Taming the Elements:
An Analysis of Temperature Derivatives and Spatial Basis Risk in the Netherlands

MSc Thesis

written by

Abderrahmane El Aarfaoui

2322668

under the supervision of **Dr. Lisa Sheenan** and submitted to the Board of Examiners in partial fulfillment of the requirements for the degree of

MSc in Banking and Finance

Utrecht School of Economics

Universiteit Utrecht

2023

To my parents

Acknowledgement

I would like to extend my heartfelt appreciation to Dr.Lisa Sheenan, for their invaluable guidance and unwavering support. Their extensive knowledge and encouragement played a pivotal role in enabling me to successfully accomplish this research and compile this thesis.

Taming the Elements:

An Analysis of Temperature Derivatives and Spatial Basis Risk in the Netherlands

Abstract

The key aim of this thesis discuss and apply a pricing model for temperature derivatives with payoffs deriving from Dutch average daily temperatures. We first run a thorough analysis of 52 years of daily average temperatures in nine dutch cities and use a Ornstein-Uhlenbeck process with seasonal volatility to capture the features of daily average temperatures. We use explicit pricing formulas and a Monte Carlo simulation to approximate the price of HDD and CAT options and propose an approach to estimate the market price of risk that relies on the Newton-Raphson method. The thesis also investigates the effects of spatial basis risk in the Netherlands and suggests trading a new security which combines an exchange traded derivative and a basis derivative, along with a discussion on accounting for this security in a profit maximization framework.

Contents

Acronyms

ACF Autocorrelation Function.

ADF Augmented Dickey–Fuller.

AIC Akaike Information Criterion.

AR Autoregressive.

CAT Cumulative Average Temperature.

CDD Cooling Degree Day.

CME Chicago Mercantile Exchange.

DAT Daily Average Temperature.

EU European Union.

FLD Full Load Days.

FLH Full Load Hours.

HDD Heating Degree Day.

IEA International Energy Agency.

JB Jarque Bera.

LNG Liquefied Natural Gas.

OU Ornstein–Uhlenbeck.

PACF Partial Autocorrelation Function.

Introduction

As climate change and increasing global temperatures are becoming an increasingly pressing global issue, this thesis delves into the intricate market of temperature derivatives in the Netherlands, seeking to shed light on the complexities of this vital department of risk management. Building upon the importance of risk management in today's economy, recent events such as the Russian Federation's *"special military operation"* have caused a surge in energy price volatility, which further emphasizes the significance of understanding temperatures derivatives.

The energy crisis spurred by this event is characterized by unstable and unpredictable behaviour of energy markets, particularly those of liquid and gaseous fuels. Given that electricity generation technologies, and therefore supply and price formation, rely primarily on natural gas, the exogenous shock to energy markets has caused utility bills to soar to unprecedented levels, marking an unprecedented increase in gas prices north of 150% July 2021 and July 2022 in the EU (Council of the EU & the European Council, 2023a). To address this challenge and promote greater stability and predictability in energy prices and costs, the European Commission proposed a reform of the European electricity market through regulations to "*improve the Union's electricity market design*" and to "*improve the Union's protection against market manipulation in the wholesale energy market*" (European Commission, 2023). Nevertheless, because of high gas grid congestion and storage fullness at the north-west European hubs, the Title Transfer Facility (TTf) spot price temporarily fell to €30/MWh in late October but rose back to reach a peak of €150/MWh in early December and dropped back again to €70/MWh. The downward trajectory continued until mi-May 2023. Although, the EU reached gas storage levels that puts it 5% or 5 bcm above its 5-year average, and that gas prices began to normalize, the IEA warns against overly optimistic predictions amid the existence of significant risks for 2024. That is, 5 bcm is the equivalent of "*just two days of EU gas demand during a cold spell*" (International Energy Agency, 2022b).

This research considers another factor that contributes to the unpredictability of energy prices: weather, which stems from the fact that past research shows that energy demand depends greatly on temperature (Engle et al., 1992; Li & Sailor, 1995; Stoft, 2002; Zanotti et al., 2003). In addition, climate conditions have had an aggravating effect on energy prices in the EU. According to the Council of the EU &

the European Council (2023b), *"Heatwaves during summer 2022 have put additional pressure on energy markets"*, which experienced both high demand for cooling and low supply because of drought. Various factors influence energy supply and demand in the Netherlands. Among the most significant of these are weather patterns and temperature fluctuations. These have long been recognized as key drivers of energy demand. According to the International Energy Agency, the average surface temperature in the Netherlands rose by 2.3°C between 1901 and 2020, a rate of warming faster than the world's average in the 20$^{\text{th}}$ century. The agency's climate hazard assessment argues that the country is enduring more frequent and intense heatwaves, with higher expected temperatures during Dutch summers. Furthermore, the Netherlands is facing important shifts in the energy heating and cooling demand, therefore shrinking the number of Heating Degree Days (HDDs) and extending the number of Cooling Degree Days (CDDs)[1]. These climate change-induced shifts are thus poised to reduce energy consumption during winter, while boosting electricity consumption in the summer. Moreover, given the Dutch energy policy's robust emphasis on augmenting the electrification of buildings, the projected surge in electricity demand for cooling during summer season could stimulate stress on the electricity grid (International Energy Agency, 2022a).

Nonetheless, with the advent of innovative financial instruments such as temperature derivatives, the dynamics of energy markets have undergone a notable transition recently. Temperature derivatives are financial instruments that allow market participants to manage their exposure to market volatility by hedging against the risk of climate hazards, notably temperature fluctuations.

Climate change is indubitably a major contributor to the rise of diverse climate patterns in many regions around the world, including the Netherlands, which has far-reaching implications for the pricing and efficiency of energy. As such, investors interested in mitigating temperature risk face an intricate challenge. That is to say, trading temperature derivatives in the Netherlands is presently confined to one reference city, Amsterdam. As it will be shown at a later stage in this research, investors involved in trading temperature derivatives attempt to hedge away price risk. Clearly, this poses an arduous geographical or spatial basis risk challenge for parties who wish to hedge against temperature fluctuations in other Dutch cities. In the context of temperature derivatives, basis risk arises when the underlying asset or a *random variable*[2] in general, that a financial instrument is based on, does not exactly match the hedging needs of the hedger. In the case of temperature derivatives, it occurs when the reference city's temperature does not accurately reflect the temperature at the actual hedging location. For instance, in a given winter

[1] HDDs and CDDs indicate the deviation of daily mean temperatures from a reference temperature of 18°C or 65°F.

[2] Temperature does not match the definition of what an asset is.

season, the temperature of the city of Groningen may or may not be the same as Amsterdam's, geographical basis risk arises if it is not. Naturally, basis risk is greater when the expected temperatures of both cities diverge. For example, say a European natural gas importer is involved in a long-term contract of 15 years which binds them to pay some 90% of the contracted gas. The importer has distributors active in the Dutch market for space heating. However, suppose the city of Groningen experiences a warm winter, then the importer faces basis risk due to the absence of *publicly* traded temperature derivatives for Groningen, combined with a warm winter season that resulted in low demand for heating in the city. Consequently, the importer faces the challenge of having excess natural gas, which poses a considerable financial risk known as quantity risk. However, they have the option to mitigate this risk by investing in gas infrastructure, which is essential for managing seasonal variations in demand and serve as means to hedge against quantity related risks (Zweifel et al., 2017). This oversimplified scenario highlights the importance of managing basis risk in energy markets particularly in the face of diverse climate patterns that may have an impact on local demand. Moreover, it underscores the need for sophisticated financial and *climate* engineering techniques, including the development of localized temperature derivatives, to efficiently hedge against basis risk and circumvent the impact of climate-related risks on energy markets and consumers.

Nevertheless, risk managers often find themselves in awkward situations, forced to either hedge using the temperature of the reference city, Amsterdam for instance, or buy over-the-counter (OTC) temperature derivatives from private contractors. While the latter option offers highly localized temperature hedging thus eliminating spatial basis risk, it brings about concerns on credit or default risk, transparency, and fairness. In contrast, trading exchange listed weather derivatives, eliminates credit risk since the payoff is guaranteed by a third party such as the Chicago Mercantile Exchange (CME), and offers liquidity and fairly priced contracts. Furthermore, a method to mitigate basis risk involves investing in basis derivatives, which are based on the disparity or distance between between the weather index of a reference weather station and the specific location of the first seeking to hedge against weather-related exposure (Alexandridis & Zapranis, 2012).

