calculating an AQI based on China's national air quality standards, which deviate from the most widely used US EPA AQI. This AQI is then employed to calculate an air quality deviation index to be used as settlement reference, which unfortunately does not reflect the pollution *level* well, and hence is less relevant when assessing the severity of pollution during a specific time period. An attempt to improve upon this can be found in [Liu et al., 2019], where a new index is suggested. The authors use the average AQI as settlement reference for contracts covering periods such as week, month or quarter, which is more consistent with existing indices from weather markets. They base the study on the national air quality standard from China, and describe daily AQI with an Ornstein–Uhlenbeck model with a seasonal function, similar to the temperature model in Benth and Saltyte Benth [2007], but with an AR-GARCH term. As the AQI is modeled directly without any transformation, the estimated seasonal function implies the series may take negative values, which is not consistent with the physical properties of pollutant concentrations.

The present study is based on the US EPA AQI, to contribute to the standardization of air quality reporting and promote transparency in the design of derivatives contracts for pollution risk management. We construct a settlement index that is similar to those in the weather market, as these are relevant and well known by market participants. The proposed AQI model involves a combination of a deterministic seasonal function and a continuous autoregressive model with a geometric model specification, to ensure that the predicted pollutant concentrations are non-negative. This model is used for pricing, and compared to alternative methods.

[2] The idea for the present study was first introduced at the EcoSta 2018 conference in Hong Kong.

3 Air pollution derivatives

3.1 The Air Quality Index

Poor air quality is caused by the release of pollutants such as gases, chemicals or airborne particles into the atmosphere. The dosage received by an individual is determined by the pollutant concentration level, and the time of exposure. The World Health Organization released it's first air quality guidelines in 1987, and these have been revised on several occasions, based on systematic reviews of a rapidly growing research literature of the adverse health effects [WHO, 2006]. In order to improve air quality management and report status to the public in a clear and concise manner, numerous countries have developed national air quality standards, see for example [Kuklinska et al., 2015] for a review of US and EU policies. The standards typically include a framework for converting concentration levels measured over a specific exposure period to a standardized scale that can be used for reporting and issuing health recommendations.

The Air Quality Index (AQI) developed by the US Environmental Protection Agency is frequently used for international reporting and comparisons, [EPA, 2009], [EPA, 2021]. The index covers the so-called *criteria air pollutants* from the US Clean Air Act. These are ground-level Ozone (O_3), Particulate Matter ($PM_{2.5}$ and PM_{10}), Carbon Monoxide (CO), Sulfur Dioxide (SO_2) and Nitrogen Dioxide (NO_2). Concentration measurements for each of the index constituents are converted to AQI values based on pollutant specific threshold values. Given a concentration measurement C, the AQI value can be obtained with

$$AQI = \frac{(AQI_{high} - AQI_{low})(C - C_{low})}{(C_{high} - C_{low})} + AQI_{low} \tag{3.1}$$

where C_{low} and C_{high} are the concentration breakpoints closest to C and AQI_{low} and AQI_{high} the corresponding AQI breakpoints. The index value is calculated for each of the criteria air pollutants, and rounded to the nearest whole number. The dominant pollutant with the highest AQI value is reported. For each hour, the following evaluation takes place

$$AQI = \max\left\{AQI_{O_3},\, AQI_{PM2.5},\, AQI_{PM10},\, AQI_{CO},\, AQI_{SO_2},\, AQI_{NO_2}\right\} \tag{3.2}$$

A detailed AQI calculation table for all of the criteria air pollutants is available online in the US EPA AQI technical documentation. For sake of illustration, we provide the PM 2.5 breakpoints in Table 1 and use them to convert a concentration measurement to the index value.

As an example, consider a measurement of $PM_{2.5}$ levels. The 24-hour average concentration is registered at $59\mu/m^3$. The AQI value is $(200 - 151)(59 - 55.5)/(150.4 - 55.5) + 151 \approx 153$. This degree of air pollution is considered *Unhealthy*.

Table 1: AQI calculation for PM 2.5, micrograms per cubic meter (μ/m^3)

AQI	$PM_{2.5}$ (μ/m^3)	Evaluation
0-50	0-12	Good
51-100	12.1-35.4	Moderate
101-150	35.5-55.4	Unhealthy for Sensitive Groups
151-200	55.5-150.4	Unhealthy
201-300	150.5-250.4	Very Unhealthy
301-400	250.5-350.4	Hazardous
401-500	350.5-500.4	Hazardous

AQI above 300 is classified as *Hazardous*. According to the US EPA guidelines *"values over 300 trigger health warnings of emergency conditions. The entire population is even more likely to be affected by serious health effects"*, [EPA, 2009]. Prolonged periods at these levels have led to government imposed lock-downs in several countries, forcing parts of industry, airports, schools and public transport to halt, see for example the Beijing [An et al., 2007] and Delhi [Kumar et al., 2015] episodes. Air pollution is closely linked to the rapid economic development and urbanization over the last decades, but also to the topography and meteorology of the regions where urbanization is taking place. An extensive review of air pollution in mega cities in China is presented in [Chan and Yao, 2008]. Further enquiries regarding underlying causes and resulting health effects can be found in [He et al., 2002], [Chen et al., 2013]. In the study performed by [Rohde and Muller, 2015], hourly measurements from a large number of sites during the period April-August 2014 are used to create pollution maps for eastern China. They report that 38% of the population experienced average concentrations that were unhealthy according to the US EPA standard, and estimate that air pollution is a contributing factor to approximately 17% of all deaths in China. When the AQI reaches extreme levels during pollution incidents, other more short term challenges may appear. Airport activity and general traffic is hindered due to poor visibility, parts of industry is forced to shut down, and a large proportion of the public will suffer from acute respiratory problems. These are the *hazardous* 300-level AQI events, which have led to lock downs on numerous occasions in several of the large cities in Asia.

It is worth noting that the reported pollution concentration is highly dependent on the location of the measuring device, and one may observe variation across metering stations placed relatively close to each other. For this reason, the overall pollution level in a city is assessed by including measurements from multiple stations, located in different types of environments. There may also be strong seasonality in air pollution levels, due to social factors such as increased need for heating during wintertime. The official AQI values for a given area are typically published by public institutions and environmental protection agencies.

3.2 Settlement index for air pollution derivatives

The construction of settlement references based on trustworthy and publilcy available AQI measurements is a crucial initiative towards ensuring accurate and reliable trading of derivatives contracts. The index should be based on data from multiple monitoring stations, to certify readings are representative of the overall air quality in a specific region. This requirement is also expected to reduce the likelihood of manipulation and bias, as the values from different stations can be cross-verified and compared. Following standard practice from temperature markets [Alaton et al., 2002], [Benth and Saltyte Benth, 2005], we base the settlement reference on the daily AQI value. Here, we calculate the daily average with

$$AQI(t) = \frac{1}{24}\sum_{i=1}^{24}(AQI_i) \tag{3.3}$$

for the hours $i = 1, \dots 24$ of day $t = 1, \dots, n$. The index to be constructed have to meet certain requirements to be suited for financial risk management. In order to be relevant, it must be a representative and accurate reflection of the source of risk to be protected against for market participants. The calculation methodology should be objective, transparent, easily accessible to all parties involved, and free from bias and manipulation. In order to achieve this, we suggest to primarily focus on the time limited pollution incidents, using well-established and familiar concepts from existing weather markets. The 300-level AQI events represent extraordinary circumstances, where a larger proportion of the public and industry is likely to face difficulties. These are also the levels where where governments typically intervene, as policy actions are informed by AQI-frameworks and related monitoring programs. In contrast to [Xue et al., 2019], [Liu et al., 2019], we will not attempt to define new index types specific for the purpose of pollution risk management, but rather implement settlement references similar to the CDD used in the temperature market. This threshold based approach is well suited for capturing effects from pollution incidents, it is well known among market practitioners, and there is a rich research literature on pricing methods. Furthermore, as the US EPA AQI is considered the *de facto* standard for international pollution reporting, we also prefer using this over national frameworks, in order to promote standardization and improve transparency.

The calculation of the proposed *Hazardous Air Quality Index (HAQI)* score is based on the daily average AQI for a specific location. If the reported value exceeds a pre-specified limit, the *Hazardous* 300 level, the severity of the breach is calculated by subtracting 300 from the measured value. The index value for a specific period, such as a week, is then calculated as the sum of all the daily HAQI scores for that period. For a specific time interval $[\tau_1, \tau_2]$, the

cumulative Hazardous Air Quality Index is thus given by

$$HAQI(\tau_1, \tau_2) = \sum_{t=\tau_1}^{\tau_2} max\{AQI(t) - 300, 0\} \tag{3.4}$$

where $AQI(t)$ is the daily mean value, and the measurement period is typically of fixed length, like *week* or *month*. The index may be used as underlying settlement reference for various derivatives contracts, such as futures or options. In this work we will focus on the latter. Options written on the HAQI will give the contract holder a right, but not an obligation to exercise, and the payoff received will depend on the realized HAQI, an agreed strike value for the index, K, and the contract size, α. The buyer of a call option with strike K and contract size α will have the following profit function

$$\pi_b = \alpha \max\{HAQI(\tau_1, \tau_2) - K, 0\} - C_0(\tau_1, \tau_2; K) \tag{3.5}$$

where the first term is the payoff from the contract upon the time of settlement $t \geq \tau_2$, and $C_0(\tau_1, \tau_2; K)$ is the option premium paid. In the event that the HAQI ends up below the strike value, the option expires worthless, and the option holder is left with a loss equal to the premium. For the option seller, the profit is given by

$$\pi_s = C_0(\tau_1, \tau_2; K) - \alpha \max\{HAQI(\tau_1, \tau_2) - K, 0\} \tag{3.6}$$

For for the purpose of illustration, we will specify a HAQI based option contract for the city of Xi'an, China as an example. The contract size, which serves as a scaling factor for the payoff received, is set to $\alpha = 100$ USD. We consider the HAQI for February 2024, assuming an agreed index strike value $K = 200$.

Contract type:	Call option
Settlement index:	HAQI
Location:	Xi'an, China
Period:	Jan-2024
Strike value:	200
Contract size:	100 USD

The holder of this contract will have the right, but not the obligation to exercise at maturity. This would be attractive if the accumulated $HAQI > 200$ for the period 1. January 2024 to 31. January 2024. Given the contract parameters, the value of the option is a function of the stochastic HAQI. What would be a fair price for this instrument?

3.3 Approaches to pricing in incomplete markets

In a complete market, where there is an active market for the underlying asset, the value of a derivative contract can be determined by evaluating the costs of forming a replicating portfolio with an identical payoff structure, see [Bjork, 2009]. As an example, consider a call option on an actively traded asset with price $S(t)$. In the absence of arbitrage, the premium at time t should be equal to the discounted value of the expected payoff at time of settlement T under the risk neutral probability measure

$$C_t(\tau_1, \tau_2; K) = e^{-r(T-t)} \mathbb{E}_Q \big[max\{S(T) - K, 0\} \mid \mathcal{F}_t \big] \tag{3.7}$$

In an incomplete market, however, this type of arbitrage argument cannot be used and alternative valuation strategies are necessary. Weather derivatives fall into this category, and several methods have been proposed for their valuation. According to the literature review by [Schiller et al., 2012], the methods for valuing temperature derivatives can be divided into three groups: burn analysis, index modeling, and daily temperature modeling. [Benth and Saltyte Benth, 2012], add a fourth category: utility-based approaches. As these methods may also be applicable for AQI-based derivatives, we will explore them further in the following. We will also comment briefly on potential uses of climate models and machine learning techniques. In order to identify alternative pricing strategies for derivatives written on the AQI, it is necessary to consider a range of methods and techniques, and their strengths and weaknesses.

The historical pricing method, also known as burn analysis, is a simple actuarial technique that uses past observations to determine the payoff of a specific contract. The price of the instrument is obtained by taking the average of the discounted value of these payoffs, while also accounting for a market risk premium [Jewson and Brix, 2005]. Using this approach, we would need to collect historical AQI values, calculate the realized HAQI´s and evaluate contract payoffs for a given strike value. The option premium is found by taking the average of the discounted payoffs. This method requires access to a reasonably sized historical time series and relies on the data being representative of future values. This can be problematic for air pollution data, as many available time series are short, and the data generating process may change due to efforts to reduce emissions or stricter regulations. Despite these limitations, historical pricing is still a relevant benchmark, and it is straightforward to calculate.

Index value simulation is a second type of actuarial approach to pricing, which involves the use of statistical models to simulate the underlying settlement reference for the derivative [Jewson and Brix, 2005]. By generating a large number of scenarios for the index, the expected value of the contract may be calculated, taking into account the distribution of possible outcomes and the probabilities of each scenario occurring. This method is useful in cases where there

is limited historical data available, as it allows the use of models to generate scenarios based on current and expected conditions, and additional factors such as the effects policy change may be incorporated. The large number of scenarios considered may also offer more stable price estimates compared to the historical approach [Schiller et al., 2012]. Finding a suitable distribution for the AQI to be used in the simulation method can be a challenge with limited data. By considering alternative distributions and validating their accuracy, we will be able to generate results that may be compared with competing approaches, to inform the pricing of AQI-based derivatives contracts.

The method of daily temperature modelling involves constructing a model for the daily average temperature and using it to simulate a large number of temperature series. These simulated series can then be used to value weather derivative contracts. This approach offers the advantage of pricing all contracts at a single location with a single model, ensuring consistency in prices. The daily temperature modelling method has become a popular approach in the weather derivatives market, with numerous studies and alternative model specifications proposed in the literature. Some examples of these studies include work by [Benth and Saltyte Benth, 2012], [Alaton et al., 2002] and [Dornier and Queruel, 2000]. In order to apply this approach for pollution derivatives, we will need to develop a stochastic model that describes the dynamics of the daily AQI. An important point to underline in this regard is that the pollutant concentration is a non-negative value, due to its physical properties. Hence, we should avoid using the common model specifications for temperature, where the variable is allowed to take negative values. This is particularly important if a falling trend is present in the reported AQI values. After fitting the model, we sample a large number of realizations and calculate contract payoffs. The option premium is found by taking the average discounted value from the process simulations. This approach may also allow us to test alternative assumptions regarding potential pollution trends due to government interventions and pollution reducing efforts. A potential downside of this approach is the reliance on the quality of the model, and it may be perceived as complex and difficult to implement by practitioners.

The utility-based approach to pricing weather derivatives relies on economic models that uses the concept of utility theory to value weather-related risks. The method involves modeling the relationship between weather events and the economic outcomes they generate in various scenarios. In [Cao and Wei, 2004], temperature is included as a source of uncertainty in the economic environment in an equilibrium pricing framework [Lucas, 1978], and the relationship between weather and the overall economic output is evaluated. The analysis suggest that the market price of risk for temperature derivatives is significant, indicating the importance of considering weather-related risks in financial decision-making. However, it relies on a number of assumptions and can be complex to implement. These methods require a good understanding of the underlying economics and the relationship between weather and the economy. Associations

between air pollution levels and economic activity have been subject to investigation in several studies, but it would be complex and difficult to implement as a foundation for pricing in practical applications. Other studies that may be placed in utility-based approach include [Davis, 2001], [Platen and West, 2004] and [Brockett et al., 2006].

Another approach to consider for pricing temperature derivatives is based on physical temperature models and weather forecasts, which utilizes expert knowledge and atmospheric models to estimate the impact of weather events on temperature, see Jewson and Brix [2005]. This approach has the advantage of incorporating a more detailed understanding of the physical processes that drive temperature patterns, but may be more complex and resource-intensive than other methods. There are also several projects focusing on ambient air quality that may serve useful, see for example the Community Multiscale Air Quality (CMAQ) modelling system, [Appel et al., 2017], [Byun and Schere, 2006]. For an overview of approaches to air pollution modelling, the reader is referred to [Zannetti, 2013] and [Daly and Zannetti, 2007]. Machine learning techniques can also be applied for pricing. By training models on historical weather data, they may be utilized to make predictions that can be used to calculated the expected contract payoffs to be discounted in order to find the derivative value.

In the following we will focus on the historical approach, index simulation and stochastic modelling. By comparing the results of these three methods, the analysis will provide insights into the strengths and limitations of each approach and inform the further development of pricing models for AQI-based derivatives contracts.

3.4 Stochastic Air Quality Index model

Empirically, the AQI share many traits with meteorological time series, such as temperature. However, the pollutant concentrations and their corresponding AQI values will by their physical properties always be non-negative. For this reason, we will not adopt the commonly used approach in temperature markets and model the series directly, [Alaton et al., 2002], but rather opt for the following specification for the AQI:

$$AQI(t) = e^{\Lambda(t)+X_1(t)} \tag{3.8}$$

Here, $AQI(t)$ is the daily air quality index, $\Lambda(t)$ is a deterministic function describing the trend and seasonality, and $X_1(t)$ is a stochastic process modeling the AQI dynamics. Similar models have been used for wind speed [Benth and Saltyte Benth, 2009] and electricity prices Benth et al. [2008]. This geometric specification is particularly relevant, given the fact that urban air pollution is directly impacted by human behaviour, and countries struggeling with high levels

actively attempt to reduce pollution. Models that may allow a falling trend in the AQI to produce negative index values should be avoided. We will focus on the log-transformed index

$$P(t) = \Lambda(t) + X_1(t) \tag{3.9}$$

and model the seasonality for $P(t)$ with the truncated Fourier series

$$\Lambda(t) = a_0 + a_1 t + \sum_{i=1}^{I_1} b_i \, cos(\frac{2i\pi t}{365}) + \sum_{j=1}^{J_1} c_j \, sin(\frac{2j\pi t}{365}) \tag{3.10}$$

where a potential trend may be reflected in the parameter a_1. Seasonal effects on yearly, half-yearly, or higher scales may be identified by increasing I_1 and J_1 and evaluating model fit. Similarly to [Benth et al., 2007], [Benth and Saltyte Benth, 2009] and others, we use a continuous-time autoregressive (CAR) model for the stochastic component. Let $\mathbf{X}(t)$ be a stochastic process in $\mathbb{R}^p$ for $p \geq 1$ defined by the vectorial Ornstein-Uhlenbeck equation

$$d\mathbf{X}(t) = A\mathbf{X}(t)dt + \mathbf{e}_p \sigma(t) dL(t) \tag{3.11}$$

where $L(t)$ is a Lévy process, $\mathbf{e}_k$ is the kth unit vector in $\mathbb{R}^p$, $k = 1, ..., p$ and $\sigma(t) > 0$ is a real-valued and square integrable function. A is the $p \times p$ matrix