Although Amsterdam is one of the few European cities for which temperature derivatives are traded on the CME, almost no research exists that looks into fair pricing of temperature derivatives in the Netherlands. While this double edged sword provides the opportunity to be one of the first empirical studies to investigate the subject, it also offers no indications on what models work best for Dutch data, nor does it indicate what premises to take into account in the matter. Yet, because of the Netherland's reliance on natural gas, it is an essential topic. As the country transitions towards renewable energy sources, current volatility in energy prices may persist for an lengthy period of time, negatively impacting the Dutch economy and

energy security. Taming the elements is about providing risk management insight for energy companies and industries exposed to temperature risk in general, but also the government, to mitigate these price fluctuations and march towards a more efficient, resilient, and sustainable use of energy in the context of wars and crises in general. That is, a more liquid temperature derivatives market will enable energy companies to hedge against low demand in multiple locations in the country, therefore lowering their losses and reducing the costs for the parties down the chain.

Finally, the inherent discrepancy between the temperature of Amsterdam and that of other Dutch cities, provoked by climate change, constitutes a basis risk concern and requires a more comprehensive understanding of the unique temperature profile. Against this backdrop, this thesis seeks to provide a comprehensive analysis of the dynamics of temperature derivatives and basis risk in different cities in the Netherlands. Therefore, drawing on a range of quantitative methods such as time series analysis, actuarial and arbitrage free pricing approaches, and Monte Carlo simulations, the thesis examines the key drivers of temperature derivatives pricing, analysing the impact of a variety of factors such as weather patterns, seasonal trends, and market sentiment. The research also explores implications of these dynamics for market participants, examining the challenges associated with managing basis risk in the Netherlands.

The thesis comprises eight chapters. The first Chapter covers the basics of electricity, gaseous fuels, through introducing the fundamental concepts and gradually building up to the weather derivatives market, where it is explained that the temperature derivatives market derives from the existence of those markets as a way to hedge against climate risk. The second Chapter, consists of an extensive review of the literature, where pricing approaches and basis risk assessment methods are discussed. In Chapter 3, the temperature data is presented and the methodology of the analysis is introduced, namely, modeling techniques of seasonality, temperature, and volatility. In Chapter 4 we conduct a basic data preprocessing to clean the dataset and undertake time series analysis to describe the data both numerically and graphically. That is, we apply classical methods to identify trends and seasonalities but also modern techniques such as Prophet and MSTL. Chapter 5 develops a stochastic model that accurately describes the DAT. Specifically, the model analytically describes the behavior of the DAT process in the Netherlands, which will be used, on the one hand, in Chapter 6 to simulate future DATs under the real world measure $\mathbb{P}$ measure, and to derive pricing formulas for HDD and CAT indices under the risk neutral measure $\mathbb{Q}$ in Chapter 7, on the other hand. Moreover, Chapter 7 proposes a new way of estimating the market price of risk of weather derivatives which is conditional on the existence of a market price of the derivatives. Ultimately, Chapter 8, investigates spatial basis risk. More specifically, we discuss the effects of this risk and verify whether risk managers should worry about the risk that arises

from the choice of the where the derivative contract in written.

Surprisingly, although the Netherlands is a comparatively small country, the results indicate that the temperature options don't have the same price for a given seasonal strip[3]. Additionally, in the Netherlands, basis risk constitutes a major risk as it is found that the differences in the seasonal temperatures can induce major losses if the choice of the weather station is inaccurate. Nevertheless, Chapter 8 proposes a way to hedge the error that may arise from the choice of the wrong station by trading a product where the underlying is a combination of one exchange traded temperature derivative and one, or more, OTC basis temperature derivative.

[3]A seasonal strip is a contract period on which temperature derivatives are written.

Theoretical Framework

The weather derivatives market predominantly comprises energy companies as its primary participants. In fact, this market emerged from contracts created by energy companies to hedge against meteorological scenarios such as warm winters and cold summers. This chapter provides a detailed overview of gas and electricity markets as explained by Zweifel et al. (2017), and an introduction to the weather derivatives market following the theory the work of Alexandridis and Zapranis (2012) and Jewson and Brix (2005).

1.1 Gas Markets

Natural gas is the third primary energy source after crude oil and hard coal. It is crucial to the development of highly fuel-efficient industries with the lowest CO_2 emissions. Nonetheless, the gaseous fuels market is subject to constraining supply and demand properties. The lack of long-term contracts impacts the latter, whilst the former depends on volatile seasonal space heating demand which is linked to temperature levels.

1.1.1 Gas Infrastructure

The formula for an ideal gas is given by its pressure P, volume V, and temperature measure τ and a constant c

$$\frac{PV}{\tau} = c$$

1.1.1.1 Properties of Gas

Low-density methane makes up a significant share of natural gas. As Table A.1 shows, there are two types of natural gas, with different densities, upper heating values, and low heating values. Liquid gas primarily consists of byproducts derived from oil refining, such as propane and butane. As they are heavier than air, these components can be liquefied at moderate pressure and and distributed in pressurized containers. In contrast, town gas primarily consists of hydrogen and carbon monoxide (Zweifel et al., 2017).

1.1.1.2 Reserves and Extraction

The world's natural gas reserves are unevenly dispersed around the globe. In fact, the area of the world that extends from the Middle east through the Caspian Region to the north of the Russian Federation is known as the Strategic Energy Ellipse and comprises more than half of conventional natural gas reserves (see B.1b).

Table 1.1: Production and Reserves of Natural Gas

	Natural gas reserves (2020)		**Natural gas production (2021)**	
	Trillion (m^3)	**Share (%)**	**Billion (m^3)**	**Share (%)**
Russia	37.4	19.9	701.7	17.4
Iran	32.1	17.1	256.7	6.4
Qatar	24.7	13.1	177.0	4.4
United States	12.6	6.7	934.2	23.1
Canada	2.4	1.3	172.3	4.3
China	8.4	4.5	209.2	5.2
Saudi Arabia	6.0	3.2	117.3	2.9
United Arab Emirates	5.9	3.2	57.0	1.4
Netherlands	0.1	0.1	18.1	0.40
Germany	*	**	4.5	0.1
World	188.1	100	4036.7	100

** less than 0.05*

*** less than 0.05%*

Source: (BP, 2021) (BP, 2021)

1.1.2 Natural Gas Economy

The discovery of the large Groningen gas field in the Netherlands in 1959 was the milestone in the history of European gas markets as it served as a catalyst in connecting to other fields in Siberian, the North Sea, and North Africa through high pressure pipeline connections.

1.1.2.1 Gas Transportation

So far as long distances are concerned, natural gas is transported via high-pressure pipelines. Zweifel et al. (2017) explains that the rate Q at which this resource flows typically relies on the pressures P_1 and P_2 at the pipeline's start and end points, the diameter d of the pipeline and the distance l to be traversed

$$Q \sim \sqrt{\frac{P_1^2 - P_2^2}{l/d^2}}$$

It is noteworthy that pipelines benefit from economies of scale, but their transportation capacity often surpasses the market potential and the financial capacity of one

company. As a result, several companies could build and use the same pipeline in co-ownership while still being competitors. Additionally, the natural gas value chain is characterized by the hold-up problem where the absence of vertical integration between two monopolistic participants in a specific market results in double marginalization and hence a welfare loss. In other words, all market participants are better-off with cooperative behaviour since the profits are super-additive and costs are sub-additive.

An alternative to gas pipelines is Liquefied Natural Gas (LNG). For offshore transportation, the latter is less costly than the former at distances beyond 2000 km, while onshore transportation using LNG technology becomes more competitive at distances exceeding 3000 km. Nevertheless, Cayrade (2004) explains that, the process of transporting LNG not only consumes around one third of the energy contained in the chain itself, but also involves significant investments. For instance, a complete LNG chain connecting Port Said to the port of Cartagena includes a liquefaction plant at the export harbor, requiring an investment of approximately \$900 million. Additionally, there are two fleets of LNG vessels, which require an investment outlay of \$360 million, and a regasification facility in Cartagena, with a cost of \$320 million.