$$A = \begin{bmatrix} 0 & 1 & 0 & \dots & 0 \\ 0 & 0 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & 1 \\ -\alpha_p & -\alpha_{p-1} & -\alpha_{p-2} & \dots & -\alpha_1 \end{bmatrix} \tag{3.12}$$

where $\alpha_k, k = 1, ..., p$ are constants. In order to represent $\mathbf{X}(t)$ explicitly, the multi-dimensional Ito's Lemma is applied to (3.11), and we get:

$$\mathbf{X}(s) = \exp(A(s-t))\mathbf{X}(t) + \int_t^s \exp(A(s-u))\mathbf{e}_p \sigma(u) dL(u) \tag{3.13}$$

for $s \geq t \geq 0$ and $\mathbf{X}(t) \in \mathbb{R}^p$. Then, the model for the log-transformed daily air quality index dynamics can be specified as

$$P(t) = \Lambda(t) + X_1(t) \tag{3.14}$$

$$X_1(t) = \mathbf{e_1}^T \mathbf{X}(t) \tag{3.15}$$

where $\mathbf{e_1}^T$ denote the transpose of $\mathbf{e_1}$, $\Lambda(t)$ is the seasonal function in (3.10) and X_q is the

qth coordinate of the vector $\mathbf{X}$, $q = 1, ..., p$. Here, $X_1(t)$ is a continous autoregressive model of order p, denoted CAR(p), and the special case of $p = 1$ corresponds to an Ornstein-Uhlenbeck process. In Benth [Benth et al., 2008] it is demonstrated that $X_1(t)$ can be approximately represented by an AR(p)-dynamics by utilizing a Euler discretization scheme. Considering the example where $p = 3$, they present the connection between the parameters of the CAR(3) and AR(3) models

$$\begin{aligned} 3 - \alpha_1 &= \beta_1 && (3.16) \\ 2\alpha_1 - \alpha_2 - 3 &= \beta_2 && (3.17) \\ \alpha_2 + 1 - (\alpha_1 + \alpha_3) &= \beta_3 && (3.18) \end{aligned}$$

where $\beta_1, \beta_2, \beta_3$ are the paramters of the discrete model. This relationship allows us to compute the parameters $\alpha_1, \alpha_2, \alpha_3$ for the CAR(3) model from the fitted AR(3) model. Generally, a CAR(p) process is stationary if the eigenvalues of the matrix A all have negative real parts, [Benth et al., 2008]. Finally, to account for potential seasonal variation in volatility, a truncated Fourier series is fitted for the squared model residuals

$$\sigma^2(t) = a_0 + \sum_{i=1}^{I_2} b_i \, cos(\frac{2i\pi t}{365}) + \sum_{j=1}^{J_2} c_j \, sin(\frac{2j\pi t}{365}) \tag{3.19}$$

Following the principles in [Benth et al., 2007], [Benth and Saltyte Benth, 2009] we estimate the model for $AQI(t)$ in a step-wise fashion. First, the seasonal function in (3.10) is fitted on the log-transformed data via least squares and used to deseasonalize $P(t)$. Second, we fit an AR(p)-model to the deseasonalized series $P(t) - \Lambda(t)$, calculate the CAR(p) parameters and check for stationarity. Finally, $\sigma(t)$ is fitted for the residuals, in order to capture any seasonal volatility in the air pollution dynamics.

4 Stochastic modelling of the daily Air Quality Index

In the following we conduct an empirical analysis of daily AQI observations from three severely polluted cities in China: Anyang, Shijiazhuang and Xi'an. The analysis of data from three different locations provides several benefits in understanding the dynamics of air pollution and developing models to price AQI derivatives, as different cities may have individual sources and intensities of pollution. AQI dynamics may vary based on factors such as population density, geography, industrial activity, and transportation patterns. After a brief descriptive analysis evaluating summary statistics and distributions for pollution levels, we fit the model presented in previous section in a step-wise fashion.

Figure 1: Location of the selected cities

4.1 Description of data

The data used in this analysis is collected from a pool of metering stations located in Anyang, Shijiazhuang and Xi'an. The observations consists of daily average AQI values for each of the five index constituents during the period 1. January 2014 to 15. January 2023. The dominant pollutant on a specific day is selected as the city AQI, leaving us with three time series covering 3275 days for further analysis. Missing observations (approximately 0.3%) are imputed using linear interpolation, and the AQI values are reported in the $0 - 500$ range, in accordance with the US Environmental Protection Agency guidelines, [EPA, 2009], [EPA, 2021]. An overview of the data sources can be found in Table 8 in the Appendix.

The descriptive statistics of the daily average AQI values in our analysis reveal several concerning trends. The mean values for Anyang, Shijiazhuang, and Xi'an are 161.2, 168.4, and 141.8, respectively. These values are all well above the recommended levels according to the US EPA AQI guidelines, which set a maximum AQI value of 100 as the threshold for good air quality. This suggests that the air quality in these cities is consistently poor, with high levels of pollutants that can have significant impacts on public health and well-being. The standard deviations for the three cities are 68.6, 80.3, and 63.9, respectively. These values indicate that there is a considerable amount of variability in the daily AQI values across each location. The skewness values also reflect that the distributions have longer upper tails, and the reported third quartile values are all above the *Unhealthy* 150 level of the US EPA AQI. For two of the cities, this is also the case for both mean and median values.

Table 2: Descriptive statistics for $AQI(t)$ and $P(t)$ per city

	Anyang	Shijiazhuang	Xi'an	ln(Anyang)	ln(Shijiazhuang)	ln(Xi'an)
Obs.	3,302	3,302	3,302	3,302	3,302	3,302
Mean	161.2	168.4	141.8	5	5	4.9
Std.dev.	68.6	80.3	63.9	0.4	0.4	0.4
Skewness	1.8	2	1.7	0.1	0.5	0
Excess kurtosis	5	4.8	5.2	1	0.6	0.6
Min	23	36	16	3.1	3.6	2.8
1 Quartile	118	118	98	4.8	4.8	4.6
Median	151	153	132	5	5	4.9
3 Quartile	184	189	167	5.2	5.2	5.1
Max	500	500	500	6.2	6.2	6.2

Figure 2 presents the pollution distribution plots for the three cities. Reported AQI values have a long upper tails, and there are multiple observations above the *Hazardous* 300 level. The log-transformation produces fairly symmetric distributions. We also evaluate the Box-Cox transform

$$y^{(\lambda)} = \begin{cases} (y^{\lambda} - 1)/\lambda & \text{if } \lambda \neq 0 \\ \ln y & \text{if } \lambda = 0 \end{cases} \tag{4.1}$$

where $y^{(\lambda)}$ is the transformed variable and the optimal λ is found via maximum likelihood estimation. $ln(y)$ is a special case for $\lambda = 0$.

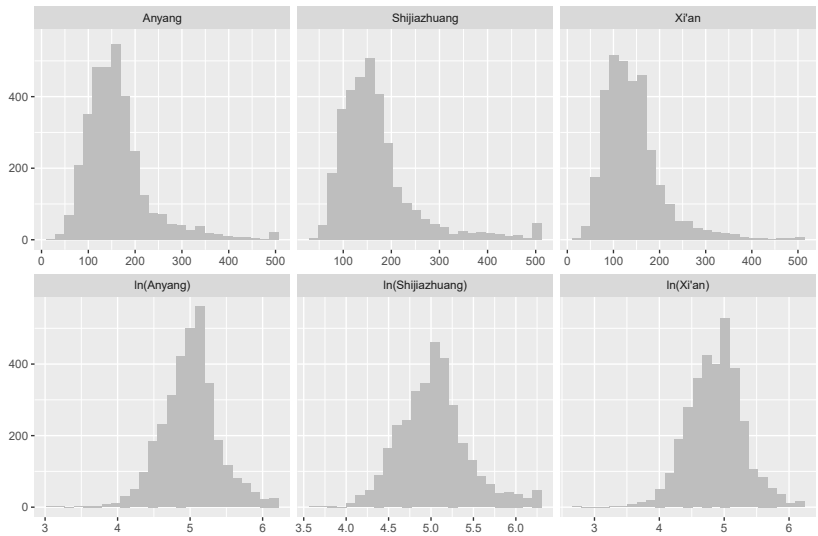

Figure 2: Histograms for $AQI(t)$ and $P(t)$ per city

This technique was used by [Benth and Saltyte Benth, 2009] when analyzing time series for daily average wind speeds for the purpose of pricing wind futures contracts. Their data share several of the properties also seen in AQI measurements, such as seasonality and non-negativity. By applying the Box-Cox transformation, we are effectively considering an extended version of the model described in Equation (3.8)

$$AQI(t) = \begin{cases} [\lambda(\Lambda(t) + X_1(t)) + 1]^{1/\lambda} & \text{if } \lambda \neq 0 \\ e^{\Lambda(t)+X_1(t)} & \text{if } \lambda = 0 \end{cases} \tag{4.2}$$

We estimate the optimal λ to be 0.003, -0.223 and -0.051 for Anyang, Shijiazhuang and Xi'an, respectively. The two parameters with smallest absolute values cannot said to be significantly different from zero, hence the log-transform is applied. With Equation (4.2), the $AQI(t)$ cannot be guaranteed be positive, unless $1/\lambda$ is an even number [Benth and Saltyte Benth, 2009]. A pragmatic solution for the city of Shijiazhuang could be to set λ to the closest value producing an even number, which is $\lambda = -0.25$. This does not produce significant changes in symmetry of the distribution, and we decide to stay consistent and use the log-transformation for all cities.

In order to implement the index simulation strategy for derivatives pricing in Section 5.1, we need to explore possible distributions for the AQI. The Gamma, Log-normal and Weibull distributions are commonly used in finance and other fields to model phenomena that exhibit long tails or skewness. Capturing seasonal effects in AQI data is an important consideration when developing pricing models for derivatives contracts, and it may be sensible to fit distributions for each month individually. This will be explored further in Section 5.1.

4.2 Trend and seasonality

The daily AQI values in Figure 3 show a strong seasonal pattern, where both the level of pollution, and the variance appear higher during the winter season, which is similar to the findings for temperature [Benth and Saltyte Benth, 2007] [Alaton et al., 2002] and wind [Benth and Saltyte Benth, 2009]. This is not surprising, since fuel burning for heating during winter is one of the causes of air pollution, and the fact that meteorology plays an important role in the concentration, distribution, and transport of pollutants in the atmosphere [Guttikunda and Gurjar, 2012]. In order to model the seasonality and potential trends in the ambient air pollution, we estimate the parameters of (3.10) on the log-transformed series $P(t)$. The mean functions are fitted with least squares minimization using $I_1 = J_1 = 4$. Simpler models have been evaluated, but they produced a weaker description of the seasonality for some cities, and we keep the same number of terms for all. The parameter estimates are reported in Table 3.

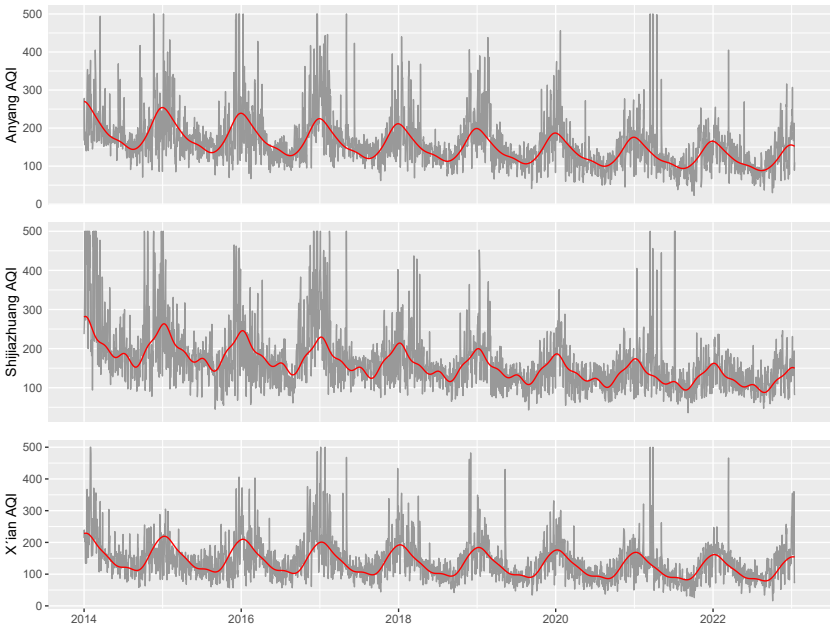

Figure 3: Daily AQI and fitted seasonal function per city

All of the estimated seasonal functions have a statistically significant negative trend parameter, suggesting that efforts made to reduce pollution may have led to improvements during the period. COVID-19 lockdowns might also played a role for the cities affected [He et al., 2020], [Xu et al., 2021]. When constructing a model to capture the seasonality of the AQI, it is essential to consider the fundamental physical properties of pollutant concentrations. Any model that allows the predicted AQI to take negative values over time would be in conflict with these physical properties, see for example [Liu et al., 2019]. There, the estimated seasonal variations were also not well-described, requiring the authors to add a more complex AR-GARCH term for the stochastic component.

As pollution levels are influenced by a range of social factors, including human behavior and activity patterns, a further extension of Equation 3.10 could be to include calendar effects in the seasonal function, to capture potential variation due to weekends and holidays. We have fitted models with indicator variables for weekends, but the estimated parameters were not found to be significant, and further exploration of calendar effects is left for future studies.

Table 3: Parameter estimates for the seasonal function $\Lambda(t)$

	Anyang	Shijiazhuang	Xi'an
(a_0) Constant	5.281 ***	5.347 ***	5.062 ***
(a_1) Trend	−0.0002***	−0.0002***	−0.0001***
(b_1) $cos(2\pi t/365)$	0.253 ***	0.204 ***	0.306 ***
(c_1) $sin(2\pi t/365)$	0.080 ***	0.059 ***	0.120 ***
(b_2) $cos(4\pi t/365)$	0.062 ***	0.074 ***	0.066 ***
(c_2) $sin(4\pi t/365)$	−0.028 ***	−0.023 ***	−0.011
(b_3) $cos(6\pi t/365)$	−0.003	−0.020 **	−0.004
(c_3) $sin(6\pi t/365)$	−0.00005	−0.006	−0.025 ***
(b_4) $cos(8\pi t/365)$	0.006	0.030 ***	−0.002
(c_4) $sin(8\pi t/365)$	−0.004	0.032 ***	0.018 **
R^2	0.418	0.357	0.405
Adjusted R^2	0.417	0.356	0.403
Residual Std. Error (df = 3292)	0.298	0.328	0.323
F Statistic (df = 9; 3292)	263.114 ***	203.420 ***	248.521 ***

*p<0.1; **p<0.05; ***p<0.01

4.3 Stochastic component

The seasonality of the data is described by the proposed deterministic function. In the subsequent step, our focus shifts towards modeling the stochastic component. To achieve this, we begin by removing the linear trend and the seasonal element from the data.

$$X(t) = P(t) - \Lambda(t) \tag{4.3}$$

In order to analyze the variability of $X(t)$, we plot the autocorrelation function (ACF) and the partial autocorrelation function (PACF) of the deseasonalized data, as shown in Figure 4. The elimination of the seasonality features from the data is evident, leaving behind an autoregressive structure for further analysis. Proceeding with the methodology outlined in Section 3.4, we estimate models to characterize the dynamics of $X(t)$.

The PACF plot indicate that an AR(2) process would sufficiently capture the autoregressive structure of the Shijiazhuang and Xi'an series. However, the Anyang data seem to require an additional lag. Similarly to [Benth and Saltyte Benth, 2009], we consider alternative ARMA(p, q) specifications for $p \leq 5$, $q \leq 5$, and evaluate the Akaike's Information Criterion (AIC). Several of the candidates have similar AIC, and we fit models both with and without a moving average term, see Table 11 in the Appendix for AQI simulations using ARMA(3,1). Results are not materially impacted , and we select the more parsimonious AR-structure for the remainder of

the analysis. To maintain consistency, we employ an AR(3) process for all three cities.

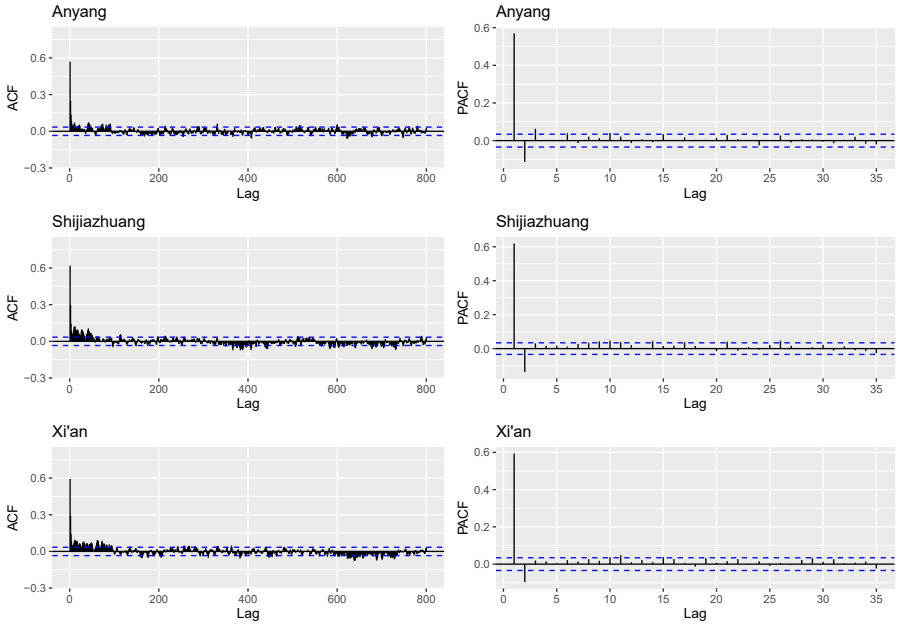

Figure 4: ACF (left) and PACF (right) for the deseasonlized $X(t)$

From here, we assume the dynamics of the deseasonalized data is described by

$$X(t) = \beta_1 X(t-1) + \beta_2 X(t-2) + \beta_3 X(t-3) + \sigma(t)\epsilon(t) \tag{4.4}$$

where $\beta_1, \beta_2, \beta_3$ are constants, $\epsilon(t)$ are i.i.d standardized random variables and $\sigma(t)$ is defined by Equation (3.19). We fit the AR(3) models using the R programming language, and calculate the CAR(3) parameters according to the Euler scheme described in Section 3.4. Estimation results are presented in Table 4.

Table 4: Parameter estimates for AR(3) and CAR(3) models.