1.1.2.2 Gas Pricing

European natural gas markets have been governed by long-term take-or-pay contracts (ToP contracts) of 15 to 30 years' duration. These contracts require domestic importers to import roughly 90% of the contracted quantity of gas, regardless of unforeseen changes in demand, such as adverse weather conditions, whilst foreign producers base their prices on heating oil, coal, and heavy oil in order to ensure gas competitiveness in electricity markets. Consequently, these contracts allocate risk such that importers bear quantity risk while producers deal with price risk.

The early 21^{st} century witnessed European gas markets' liberalization, thereby instigating gas trade on spot and futures markets. Trading occurs on physical gas hubs with concentrated pipelines such as Zeebrugge hub in Belgium, or virtual gas hubs like the Title Transfer Facility in the Netherlands.

Due to the liquidity of gas markets, the spot price of gas hubs is an alternative reference price to heating oil for gas contracts. Although the prices of the two candidate references are correlated, seasonality affects gas prices more than prices of heating oil in that the former exhibits price spikes in excessively cold winters and hot summers, owing to substantial gas storage costs used to meet increased demand.

1.1.2.3 Third Party Access

Wholesale gas and electricity markets share several grid-related facets. However, gas markets have some unique aspects. The gas year starts on October 1 at 6.00 a.m, coinciding with the start of the heating season and the space heating market, and concludes on October 1st of the following year at 5:59 a.m. Nevertheless, as shown in Table 1.2 demand profiles of some consumers is significantly higher than that of others. Moreover, the minimum trading unit is a block of 1 MWh. Spot markets trade day ahead contracts, while futures markets trade monthly, quarterly, and yearly contracts.

Natural gas demand from end-users for space heating has strong seasonality. Full Load Hours (FLH) or Full Load Days (FLD) are used to measure consumption as follows

$$\text{FLH} = \frac{\text{Gas sales } (\text{m}^3/\text{a})}{\text{max.load } (\text{m}^3/\text{h})} \quad \text{and} \quad \text{FLD} = \frac{\text{Gas sales } (\text{m}^3/\text{a})}{\text{max.load } (\text{m}^3/\text{d})}$$

Table 1.2: Capacity utilization by end-users of Natural gas

	FLH (h/a)	**FLD (h/d)**	**Capacity Utilization (%)**
Private households	1500–2000	60–95	16–26
Real estate companies	1800–2700	75–110	20–30
Industrial customers	2500–5000	100–210	27–58
Market average	3600	150	41
Natural gas contracts with nearby wells	3000–4000	125–167	34–46
Block delivery	8000–8760	340–365	>93

Source: Erdmann and Zweifel (2008)

Given the uncertainty due to gas demand fluctuations, expected consumption can be effectively modeled using Heating Degree (HDD) days which yields accurate prediction of daily temperatures exceeding a reference temperature.

1.2 Electricity Markets

1.2.1 Properties of Electricity

The economic point of view treats electricity as a non-homogeneous phenomenon in contrast to the physical point of view. For the latter, electric power and electric work are said to be homogeneous. For the former, however, prices are subject to

temporal and geographical fluctuations, hence making electricity a heterogeneous product (Zweifel et al., 2017).

1.2.1.1 Electricity Substitutability

Electricity has a low elasticity demand (Praktiknjo, 2014), due to its limited substitutability by other energy sources. Explicitly, aside from its minor environmental impact, it is the source with maximum energy which makes electric applications more efficient than technologies based on fossil energy. Additionally, according to the laws of thermodynamics, there is no theoretical limit to the amount of energy that can be packed into a given space. As a result, electricity, being a versatile form of energy, is able to generate high temperatures by converting electrical energy into heat energy. Moreover, as a form of energy carried by electromagnetic fields, it does not have mass associated with it, which implies that electricity can be easily controlled and manipulated, enabling efficient and quick activation and deactivation processes. Consequently, some of its features imply that auto-generating electricity is more expensive than electricity provided by the grid in circumstances like blackouts.

1.2.1.2 Electricity Storability

As a result of the highly sensitive electricity demand to time and space, loading profiles of northern and southern countries differ, with daily peaks occurring at different hours during the day depending on the season of the year. Zweifel et al. (2017) explain that in order to levelize daily load curves, suppliers typically follow these strategies:

- Storing in periods of low demand and supplying the stored electricity in periods of high demand
- Undertake load management measures to implement varying service quality
- Shifting demand from peak to off-peak times using price differentiation

1.2.2 Electricity Generation

Ineffective load management, alongside the non-storability of electricity, result in unlevelized load profiles, which hampers supply and demand balancing. Consequently, power generators have to constantly adjust their electricity production to maintain a balance between supply and demand, ensuring that there is enough electricity to meet the needs of consumers at any given time.

1.2.2.1 Power Generation Resources

Using magnetic induction, the principle behind producing electricity in power plants, a sizable power plant could generate enough electricity to power approximately 200 000 houses. The characteristics of various producing technologies are presented in Table 1.3.

The overall cost of electricity production, known as the levelized cost of electricity, is determined by three main factors: the investment required for the power generation infrastructure, the rate at which the capacity if the infrastructure is utilized, and the cost of fuel used for electricity generation. The first one is the initial capital expenditure needed to build the power generation facility, the second one refers to how much of the installed capacity is actually used over a given period of time, while the last one is the expense incurred in obtaining the necessary fuel to generate electricity. A power plant could theoretically generate up to 8760 MWh of electricity annually. Nevertheless, power plants often run at low capacity factors. For instance, at low rates of capacity utilization, coal-fired plants incur higher levelized costs tan gas-fired plants due to larger expenditure per unit of capacity. However, Table 1.3 indicates that after 3000 hours of operation, coal-fired plants' low fuel cost cause their levelized cost to fall below that of gas-fired plants.

Table 1.3: Properties of power generation

	Fuel efficiency (%)	**Investment outlay (€/kW)**	**Useful life (years)**	**Fuel cost (€/MWh$_{el}$)**
Hard coal 700 MW	38-46	1250—1800	40	25-45
Lignite 700 MW	35-43	1350-1900	40	15-25
Nuclear 1400 MW	36	2400-5000	40	10-15
Gas turbine 200 MW	28-42	450-700	20	75-100
CCGT* with 700 MW	>58	680-900	30	50-70
Hydro-power 100 MW	80-90	1500-4000	50-80	-
Wind power onshore	40-50	1000-2500	20	-
Photovoltaics 1 MW	8-13	2000-4000	40	-
Fuel cells (<100 kW)	30-50	High	∼5	60-120

* *Combined-cycle gas turbines*

Source: Erdmann and Zweifel (2008)

1.2.2.2 Power Plant Operation

In the EU there is a liberalized market for electricity. The cost of fuel used to produce an additional unit of electricity in this setting makes up the marginal cost, which is negatively impacted by fuel costs and positively affected by fuel efficiency. The EU, on the other hand, levies CO_2 emission allowances on fossil fuels used in the electricity production process.

The European Power Exchange (EPEX) is an efficient and effective day-ahead electric power exchange that operates in in Europe, where the prices for electricity are determined through the interaction of supply and demand. Participants in the market submit their bids and offers for electricity on an hourly basis. The EPEX then aggregates these orders to create a supply and demand curve for each hour. The clearing price, at which the quantity of electricity is supplied matches the quantity demanded, is determined by averaging the orders for each hour. This process contributes to establishing a fair and transparent pricing mechanism for electricity trading in the European market.

1.2.2.3 Electricity Trading

The largest European power exchanges are Nord Pool, the European Energy Exchange, and Power Exchange Central Europe, each offering different products. The electricity market offers two types of contracts. The first type involves physical delivery of electricity, while the second requires cash-settlements. The main market for physical delivery of energy within the next 24 hours is the day-ahead market. Moreover, every market is divided into control areas, which helps to alleviate congestion of the power grid. That is, when there is a congestion in the flow of electricity within the grid, the prices in different control areas are calculated and adjusted to help balance the grid.