	AR(3)			CAR(3)		
	β_1	β_2	β_3	α_1	α_2	α_3
Anyang	0.641***	-0.153***	0.063***	2.359	1.871	0.449
Shijiazhuang	0.708***	-0.158***	0.029 *	2.292	1.742	0.432
Xi'an	0.652***	-0.107***	0.018	2.348	1.803	0.437

*p<0.1; **p<0.05; ***p<0.01

All of the estimated parameters are significant at the level of 1% for the city of Anyang. This is also the case for the two first lags for the remaining cities. The third term is significant at the 10% level for Shijiazhuang, while for Xi'an it is not found to be significant.

For the CAR(3) model, the stationarity condition specifies that the eigenvalues of the matrix A should all have negative real parts, [Benth and Saltyte Benth, 2009], [Ichihara and Kunita, 1974]. This condition is fulfilled for the three models, and we conclude that they are all stationary.

4.4 Residuals

We now move on to analyze the residuals $\sigma(t)\epsilon(t)$ from the model in Equation (4.4). By evaluating the ACF plot of squared residuals presented in Figure 5, the seasonality of volatility in the residual process is revealed. Similar patterns have also been described for temperature and wind data [Benth et al., 2007], [Benth and Saltyte Benth, 2009]. In order to capture this, we estimate $\sigma^2(t)$ as described in in Equation (3.19) above.

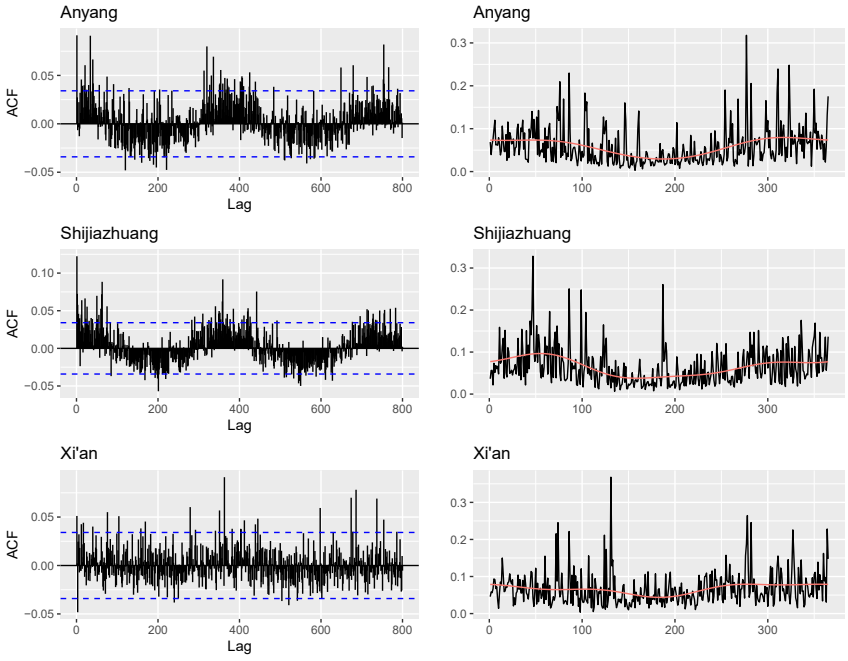

Figure 5: ACF for squared residuals (left) and empirical variance with fitted $\sigma^2(t)$ (right).

We calculate the empirical daily variance by computing the variance of the residuals for each day $d = 1, \ldots, 365$ across the year, disregarding leap years. This involves using 9 residual observations for each day of the year, corresponding to one observation for each year in the dataset. The lack of historical data is undoubtedly a limiting factor when performing such analysis, and the results should be interpreted with this in mind. We evaluate alternative numbers of terms in the Fourier series in Equation (3.19), and select a model with $I_1 = J_2 = 2$. The estimated parameters can be found in Table 5 .

Table 5: Parameter estimates for the seasonal volatility $\sigma^2(t)$

	Anyang	Shijiazhuang	Xi'an
Constant	0.059***	0.065***	0.067***
$cos(2\pi d/365)$	0.024***	0.024***	0.013***
$sin(2\pi t/365)$	−0.001	0.006 **	−0.006 *
$cos(4\pi d/365)$	−0.007 **	−0.006 **	−0.006 *
$sin(4\pi d/365)$	−0.002	0.006 *	−0.002
R^2	0.151	0.166	0.056
Adjusted R^2	0.141	0.156	0.045
Residual Std. Error (df = 360)	0.042	0.042	0.045
F Statistic (df = 4; 360)	15.988***	17.882***	5.319***

*p<0.1; **p<0.05; ***p<0.01

The results indicate that there is a difference in the magnitude of fluctuations between the cold and warm seasons, see also Figure 5. Similar to the investigation of temperature in [Benth et al., 2007], we find that the variability during the cold season is noticeably higher than during the warm season, which may be due to the increased burning of fuel for heating and meteorological factors. This effect seem to be present for all cities, but less pronounced for Xi'an.

We obtain the standardized residuals $\epsilon(t)$ by dividing out the estimated volatility function in the initial residual process $\sigma(t)\epsilon(t)$ and evaluate the distribution based on the summary presented in Table 6 . Standardized residuals have zero mean and unit standard deviation, but the moderate

Table 6: Evaluation of standardized residuals $\epsilon(t)$

	Anyang	Shijiazhuang	Xi'an
Mean	-0.0002	-0.0003	-0.0004
Std.dev	1.000	1.000	1.000
Max	5.770	6.934	6.454
Min	-5.538	-4.891	-4.757
Skewness	-0.432	-0.296	-0.345
Excess kurtosis	2.887	1.896	2.288

left skew and fat tails indicate that the standard normal may not be the ideal distribution for driving the dynamics in Equation (3.11), i.e. the Lévy process may not be a Brownian motion. In our simulation experiments, we utilized both the standard normal distribution and the normal inverse Gaussian distribution, without observing significant differences in results. Consequently, we chose to proceed with the former. Figure 6 displays an example realization from the Monte Carlo simulation, utilizing standard normal innovations.

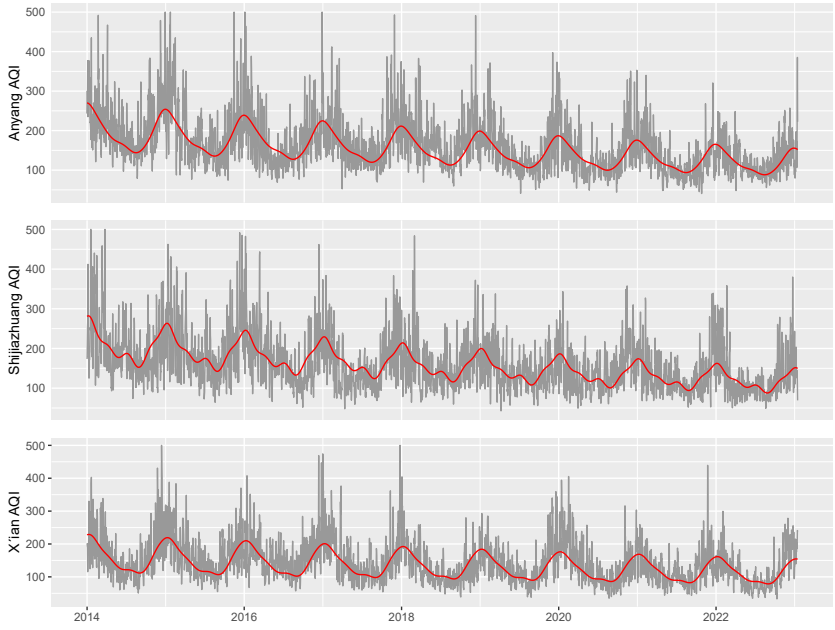

Figure 6: Daily AQI simulation with N(0,1) innovations and seasonal function per city

5 Financial pollution risk management

With three different approaches available for pricing air pollution derivatives, we are equipped to estimate the prices of options on the HAQI index. Our analysis will involve implementing historical pricing (burn analysis), index modelling and daily AQI simulation utilizing the model developed in the previous section. We evaluate and compare the results obtained by applying these methods on data from all three cities included in the study. Finally, we examine a practical use case by providing an example of how air pollution derivatives may be used to hedge against financial risk due to pollution incidents.

5.1 Pricing air pollution derivatives

Similar to weather derivatives, the underlying settlement reference for the contracts being introduced in this context is not traded. Nor have there been any transactions in the instruments themselves. A common approach when pricing contracts in weather markets is to consider the market price of risk (MPR), and represent the derivative price as the discounted expected value under the martingale measure Q characterized by the MPR, utilizing Girsanov or Esscher transformations, see for example [Alaton et al., 2002]. A simplifying assumption often used is setting $P = Q$, presuming risk neutrality. This approach will be applied here, and we keep in mind that this choice may underestimate prices. In [Benth et al., 2011], a study exploring pricing of Asian temperature risk is performed. They find that the MPR for temperature derivatives exhibits a seasonal structure, which arises from the seasonal variation in temperature volatility. By having knowledge of the relationship between MPR and seasonal variations, they compute prices for locations where there are no listed instruments available for trading in the market. An extension of the present study could be to explore similar strategies, utilizing common patterns in the seasonality of volatility in the AQI and temperature for each location. In the following, we present the findings from our analysis of contract pricing. Results are summarized in tables per pricing strategy, using the notation C^{city} for the option premium per location, assuming a HAQI strike value $K = 200$ and contract size normalized to $\alpha = 1$ for monthly contracts.

In the historical burn analysis, we use past observations from the three cities to calculate the daily HAQI index, and report mean AQI and realized contract payoffs. Results based on data from the entire period 1. January 2014 to 31. December 2022 are presented in Table 7. Due to the relatively short history available, the estimated premiums have high variability. Given the fact that the index is aiming to catch an upper tail event (daily breach of 300 AQI level) and that the trend in the AQI is falling, this is not unexpected. When only calculating on data from most recent years, both the average AQI, HAQI and expected option payoffs are considerably lower, but based on few observations. To address this limitation, alternative strategies such as index simulation and daily AQI simulation using the stochastic model are also evaluated.

In the index simulation study, we explore the use of three different distributions, namely gamma, weibull, and lognormal, for generating daily AQI values for each month separately, to capture the seasonal variation. By fitting these distributions to the historical AQI data, we create scenarios of possible AQI values and use them for pricing the options. An alternative strategy could be to simulate the settlement index directly, but such distributions would need to be fitted on a low number of observations. Our approach involves estimating $3 \times 3 \times 12 = 108$ distributions for the Monte Carlo experiments, and generating daily AQI samples for each month and calculating HAQI and contract payoffs.

Table 7: Historical average monthly AQI and option premiums using strike value 200 HAQI.

Month	AQI^{Anyang}	$AQI^{Shijiazhuang}$	AQI^{Xian}	C^{Anyang}	$C^{Shijiazhuang}$	C^{Xian}
1	227.68	240.04	207.54	222.11	620.00	154.56
2	196.65	196.89	180.68	64.89	320.44	50.89
3	172.52	180.36	165.22	51.89	137.33	50.78
4	155.05	158.21	133.56	0	0	0
5	139.35	145.76	121.96	14.56	26.89	2.44
6	136.47	141.91	106.79	0	0	0
7	122.78	146.38	103.46	0	66.67	0
8	118.59	125.67	101.45	0	0	0
9	130.99	132.82	106.63	0	0	0
10	146.33	163.50	125.69	0	123.22	0
11	183.18	185.81	163.74	58.89	171.44	12.00
12	205.32	204.27	181.72	288.56	452.44	59.89

Similarly to the historical pricing method, the results obtained from index simulation will be sensitive to which part of the available history that is used for fitting the distributions, due to the falling trend in the air pollution series. Simulation results from distributions based on the entire sample can be found in Table 9 in the Appendix, where average AQI and estimated option premiums are reported. The produced output exhibit a consistent trend, with the option contracts generating payoffs predominantly during the winter season, which includes the first and last three months of the year. Although all of them yield average AQI values that are comparable to the historical data, their tail behaviors differ. Specifically, during the winter months, the lognormal distribution generates higher option premiums than the other two distributions for most of the months. It is also the only distribution generating a small option premium greater than zero outside winter season, which also can be found in the historical data.

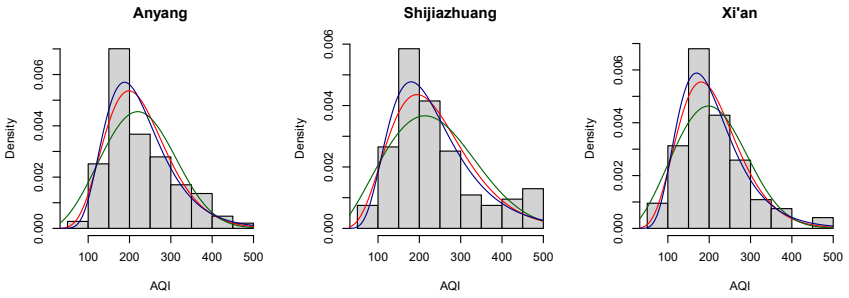

Figure 7: Fitted gamma (red), weibull (green) and lognormal (blue) for the month of January

However, the lognormal still yields significantly lower premiums than the burn approach. This may be attributed to the limited history used by the latter, where extreme values are more

prevalent in the initial years of the sample. To illustrate the differences between the three alternatives, the empirical and fitted distributions for the month of January are presented in Figure 7. Among the distributions examined, the lognormal appears to offer the best fit, while the weibull distribution exhibits difficulties in accurately describing the data, in contrast to the findings in [Alexandridis and Zapranis, 2013], where wind data is studied. All of them seem to underestimate the mass in the upper tail for the city of Shijiazhuang during the winter season.

Using a stochastic model for daily AQI simulation has several advantages for pricing AQI-based derivatives. With this approach, temporal information, such as the observed falling trends, can easily be accounted for when pricing contracts out-of-sample. The method also serve as a convenient instrument for presenting spatial variation, as the estimated trend, seasonality and volatility may be compared across different locations by evaluating the fitted models. Such comparisons would require more effort, and may be perceived as somewhat ad-hoc when using the other methods. Daily AQI simulations also enable the analyst to price contracts of arbitrary lengths with a single tool, without having to model these specifically. The burn analysis suffers under the limited availability of historical data, and the index simulation approach may require the user to make difficult choices regarding which distribution to utilize and how to make best use of scarce data.

The choice of model for the stochastic component in the daily simulation strategy remains an active area of research, however. Ornstein-Uhlenbeck type models with both normal and non-normal innovations have been extensively used in weather markets, [Alaton et al., 2002], [Benth and Saltyte Benth, 2007]. Other methods employed include autoregressive structures and GARCH for the residual process. In Table 10 and Table 11 in the Appendix, we provide a summary of the simulation results from the stochastic model using both AR(3) and ARMA(3,1) for the stochastic component. The upper panels present monthly results when all years are included in the calculation of monthly average AQI and option premiums. The middle panel displays comparable findings based on the first four years of the simulated sample, while the bottom panel depicts the output when focusing on the last four years. Both models generate small positive premiums outside the winter season for all sub samples, as seen in the burn analysis, but estimated option prices does not differ greatly between the two.

Since the decreasing trend in hazardous air quality events is prominent across all cities examined, the simulation outcomes for the most recent period reflect lower prices for options written on the HAQI. While this trend may be used to support future predictions and form price expectations based on the lower panels of Table 10 and Table 11, it should be noted that this period coincides with the COVID-19 pandemic, and caution is required when interpreting results. In Figure 8, we present results from the Monte Carlo simulations using the stochastic model driven by normally distributed innovations. We calculate a confidence interval based on the 1% and 99% percentiles

for the monthly AQI, and add historical observations, for the city of Xi'an. Such analysis may be informative when evaluating the development over time, and related uncertainty.

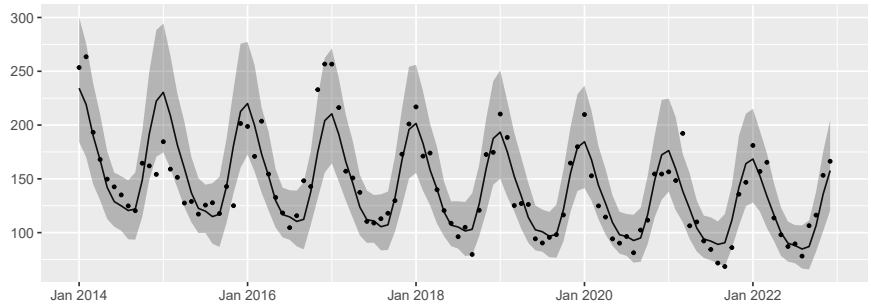

Figure 8: Monthly AQI (points) and simulated confidence intervals for the city of Xi'an.

Based on our analysis, the level of ambient air pollution as reported by the AQI appears to have followed a falling trend, but with large fluctuations. It is important to keep in mind that the study has been performed using a rather limited set of historical data, and the most recent years include the COVID-19 pandemic, which may not be representative of typical future conditions. Given these uncertainties, it can be challenging to predict what the future will bring, but pollution levels are likely to stay well above recommended levels for the cities in this study, in particular during the winter season.

5.2 Hedging against air pollution risk

In the following, we will provide a practical example of how air pollution derivatives may be used to protect against financial losses due to pollution events. Consider an organizer of a major sports event. The organizer fears unusually high air pollution levels may affect ticket sales and revenue. It decides to procure a HAQI call option for the month of January, the time the event is taking place. The details of the contract are

Contract type:	Call option
Settlement index:	HAQI
Location:	Xi'an, China
Period:	Jan-2024
Strike value:	200 HAQI
Contract size:	100 USD

In the event that the pollution level gets so severe that this level is breached, the organizer

expects fewer tickets to be sold, or even a cancellation of the games, but the call option will generate a payout. Assume the index ends up at 400 HAQI, and that the option premium paid was 4.500 USD. The net financial profit for the organizer would then be

$$\begin{aligned}\text{Profit} &= \text{Contract size} \times \text{Max}(\text{HAQI} - \text{Strike value}, 0) - \text{Premium} \\ &= 100 \times \text{Max}(400 - 200, 0) - 4.500 \\ &= 15.500 \text{ USD}\end{aligned}$$

Based on an assessment of the financial exposure and potential loss, the organizer can calculate the number of contracts to purchase for risk mitigation. From the option premium and agreed strike value, we can also calculate a break-even level for the HAQI index:

$$\begin{aligned}100 \times (\text{HAQI}^* - 200) - 4.500 &= 0 \\ \text{HAQI}^* &= (4.500/100 + 200) = 245\end{aligned}$$

This HAQI index level can be breached in several ways, for example with five days at 350 AQI during a pollution incident. AQI above this level has been observed 49 days in our data set for Xi'an. Now, consider some alternative outcomes. In situations where the HAQ index ends above the calculated break even level, the option contract holder receives a net positive payoff from the transaction, as illustrated above. In the event that the index falls between the strike and the break-even value, the contract will still be exercised, but the premium paid will not be fully covered by the payout. Finally, if the risk the organizer insures against does not materialize and HAQ < 200, the option contract expires worthless.