The intraday market serves to ensure a demand and supply equilibrium. It enables trading of electricity at prices that closely reflect real-time conditions such that equilibrium of supply and demand is achieved. In other words, this market provides a platform for participants to trade electricity on short notice, based on the most up-to-date information about the market dynamics. This helps ensure that the market responds efficiently to any fluctuations or changes in electricity supply or demand that occur during the day (Alexandridis & Zapranis, 2012).

1.3 Weather derivatives Market

Insomuch as energy and temperature are correlated, the energy industry relies on weather derivatives to hedge price and quantity risk of energy demand. Weather derivatives are contracts that have a duration, weather station, weather variable,

payoff function, premium, measurement agency, settlement agent, back-up station, settlement date, and an index, and energy companies are one the biggest users of these derivatives (J. C. Hull, 2003).

1.3.1 Weather Hedging Rationale

The weather derivatives market is a complement to the energy markets in that analysis of energy demand often factors in the unpredictability of weather conditions. In fact, in the microeconomic modeling of energy demand, aside from calendar effects, the business cycle, and fluctuations in income, temperature fluctuations are included in the short term factors affecting this demand (seeB.3). Empirically, Ranson et al. (2014) showed that energy consumption for heating and cooling purposes tends to be higher at both extremely low and extremely high temperatures, while it is relatively lower at moderate temperatures. This characteristic, is further supported by evidence from various studies provided by Ranson et al. (2014). Not only that, most of those studies predict two significant effects of climate change on energy demand. First, there is an expected reduction in the demand for heating energy, meaning that less energy will be required for heating purposes in certain regions. This is likely due to projected increases in average temperatures, particularly in colder climates, which would reduce the need for heating systems. Second, the studies predict a notable increase in demand for cooling energy. As climate change leads to higher temperatures, especially in warmer regions, there will be an elevated need for cooling systems to litigate the impacts of heatwaves, This higher demand for cooling is expected to offset, or even surpass, the reduced demand for heating in terms of overall energy consumption (De Cian et al., 2013; Dowling, 2013; Hamlet et al., 2010; Holland et al., 2011; Isaac & Van Vuuren, 2009).

Traded products on energy markets have unique physical features that expose them to weather changes as shown in Table 1.4. In contrast, managing weather risk can reduce different costs. For instance, a hydroelectric power generation company using weather derivatives reduces its profit fluctuations which lowers the borrowing interest rate, market valuation, and default risk.

Weather derivatives serve the same purpose as insurance contracts or catastrophe bonds(CAT)[1]. Yet, aside from contractual differences, weather derivatives require regular mark-to-market valuations. That is, Buckley et al. (2002) explain that a weather derivative can be treated as either a hedging instrument so that payoff is accounted for as an ordinary income, or as a derivative instrument such that payoff is accounted for as a capital gain or loss. The latter involves mark-to-market

[1]CAT bonds' payoff depends on a disaster's realization. Investors take on a *catastrophic* risk where the interest and the principal diminish if a particular category of *catastrophic* insurance claims exceed a certain amount.

accounting in that the security is considered a capital asset for which capital gain or loss depends on changes of its market value[2], and is taxed as such, whereas the former is subject to corporation tax. Additionally, insurance contracts are based solely on great catastrophes such hurricanes or floods, whereas weather derivatives are suitable for highly likely and frequent weather conditions. Researchers such as Field et al. (2012, Ch. 3) expound that the frequency, intensity, spacial extent, duration, and timing of weather extremes are altering due to climate change. Therefore, industries at risk are better-off using weather derivatives as they are more flexible.

Table 1.4: Weather exposures of industries

Industry	**Variable**	**Risk**
Agriculture	Temperature Rainfall	High crop loss
Airline	Wind Frost	Cancellation of flights high operational costs
Construction	Temperature Snowfall Rainfall	Delayed schedules
Energy	Temperature Rainfall	High crop loss
Hydroelectric power generation	Precipitation	Lower revenue during drought periods
Transportation	Wind Snowfall	Cancellation of shipping services

Source: Alexandridis and Zapranis (2012).

1.3.2 Risk Considerations

Weather derivatives are used to hedge against volumetric risks. In essence, there is a strong relationship between hedged volumes energy demand and supply and the underlying weather index. In Cao et al. (2003), it is shown that long-term supply is affected by warming trends. Hence, it is clear that weather derivatives serve the purpose of hedging against volume risk. Since the payoff of a weather derivative does not depend on the amount of money lost but on the weather index, the payoff will not compensate exactly for the money lost (Jewson & Brix, 2005). Weather risk is unique in that it is extremely localized and difficult to predict accurately despite advances in meteorological science. As a result, trading weather derivatives carries significant risk. Namely, basis risk and counter party default risk.

The introduction of exchange traded weather derivatives improved the transparency and fairness of this market. Nonetheless, weather contracts are based on a reference temperature in a specific measurement site such as Amsterdam, London,

[2]This applies only to exchange-traded derivatives. The capital gain or loss is, otherwise, accounted for when the derivative is expired or exercised.

and New York. Investors outside these cities face relative spatial basis risk. On the other hand the OTC market offers more accurate hedging with customized weather contracts, which are not priced fairly nor are they guaranteed against credit risk.

Spatial basis risk arises from the location of the weather station relative to the hedging company. Hence, the risk shrinks when financial loss is highly correlated with weather, and when contracts based on an optimal location are used for hedging.

1.3.3 Weather Derivatives Trading

Weather derivatives are traded on both primary and secondary markets, with primary market trades typically taking place OTC. In the secondary market, a considerable amount of trading is also conducted OTC through voice-brokers. However, due to the growing demand for weather derivatives, an organized market was established in the Chicago Mercantile Exchange (CME). The CME provides investors with standardized contracts that can be bought and sold more easily, providing more liquidity and transparency in the trading of weather derivatives.

1.3.3.1 Traded products

This work is centered around temperature derivatives. Degree days, average temperatures, and cumulative averages are the most commonly used indices in temperature-based contracts.

Degree Days The Heating Degree Day (HDD) index is the number of degrees of Daily Average Temperature (DAT) that are below a reference temperature, while Cooling Degree Days (CDD) index is the number of degrees of DAT that are above a reference temperature. The daily average temperature is calculated as follows

$$DAT_n := \overline{T}_n = \frac{T_n^{Max} + T_n^{Min}}{2}, \quad \forall n \in \{1, 2, \dots N\} \tag{1.1}$$

where T_n^{Max} and T_n^{Min} are the maximum and minimum temperatures, respectively. Therefore, HDDs are used to measure the demand for heating in winter periods, and CDDs are a measure of the demand for energy used for cooling in summer periods. That is, for each $n \in N$

$$HDD_n = \sum_{n=1}^{N} max(0, T_{ref} - DAT_n) \tag{1.2}$$

$$CDD_n = \sum_{n=1}^{N} max(0, DAT_n - T_{ref}) \tag{1.3}$$

For a given day, summing up its respective HDDs and CDDs yields the magnitude of the deviation of its average temperature from the baseline (Jewson & Brix, 2005).