6 Summary and conclusion

In this study, we explore the financial impact of air pollution and the potential for establishing a market for derivatives contracts based on publicly available air pollution data for risk mitigation. We emphasize the importance of a reliable and trustworthy settlement reference based on internationally accepted standards, such as the US Environmental Protection Agency's Air Quality Index (AQI), and recommend implementing indices similar to those used in existing weather markets. Our findings suggest that established methods for pricing weather derivatives can also be applied for pricing contingent claims based on the AQI. We demonstrate that the proposed model, which combines a seasonal AQI function and stochastic component described by a CAR(3) process with seasonal volatility, is suited for describing the daily AQI dynamics. The model is used for pricing option contracts via Monte Carlo methods, and compared to historical analysis and index simulation. These pricing strategies are concluded to be more

vulnerable to the scarce data available in this area, and further exploration of the stochastic modelling approach for pollution derivatives pricing is recommended.

Acknowledgements

This project has received financial support from the Research Council of Norway, via the industrial Ph.D. scheme, and the Norwegian Finance Market Fund.

References

Peter Alaton, Boualem Djehiche, and David Stillberger. On modelling and pricing weather derivatives. *Applied mathematical finance*, 9(1):1–20, 2002.

Antonis Alexandridis and Achilleas Zapranis. Wind derivatives: Modeling and pricing. *Computational Economics*, 41(3):299–326, 2013.

X. An, T. Zhu, Z. Wang, C. Li, and Y. Wang. A modeling analysis of a heavy air pollution episode occurred in beijing. *Atmospheric Chemistry and Physics*, 7(12):3103–3114, 2007. doi: 10.5194/acp-7-3103-2007. URL `https://acp.copernicus.org/articles/7/3103/2007/`.

K. W. Appel et al. Description and evaluation of the community multiscale air quality (cmaq) modeling system version 5.1. *Geoscientific Model Development*, 10(4):1703–1732, 2017. doi: 10.5194/gmd-10-1703-2017. URL `https://gmd.copernicus.org/articles/10/1703/2017/`.

Fred Espen Benth and Jurate Saltyte Benth. Stochastic modelling of temperature variations with a view towards weather derivatives. *Applied Mathematical Finance*, 12(1):53–85, 2005.

Fred Espen Benth and Jurate Saltyte Benth. The volatility of temperature and pricing of weather derivatives. *Quantitative Finance*, 7(5):553–561, 2007.

Fred Espen Benth and Jurate Saltyte Benth. Dynamic pricing of wind futures. *Energy Economics*, 31(1):16–24, 2009.

Fred Espen Benth and Jurate Saltyte Benth. Weather derivatives and stochastic modelling of temperature. *International Journal of Stochastic Analysis*, 2011, 2011.

Fred Espen Benth and Jurate Saltyte Benth. *Modeling and pricing in financial markets for weather derivatives*, volume 17. World Scientific, 2012.

Fred Espen Benth, Saltyte Benth Jurate, and Steen Koekebakker. Putting a price on temperature. *Scandinavian Journal of Statistics*, 34(4):746–767, 2007.

Fred Espen Benth, Jurate Saltyte Benth, and Steen Koekebakker. *Stochastic modelling of electricity and related markets*, volume 11. World Scientific, 2008.

Fred Espen Benth, Wolfgang Karl Härdle, and Brenda López Cabrera. Pricing of asian temperature risk. In *Statistical Tools for Finance and Insurance*, pages 163–199. Springer, 2011.

Tomas Bjork. *Arbitrage Theory in Continuous Time*. Oxford University Press, 3 edition, 2009. URL `https://EconPapers.repec.org/RePEc:oxp:obooks:9780199574742`.

Patrick L Brockett, Mulong Wang, Chuanhou Yang, and Hong Zou. Portfolio effects and valuation of weather derivatives. *Financial Review*, 41(1):55–76, 2006.

Daewon Byun and Kenneth L Schere. Review of the governing equations, computational algorithms, and other components of the models-3 community multiscale air quality (cmaq) modeling system. *Applied mechanics reviews*, 59(2):51–77, 2006.

Melanie Cao and Jason Wei. Weather derivatives valuation and market price of weather risk. *Journal of Futures Markets: Futures, Options, and Other Derivative Products*, 24(11):1065–1089, 2004.

Chak K Chan and Xiaohong Yao. Air pollution in mega cities in china. *Atmospheric environment*, 42(1):1–42, 2008.

Zhu Chen, Jin-Nan Wang, Guo-Xia Ma, and Yan-Shen Zhang. China tackles the health effects of air pollution. *The Lancet*, 382(9909):1959–1960, 2013.

Aaron Daly and Paolo Zannetti. Air pollution modeling–an overview. *Ambient air pollution*, pages 15–28, 2007.

Mark Davis. Pricing weather derivatives by marginal value. *Quantitative finance*, 1(3):305–308, 2001.

Fabien Dornier and Michel Queruel. Caution to the wind. *Energy & power risk management*, 13(8):30–32, 2000.

EPA. A guide to air quality and your health. *United States Environmental Protection Agency*, 2009.

EPA. Reviewing national ambient air quality standards (naaqs): Scientific and technical information, 2021. URL `https://www.epa.gov/naaqs`.

Bhola Ram Gurjar, A Jain, A Sharma, A Agarwal, P Gupta, AS Nagpure, and J Lelieveld. Human health risks in megacities due to air pollution. *Atmospheric Environment*, 44(36):4606–4613, 2010.

Sarath K Guttikunda and Bhola R Gurjar. Role of meteorology in seasonality of air pollution in megacity delhi, india. *Environmental monitoring and assessment*, 184:3199–3211, 2012.

Guojun He, Yuhang Pan, and Takanao Tanaka. The short-term impacts of covid-19 lockdown on urban air pollution in china. *Nature sustainability*, 3(12):1005–1011, 2020.

Kebin He, Hong Huo, and Qiang Zhang. Urban air pollution in china: current status, characteristics, and progress. *Annual review of energy and the environment*, 27(1):397–431, 2002.

Kanji Ichihara and Hiroshi Kunita. A classification of the second order degenerate elliptic operators and its probabilistic characterization. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 30(3):235–254, 1974.

Stephen Jewson and Anders Brix. *Weather derivative valuation: the meteorological, statistical, financial and mathematical foundations*. Cambridge University Press, 2005.

Karolina Kuklinska, Lidia Wolska, and Jacek Namiesnik. Air quality policy in the us and the eu–a review. *Atmospheric Pollution Research*, 6(1):129–137, 2015.

Prashant Kumar, Mukesh Khare, Roy Harrison, William Bloss, Alastair Lewis, Hugh Coe, and Lidia Morawska. New directions: air pollution challenges for developing megacities like delhi. *Atmospheric Environment*, 122:657–661, 2015.

SY Li and XW Zhu. Contract design of haze index options and pricing by monte carlo simulation. *J. Syst. Eng. Theor. Pract*, 36:2477–2488, 2016.

Zhuoxin Liu, Laijun Zhao, Chenchen Wang, Yong Yang, Jian Xue, Xin Bo, Deqiang Li, and Dengguo Liu. An actuarial pricing method for air quality index options. *International Journal of Environmental Research and Public Health*, 16(24):4882, 2019.

Robert E. Lucas. Asset prices in an exchange economy. *Econometrica*, 46(6):1429–1445, 1978. ISSN 00129682, 14680262. URL `http://www.jstor.org/stable/1913837`.

David Mage, Guntis Ozolins, Peter Peterson, Anthony Webster, Rudi Orthofer, Veerle Vandeweerd, and Michael Gwynne. Urban air pollution in megacities of the world. *Atmospheric environment*, 30(5):681–686, 1996.

Ioannis Manisalidis, Elisavet Stavropoulou, Agathangelos Stavropoulos, and Eugenia Bezirtzoglou. Environmental and health impacts of air pollution: a review. *Frontiers in public health*, page 14, 2020.

Miriam E Marlier, Amir S Jina, Patrick L Kinney, and Ruth S DeFries. Extreme air pollution in global megacities. *Current Climate Change Reports*, 2:15–27, 2016.

Francisco Pérez-González and Hayong Yun. Risk management and firm value: Evidence from weather derivatives. *The Journal of Finance*, 68(5):2143–2176, 2013.

Eckhard Platen and Jason West. A fair pricing approach to weather derivatives. *Asia-Pacific Financial Markets*, 11(1):23–53, 2004.

Robert A Rohde and Richard A Muller. Air pollution in china: mapping of concentrations and sources. *PloS one*, 10(8):e0135749, 2015.

Frank Schiller, Gerold Seidler, and Maximilian Wimmer. Temperature models for pricing weather derivatives. *Quantitative Finance*, 12(3):489–500, 2012.

Anders D. Sleire. Stochastic modelling of ambient air quality and pricing of air pollution derivatives. The 2nd International Conference on Econometrics and Statistics, City University of Hong Kong, 2018.