The cumulative average temperature (CAT) index is mainly used in Europe during summer periods. It is the sum of DATs over the period of a contract

$$CAT_n = \sum_{n=1}^{N} T_n \tag{1.4}$$

1.3.3.2 Pay-off functions

Weather derivatives are not different from other classes of derivatives as they can be based on calls, puts, etc. The pay-off of a weather derivative depends on the strike price K, the tick size α, ϑ is the limit defined in currency terms and C and F are the cap and floor estimated in units of the temperature index. The buyer of an option receives a payoff ϕ which is a function of the cumulative index over the contract period P

$$\phi = f(DD) \tag{1.5}$$

where heating and cooling degree seasons are defined as

$$DD = HDD^P = \sum_{n \in P} HDD_n \tag{1.6}$$

$$DD = CDD^P = \sum_{n \in P} CDD_n \tag{1.7}$$

Contracts traded OTC are usually capped by an upper limit C. For instance, the payoff of a call option with cap, given a payoff rate α is

$$\phi = \begin{cases} 0 & \text{if } DD < K \\ \alpha(DD - K) & \text{if } K \leq DD \leq C \\ \vartheta & \text{if } DD > C \end{cases} \tag{1.8}$$

Or, written more succinctly as

$$\phi = min\{\alpha(DD - K)^+, \vartheta\} \tag{1.9}$$

Therefore, the option is at the money (ATM) when K is near the historical average of DD, and out of the money (OTM) when K is far from this average. Similarly the payoff of a put option given a floor F is given by

$$\phi = \begin{cases} \vartheta & \text{if } DD < F \\ \alpha(K - DD) & \text{if } F \leq DD \leq K \\ 0 & \text{if } DD > K \end{cases} \tag{1.10}$$

Or, written concisely as

$$\phi = \min\{\alpha(K - DD)^+, \vartheta\} \tag{1.11}$$

This pay-off function has an economic value to the buyer in that it protects against very low values of the index. Additionally, Jewson and Brix (2005) explain that collars are structured as a combination of a long call and short put with the following pay-offs where one longs the call option with a cap and shorts the put option with a floor.

$$\phi = \begin{cases} -\vartheta & \text{if } \phi < C \\ \alpha(DD - K_c) & \text{if } C \leq DD \leq C \\ 0 & \text{if } K_c \leq DD \leq K_p \\ \alpha(K_p - DD) & \text{if } K_p \leq DD \leq F \\ \vartheta & \text{if } DD > F \end{cases} \tag{1.12}$$

which translates into

$$\phi = min\{\alpha(DD - K_c)^+, -\vartheta\} + min\{\alpha(K_p - DD)^+, \vartheta\} \tag{1.13}$$

Literature Review

The research on weather derivatives focuses on two key apprehensions, pricing and risk.

2.1 Weather Derivatives Pricing

As regards the fair pricing of weather derivatives, there is no model that offers a closed form formula for pricing. However, Jewson and Brix (2005) distinguish between three main pricing approaches; historical burn analysis, index modelling which is an extension of the first, and daily simulation. Burn analysis is an actuarial method that consists of assessing how a contract would have performed in past years. The approach determines the payoff of a weather derivative through accumulating degree days over a different years and estimates the price as the average annual payoffs. Some researchers such as (Davis, 2001; Dorfleitner & Wimmer, 2010; Geman & Leonardi, 2005) argue that modeling indices directly such as the HDD, CDD, or CAT index is more accurate. This methodology estimates the distribution of the index, and tests the hypothesis that the observations originate from such a distribution. If the hypothesis is not rejected, the distribution can be used to represent the index distribution (Jewson & Brix, 2005). While Geman and Leonardi (2005) find that the hypothesis of normality is rejected for December HDDs in Paris and accepted for December AccHDD indices, Jewson and Brix (2005) explain that there is a high probability that the wrong distribution is used, because of the absence of theory on what distribution is appropriate to fit the indices. In daily simulations, the models directly simulate the future behaviour of temperature. These dynamic models show greater potential than the two other methods. In fact, according to Geman and Leonardi (2005), modeling daily average temperature has an overall higher potential accuracy in that it allows, for instance, a more accurate extrapolation of extremes, and an easier incorporation of meteorological forecasts in the pricing process. In this framework, either discrete or continuous processes can be used. Moreno (2000) claims that since temperature values are recorded in discrete form, temperature should be modeled with discrete processes. He conducted a numerical comparison of a discrete mean-reverting process and an autoregressive process as suggested by Carmona (1999) obtaining an acceptable goodness of fit for both models, whereas

Bob (1998a, 1998b) is the first to suggest working with continuous stochastic models. These models are often sub-Ornstein–Uhlenbeck equations of a more general mean-reverting process in order to account for some essential attributes of temperature. Modeling temperature data using continuous processes offers powerful forecasting tools. In an approach similar to J. Hull and White (1990) research conducted by Dornier and Querel (2000) uses a modified continuous Ornstein-Uhlenbeck stochastic process that reverts back to historical mean temperature by adding a term for changes in seasonal variations, Alaton et al. (2002) use the same approach while implementing a sinusoid function to model average historical daily temperatures through incorporating seasonalities in the mean, whereas Brody et al. (2002) suggest including a fractional Brownian motion in the Ornstein-Uhlenbeck process. More recently, Zapranis and Alexandridis (2008) use neural networks to estimate the parameters in a time dependent speed of mean reversion in the Ornstein-Uhlenbeck process. Alaton et al. (2002) work on Swedish data and get a satisfactory fit of the model for the data. Nonetheless, Alexandridis and Zapranis (2012) say that it is a simplification of the real world which could lead to mispricing of options. In fact, as Benth and Šaltytė-Benth (2005) noted, Alaton et al. (2002) observe that the temperature differences are approximately normally distributed, but didn't test for it. Moreover, the normality hypothesis got rejected when Benth and Šaltytė-Benth (2005) examine Norwegian temperature data, and they subsequently put forth using a Lévy-based OU process which incorporates a Lévy noise instead of a Brownian motion as Dornier and Querel (2000) do. Some other authors such as Cao and Wei (2000) apply an econometric approach where they correct for trend and seasonality in temperature and propose a discrete auto-regressive of the temperature residuals. Similarly to the continuous process that Brody et al. (2002) suggest, results from Caballero et al. (2002) show long-range correlation in the autocorrelation function (ACF) of the temperature, therefore they propose modelling this dependence temperature time series with ARMA processes. In their study, the long memory in the ACF is captured with an autoregressive fractional integrated moving average model (ARFIMA)[1]. Additionally, Caballero et al. (2002) use these processes on the UK's DATs and find that ARFIMA models perform effectively in terms of autocovariance. Nevertheless, Jewson and Caballero (2003) argue that the ARFIMA model doesn't capture the seasonality in the ACF of temperature, and apply a special AR process called autoregressive on moving average (AROMA) to DATs. The regression analysis involves using multiple moving averages of previously detrended and deseasonalized temperatures to predict the detrended and deseasonlized temperature at time t. This model which further incorporates seasonality (SAROMA) also suggests a strong correlation of between temperature observations at multiple lags. Moreover, some researchers like Benth et al. (2007) combined the econometric

[1]This is a process that generalize ARIMA models proposed by Box et al. (2015).

approach and the continuous processes by using a continuous-time auto-regressive process.

It is clear that fairly pricing temperature derivatives in Netherlands poses a challenge in terms of selecting the most appropriate model. Nonetheless, some authors such as Tindall (2006) and Esunge and Njong (2020) combined multiple approaches from above to model the speed of mean reversion and volatility in the pricing of weather derivatives for Sydney, New York, Seattle, and Cincinnati.

2.2 Basis Risk and Credit Risk

In terms of risk concerns, hedgers using weather derivatives typically face two types of risk, namely credit risk and basis risk. Exchange traded weather derivatives mitigate the default risk of the counter-party. However, since contracts are not necessarily written in the exact location the hedger wishes to cover, they often need to bear a basis risk. This risk is of main concern, according to Manfredo and Richards (2009), the effectiveness of a hedge is contingent on the behaviour of the basis. Given the sensitive aspect of the underlying *asset*, basis risk is expected to become virtually nil upon hedging. For instance, D'arcy and France (1992) analyse hedging basis risk for insurance derivatives, namely futures based on catastrophe losses. Their research about insurance futures contracts highlights the solvency concerns that arise with insurance futures contracts, which are a comparable asset to weather derivatives. Furthermore, they proposed a an index for an insurance future that reduces risk substantially for insurers. Moreover, Major (1999) analyse the basis risk between catastrophe futures and portfolios of insured homeowners' building risks running in the context of hurricane danger. Nevertheless, weather derivatives are a relatively new financial instrument[2] without extensive research on the topic of basis risk associated with trading these instruments. With regard to rainfall derivatives, Ritter et al. (2014) investigate the effectiveness of regional diversification in reducing geographical basis risk of such products. Their analysis concludes that inferring optimal portfolio weights from a multi-site rainfall model performs better than other approaches, such as inverse distance weighting, in reducing spatial basis risk. Yang et al. (2009) explore the basis risk of Heating Degree Day (HDD) and Cooling Degree Day (CDD) weather derivatives and their hedging effectiveness in the United States' energy market. Empirically, they find that although the RMS regional and CME city weather indexes yield comparable result, risk managers could use the former as a complement to the exchange-listed indexes. Additionally, Manfredo and Richards (2009) argue that the choice of the weather station is "*less critical in managing*

[2]Weather derivatives started trading OTC in 1997. Their popularity drove the Chicago Mercantile Exchange to introduce the first exchange-traded weather futures contracts in 1999. These can be traded electronically on the CME's GLOBEX 2 system.

basis risk". Instead, they suggest that non-linear-payoff weather derivatives could be used solely or combined with linear payoff instruments to minimize the basis risk between crop yields and weather. To counter the basis risk, hedgers can either buy basis derivatives that rely on the disparity between the weather indices of two distinct locations, or buy a location specific OTC contract.

In both cases, OTC trading involves taking into consideration the probability of the counter-party failing to perform the contractual obligation, i.e. credit risk. The predictability of default has been modeled by many researchers. In fact, advances in artificial intelligence have made it easier to predict the default, Brockett et al. (1994) use a back-propagation learning algorithm based on a feed forward network to analyse the insolvency probabilities of insurers in the U.S. Furthermore, Cummins and Mahul (2003) use expected utility theory and add to it through working on an optimal hedging strategy when the insurer and the buyer have divergent beliefs about the probability of total default of the insurer.

It is clear that there is a trade-off between basis risk and credit risk when the indexed temperature and geographical location do not perfectly overlap. In, fact Golden et al. (2007) study the efficiency of combining exchange-traded weather derivatives and OTC weather derivatives to hedge spatial basis risk while accounting for credit risk. To this end, they test the effectiveness of hedging ratios using both linear and non linear hedging approaches. Their results suggest that with low basis risk and high weather risk both linear and non-linear hedging techniques are highly effective.

Overall, the literature behind temperature derivatives is relatively modest and most of it revolves around the US market. Up until now, research is centered around deriving accurate pricing formulas while the effect of geographical basis risk is often neglected. This thesis somewhat relies on the approach of Alaton et al. (2002) since the model they propose describes essential DAT features. That is, the thesis contributes to the literature by conducting the research on temperature in the Netherlands. In addition, the thesis adds to the literature by taking a different path with calculating the market price of risk, which is estimated through a root finding method. Moreover, basis risk is is deeply analysed and incorporated in the profit maximization of both energy importers and exporters.

Methodology and Data

The methods, data collection, and processing processes described bellow are applied to the cities of Amsterdam, Utrecht, Rotterdam, Eindhoven, Groningen, Maastricht, and Tilburg/Breda[1]. Daily values from the 1st of January 1971 to the 10th of February 2023 were extracted from the sources presented in section 3.2.

These cities were chosen based on population size and business attractiveness. Nonetheless, The Hague, the third biggest Dutch city, is excluded from this study because there is insufficient data to analyse.

3.1 Methodology

As discussed in the literature review, several approaches may be adopted to successfully price temperature derivatives. It was also indicated that daily modeling is more appropriate than other methods. Therefore, with reference to Jewson and Brix (2005), this research will model Daily Average Temperatures (DATs) in several steps.

3.1.1 Seasonality Modeling

In this subsection, potential sources of trends are discussed and identified. Moreover, classical detrending methods but also more recent ones, such as Multiple Seasonal-Trend decomposition using Loess (MSTL) and Facebook's Prophet, are used to detrend and remove seasonality from the dataset and compare their performance in doing so. In a second phase, the thesis will detrend and model the seasonal component as Fourier series, which is believed to yield more consistent results.

3.1.2 Temperature Modeling

Two main approaches exist to accomplish this task successfully. The first one is a time series technique that relies on models such as the Autoregressive Moving Average Model (ARMA), while the second one stochastically represents temperature as a mean reverting Ornstein-Uhlenbeck (OU) process. The stochastic representation uses a sub-formula of the OU process which is adjusted to fit assumptions about

[1]Data is recorded by KNMI's station Gilze-Rijen which lies at approximately 10km from both cities.

the underlying. Hence, this twist will be thoroughly discussed to understand the dynamics involved in the modeling.

3.1.3 Volatility Modeling

Similarly to seasonality, several papers in the literature use Fourier series to model DAT's volatility. Nevertheless, the thesis will use a comparative approach in that other methods involving parametric, local and non-parametric regression, and piecewise constant functions will be compared to each other, which will allow investigating the underlying from the perspective of different studies in the literature.

3.1.4 Temperature Options Pricing

The statistical analysis used in the previous chapters will be used to fairly price temperature options for each city in the Netherlands. To do so, some of the pricing methods that are presented in the literature will be applied to the Dutch market, namely

- Alaton approach with an Ornstein-Uhlenbeck process suggested by Alaton et al. (2002)
- Monte Carlo simulation with Ornstein-Uhlenbeck process

3.1.5 Risk Inspection

Eventually, an important part of the study will be dedicated to looking into basis risk in the Netherlands. To this end, the research will examine correlations between the detrended data and the relationship between respective DATs and the long-term historical DAT.

3.2 Data

Historical weather data is available and will be retrieved from the National Centers for Environmental Information which is a database of the National Oceanic and Atmospheric Administration, but also from the *Koninklijk Nederlands Meteorologisch Instituut* (KNMI) which is the Netherland's meteorological institute.

The reason both databases will be used is because the latter requires extra data transformation to adjust the values which is computationally expensive. Consequently, it will be used only for complementary reasons.

For every city, raw data comprises daily values for

- Maximum temperature of the day T_n^{Max}
- Minimum temperature of the day T_n^{Min}

The first step in the research consists in calculating the daily average temperature (DAT) for all cities according to the definition of HDDs and CDDs

$$DAT_n := T_n = \frac{T_n^{Max} + T_n^{Min}}{2}, \quad \forall n \in \{1, 2, \dots N\} \tag{3.1}$$

this new variable is our underlying *"asset"*.

Data Cleaning and Visualisation

The development of the weather derivatives market is contingent on continuously and accurately providing the data by national meteorological services (NMSs). Nevertheless, aside from the Swiss Meteorological Service, NMSs rarely check for data quality since they are mainly tasked with making forecasts.

4.1 Pre-processing

4.1.1 Missing Values

Each of the data subsets had missing maximum and minimum temperatures for the same two dates. These dates are 2016 October 30 and 2019 March 18. Therefore, the method suggested in Alexandridis and Zapranis (2012) was applied to fill the missing data.

Let T_n be the missing temperature on day n. First, $\overline{T}_n^y$, the average temperature for day n is calculated across the N years is as follows

$$\overline{T}_n^y = \frac{1}{N}\sum_{y=1}^{N} T_n^y \tag{4.1}$$

Second, $\overline{T}_n^d$, the average temperature of seven days before and after the day with missing data is calculated as follows

$$\overline{T}_n^d = \frac{\sum_{d=1}^{7} T_{n-d} \sum_{d=1}^{7} T_{n+d}}{14} \tag{4.2}$$

Subsequently, the missing value, $\widetilde{T}_n$, is replaced by the average of both averages as follows

$$\widetilde{T}_n = \frac{\overline{T}_n^y + \overline{T}_n^d}{2} \tag{4.3}$$

4.1.2 Erroneous Values

A common anomaly in temperature data is finding implausible values. If these absurd values are left uncorrected it results in significant mispricing of weather

derivatives.The following verification was applied to each subset to ensure the quality of the data:

- Checking that daily maximums are not lower than daily minimums
- Checking that daily temperatures are within acceptable limits for the given time of year and weather station

4.2 Underlying Dynamics

4.2.1 Warm and Cold Periods

In order to get a better understanding of the underlying's dynamics, Figure 4.1 shows DATs' behaviour in the last 13 years. The oscillatory character of DATs depicted in Figure 4.1 shows that the underlying has a seasonal cycle of one year. Moreover, descriptive statistics of temperature is provided in Table A.2. The Table shows roughly similar metrics for all the cities. All cities have platykurtic and negatively skewed DATs that fluctuate between -16.50(°C) and 31.50(°C), approximately. Nevertheless, the measurements in Table A.2 describe a period that extends from the 1st of January 1971 to the 10th of February 2023, which conceals subtle monthly and seasonal differences which are discussed in the following subsection.