Thomas Stoerk. Statistical corruption in beijing's air quality data has likely ended in 2012. *Atmospheric Environment*, 127:365–371, 2016.

Sarfraz Thind. As temperatures tumble in north america, weather derivatives warm up. *Institutional Investor, January*, 23, 2014.

WHO. *Air quality guidelines: global update 2005: particulate matter, ozone, nitrogen dioxide, and sulfur dioxide*. World Health Organization. Regional Office for Europe., 2006.

WRMA. *Weather Risk Derivative Survey*. Weather Risk Management Association, 2011.

Zhihu Xu, Ru Cao, Xin Hu, Wenxing Han, Yuxin Wang, Jing Huang, and Guoxing Li. The improvement of air quality and associated mortality during the covid-19 lockdown in one megacity of china: an empirical strategy. *International Journal of Environmental Research and Public Health*, 18(16):8702, 2021.

Jian Xue, Yan Xu, Laijun Zhao, Chenchen Wang, Zeeshan Rasool, Mingming Ni, Qin Wang, and Deqiang Li. Air pollution option pricing model based on aqi. *Atmospheric Pollution Research*, 10(3):665–674, 2019.

Paolo Zannetti. *Air pollution modeling: theories, computational methods and available software*. Springer Science & Business Media, 2013.

Appendix

Table 8: AQI data providers

City	Source
Anyang	Henan Environmental Protection Agency
Shijiazhuang	Hebei Province Environment Protection Agency
Xi'an	Shaanxi Provincial Environmental Protection Office
	Xi'an Environmental Protection Agency

Source: `https://aqicn.org/data-platform`

Table 9: Simulated average monthly AQI and option premiums using gamma (top), weibull (middle) and lognormal (bottom) and strike value 200 HAQI, $N = 1000$ random draws.

Month	AQI^{Anyang}	$AQI^{Shijiazhuang}$	AQI^{Xian}	C^{Anyang}	$C^{Shijiazhuang}$	C^{Xian}
1	226.26	235.26	209.07	111.72	330.83	50.84
2	196.88	196.56	180.78	13.02	69.85	0.99
3	172.89	179.88	165.46	0.40	9.50	0.09
4	155.09	157.56	133.67	0	0	0
5	138.94	145.78	121.98	0	0	0
6	136.54	142.20	106.66	0	0	0
7	122.76	146.71	103.61	0	0	0
8	118.74	125.61	101.61	0	0	0
9	130.89	133.08	106.39	0	0	0
10	146.09	163.24	125.91	0	2.16	0
11	183.23	185.68	163.41	4.39	15.73	0.19
12	205.72	204.50	181.47	83.25	149.63	3.83
1	226.35	235.95	208.61	109.74	383.50	45.35
2	197.22	197.42	180.86	8.73	100.17	0.57
3	171.08	180.69	164.60	0.25	10.19	0.11
4	153.55	157.37	132.84	0	0	0
5	138.14	144.81	121.67	0	0	0
6	134.77	141.67	106.46	0	0	0
7	122.32	145.81	103.01	0	0	0
8	118.81	125.89	101.12	0	0	0
9	130.35	132.49	106.78	0	0	0
10	145.63	164.30	125.64	0	3.49	0
11	183.05	185.44	163.49	2.33	20.30	0
12	204.73	204.21	181.97	103.43	183.30	2.44
1	225.69	233.41	208.07	146.72	348.33	78.81
2	196.19	194.76	180.45	27.45	94.35	5.49
3	171.82	180.48	164.28	3.53	22.39	0.39
4	155.01	158.46	133.38	0	0.20	0
5	138.88	145.20	120.94	0	0	0
6	136.51	141.94	106.96	0	0	0
7	123.04	145.84	103.44	0	0.14	0
8	118.91	125.63	101.24	0	0	0
9	131.17	133.19	107.40	0	0	0
10	148.06	162.99	126.89	0.80	9.51	0.20
11	183.82	184.76	164.04	19.57	34.28	3.29
12	204.45	200.99	180.93	111.96	176.31	11.14

Table 10: Average monthly AQI and option premiums simulated from daily AQI model. Calculations are based on all data (top), the period 2014-2018 (middle) and the period 2019-2023 (bottom) and strike value 200 HAQI, $N = 1000$ random draws using AR(3) for $X(t)$.

Month	AQI^{Anyang}	$AQI^{Shijiazhuang}$	AQI^{Xian}	C^{Anyang}	$C^{Shijiazhuang}$	C^{Xian}
1	213.25	219.21	198.28	160.42	270.36	71.86
2	199.55	199.16	184.05	75.01	146.02	21.92
3	176.42	179.14	159.67	25.11	71.55	3.62
4	155.50	163.96	140.55	2.78	14.08	0.48
5	142.14	144.75	119.89	0.08	0.30	0.03
6	133.75	145.25	107.99	0	0.18	0
7	123.26	144.76	105.88	0	0.39	0
8	118.25	126.55	101.78	0	0.08	0
9	127.53	131.86	103.66	0.12	0.47	0.01
10	150.73	161.76	127.20	4.26	15.09	0.60
11	182.09	182.55	161.82	42.62	56.40	8.59
12	207.81	203.65	188.41	125.47	140.97	44.43
1	252.57	263.01	223.79	359.61	576.99	152.02
2	230.35	233.49	204.70	156.15	290.55	43.62
3	203.90	210.00	177.57	54.25	148.16	7.53
4	179.74	192.24	156.31	6.24	30.40	1.06
5	164.37	170.13	133.09	0.17	0.68	0.07
6	154.49	170.33	120.04	0	0.40	0
7	142.40	170.10	117.58	0	0.89	0
8	136.33	148.56	112.84	0	0.18	0
9	146.68	154.17	114.67	0.26	1.05	0.02
10	173.38	188.90	140.53	9.28	32.39	1.27
11	209.80	213.87	178.87	89.91	118.33	17.15
12	239.10	237.79	208.73	256.05	287.60	86.37
1	181.04	183.53	177.21	15.76	43.62	12.97
2	169.10	165.46	163.69	5.98	19.82	2.93
3	149.50	149.06	142.10	0.76	6.63	0.27
4	131.73	136.32	124.86	0.01	0.42	0.02
5	120.44	119.89	106.88	0	0	0
6	113.33	120.69	96.20	0	0	0
7	104.60	119.97	94.34	0	0	0
8	100.53	105.06	90.78	0	0	0
9	108.79	110.23	92.78	0	0	0
10	128.69	135.38	114.03	0.14	0.65	0.04
11	155.02	151.95	145.30	2.83	3.56	1.49
12	177.03	170.00	168.33	12.55	13.17	7.30

Table 11: Average monthly AQI and option premiums simulated from daily AQI model. Calculations are based on all data (top), the period 2014-2018 (middle) and the period 2019-2023 (bottom) and strike value 200 HAQI, $N = 1000$ random draws using ARMA(3,1) for $X(t)$.

Month	AQI^{Anyang}	$AQI^{Shijiazhuang}$	AQI^{Xian}	C^{Anyang}	$C^{Shijiazhuang}$	C^{Xian}
1	214.15	218.94	199.82	168.28	263.86	86.02
2	199.70	198.77	184.87	79.72	144.71	26.93
3	176.39	178.84	160.43	26.60	72.13	4.72
4	155.61	163.89	140.73	2.83	17.24	0.74
5	142.03	144.75	120.39	0.05	0.48	0.03
6	133.68	145.21	108.10	0	0.22	0
7	123.28	145.06	106.07	0	0.44	0
8	118.30	126.76	101.63	0	0.08	0
9	127.71	131.99	103.67	0.10	0.44	0.00
10	150.74	161.29	127.37	4.87	17.00	0.48
11	181.98	182.80	161.92	44.71	57.60	9.37
12	207.20	203.90	187.88	123.83	141.58	45.36
1	253.87	261.86	227.79	378.46	554.91	183.15
2	230.96	232.08	206.53	167.32	283.92	54.28
3	204.02	208.76	179.05	57.07	147.43	10.12
4	180.28	192.14	156.80	6.33	37.46	1.56
5	164.06	169.58	134.23	0.09	1.08	0.07
6	154.30	169.94	120.23	0	0.50	0
7	142.36	170.41	118.00	0	0.99	0
8	136.30	148.85	113.00	0	0.18	0
9	146.86	154.44	115.13	0.21	0.93	0.01
10	173.18	188.40	140.45	10.40	36.90	1.02
11	209.01	213.83	179.23	92.87	122.35	18.75
12	238.53	237.92	207.88	251.97	286.98	86.86
1	181.69	183.77	177.09	16.17	45.79	14.97
2	169.17	165.95	163.77	6.31	22.39	3.15
3	149.38	149.25	142.06	1.11	6.82	0.35
4	131.54	136.21	124.77	0.01	0.53	0.07
5	120.29	120.21	106.68	0	0.01	0
6	113.47	120.85	96.19	0	0	0
7	104.54	120.29	94.30	0	0	0
8	100.66	105.17	90.44	0	0	0
9	108.91	110.05	92.51	0	0.00	0
10	128.72	134.70	114.39	0.14	0.63	0.05
11	155.39	152.72	144.80	3.41	3.71	1.31
12	176.54	170.39	168.10	12.18	15.50	8.16

Graphic design: Communication Division, UiB / Print: Skipnes Kommunikasjon AS

uib.no

ISBN: 9788230862100 (print)
9788230865873 (PDF)