Figure 4.1: Daily average temperatures

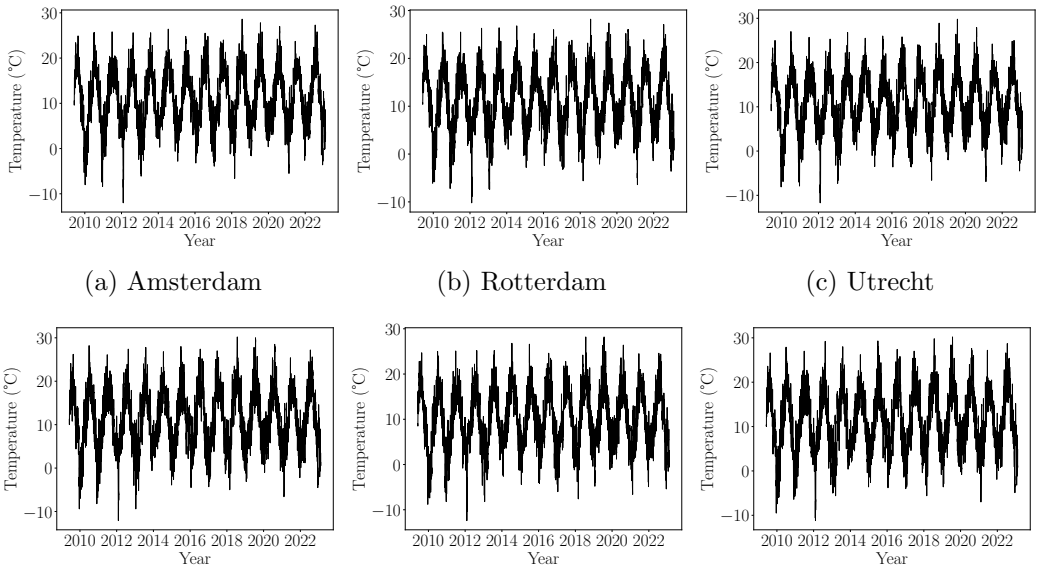

(a) Amsterdam (b) Rotterdam (c) Utrecht

(d) Eindhoven (e) Groningen (f) Maastricht

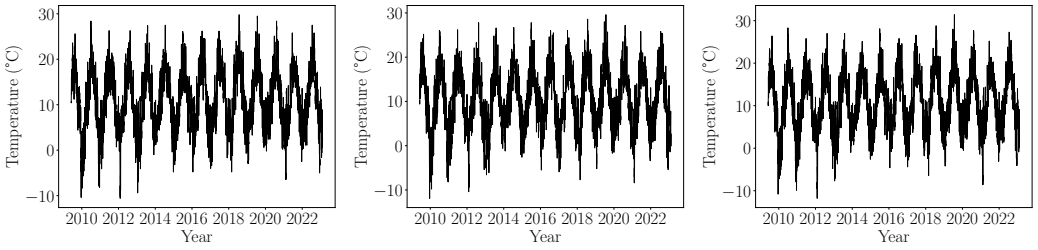

(g) Tilburg (h) Enschede (i) Nijmegen

A closer inspection of Figure 4.2 indicates that the cold period in the cities extends from October to April while the warm period begins in May and ends in September. However, Figure 4.3 presents a more comprehensive evaluation of temperature through a hierarchical ascendant classification which reveals a seasonal discrepancy between Dutch cities. The figure shows that in Utrecht, Eindhoven, Maastricht, Tilburg, Enschede, and Nijmegen the winter period starts in October and ends in April and the summer period begins in May and concludes in September, while the cities of Amsterdam, Rotterdam, and Groningen experience a lag in that the cold season goes from November to April and the warm season lasts between May and October.

The lag between summer and winter periods[1] in the Netherlands suggests that hedging temperature risk using Amsterdam's temperature derivatives may be unreliable for other cities.

Figure 4.2: Monthly average temperature

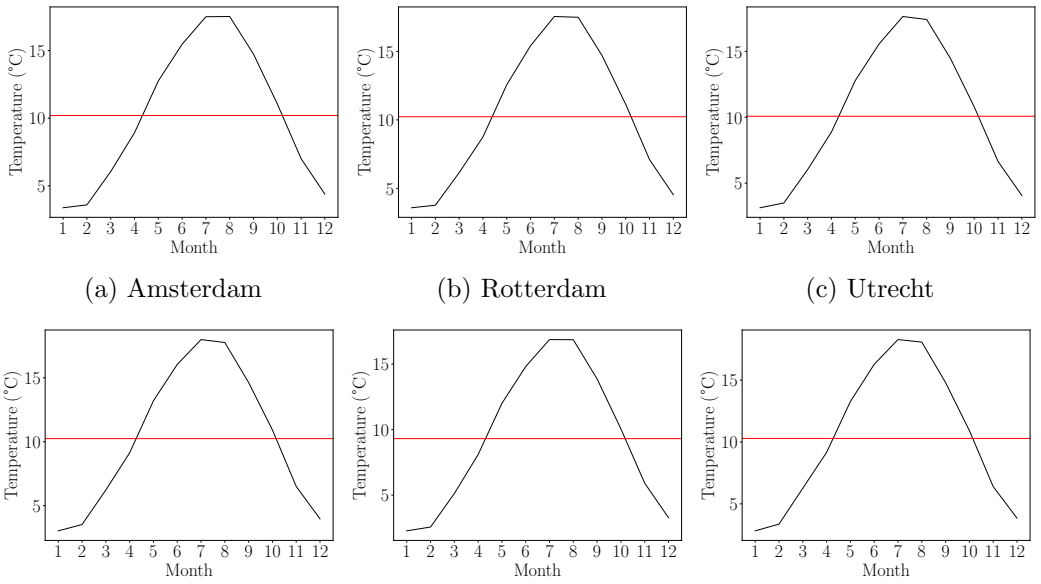

(a) Amsterdam (b) Rotterdam (c) Utrecht

(d) Eindhoven (e) Groningen (f) Maastricht

[1]Warm/cold is interchangeably used with summer/winter in this chapter

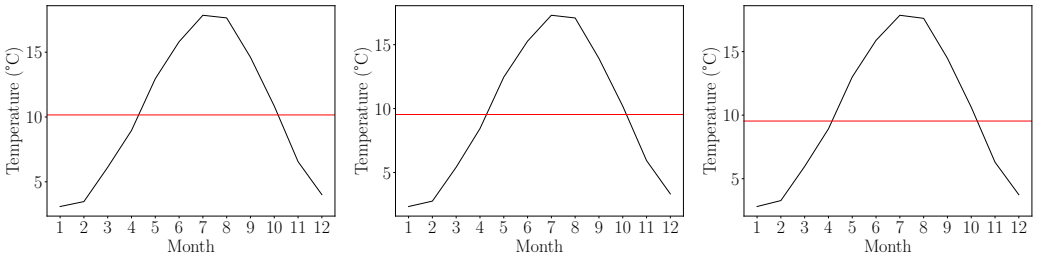

(g) Tilburg (h) Enschede (i) Nijmegen

Figure 4.3: Hierarchical classification of monthly DATs

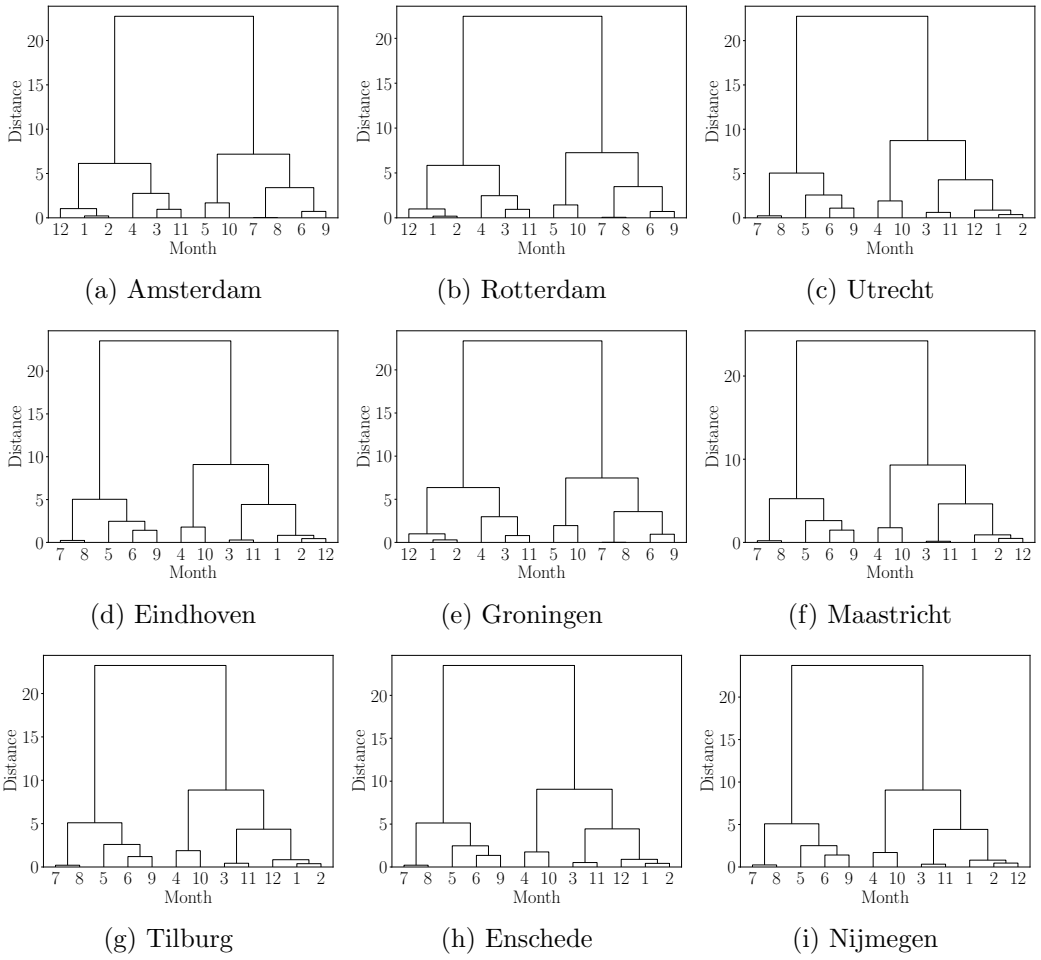

(a) Amsterdam (b) Rotterdam (c) Utrecht

(d) Eindhoven (e) Groningen (f) Maastricht

(g) Tilburg (h) Enschede (i) Nijmegen

The lag between summer and winter periods[2] in the Netherlands suggests that hedging temperature risk using Amsterdam's temperature derivatives may be unreliable for other cities. That is, on average, since 1971, hedgers based in Nijmegen or speculators using Eindhoven's temperature data are worse-off trading Schipol's HDDs or CATs because of the non-overlapping periods.

[2]Warm/cold is interchangeably used with summer/winter in this chapter

4.2.2 Bimodal Distribution

Much of what could be said about DATs' second and third standardised moments is contained in the Jarque–Bera (JB) test results. The JB test is particularly useful in financial and econometric analysis to check the assumption of normality in a dataset where

- H_0: The sample comes from a normal distribution with unknown mean and variance
- H_1: The sample does not come from a normal distribution

The JB statistic is defined as

$$\text{Jarque Bera} = \frac{n}{6}\left(s^2 + \frac{(k-3)^2}{4}\right) \tag{4.4}$$

Where n is the sample size, s is the sample skewness, and k is the sample kurtosis. Equation 4.4 indicates that the JB test compares the observed skewness and kurtosis of the data to what would be expected under a normal distribution. Hence, for all cities, the DAT's JB statistic is greater than 163 and the p-values are all zero, therefore strongly rejecting the null hypothesis. In fact, Figure 4.4 confirms that DATs have a bimodal distribution which derives from the sinusoidal character of the data. That is, the two means are centered and reflect the peaks of the summer and winter periods discussed in subsection 4.2.1 which is highlighted in Figure 4.5.

Figure 4.4: Distribution of the DATs

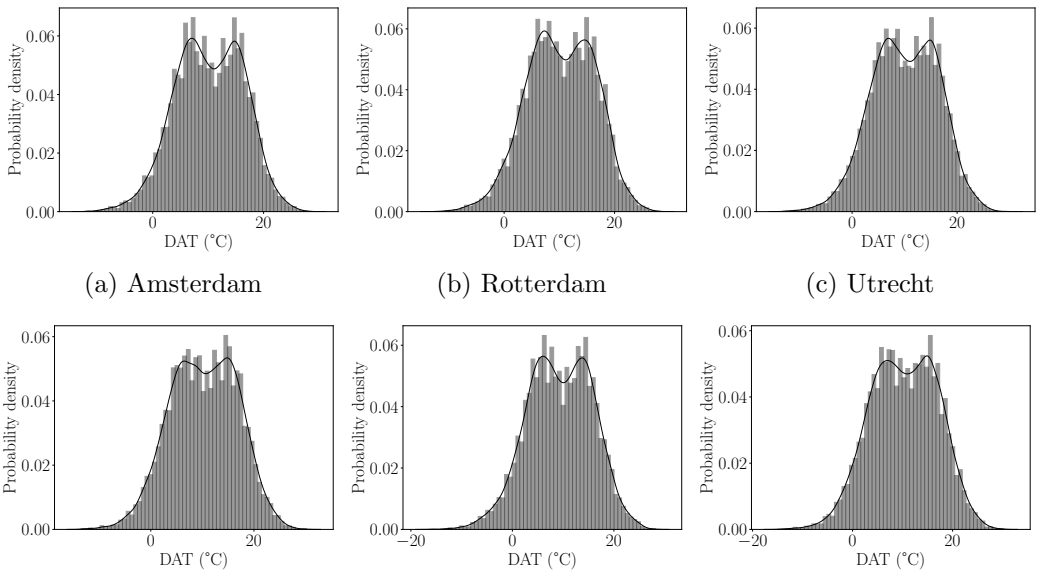

(a) Amsterdam (b) Rotterdam (c) Utrecht

(d) Eindhoven (e) Groningen (f) Maastricht

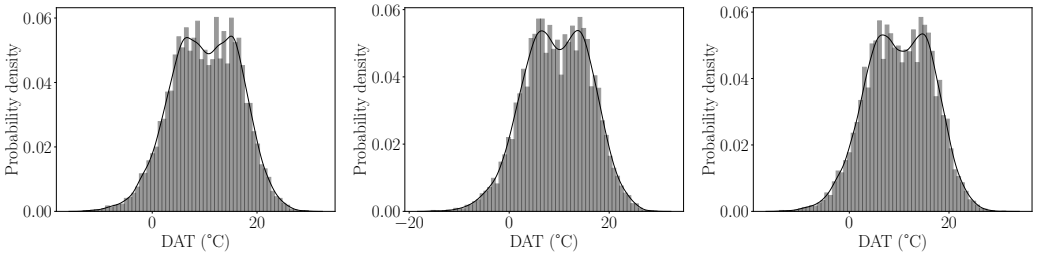

(g) Tilburg (h) Enschede (i) Nijmegen

In Figure 4.5, it is noticeable that, to a different degree, all the cities have a slightly negatively skewed cold period temperature and slightly positively skewed warm period temperature. Additionally, examining the three standards moments of the DATs, it is evident from Figure 4.6 that the cold period is characterized by a higher standard deviation than the warm period, which is in line with the results of Benth and Šaltytė-Benth (2005), Benth et al. (2007), Benth and Benth (2007), and Zapranis and Alexandridis (2008).

Figure 4.5: Distribution of winter and summer periods

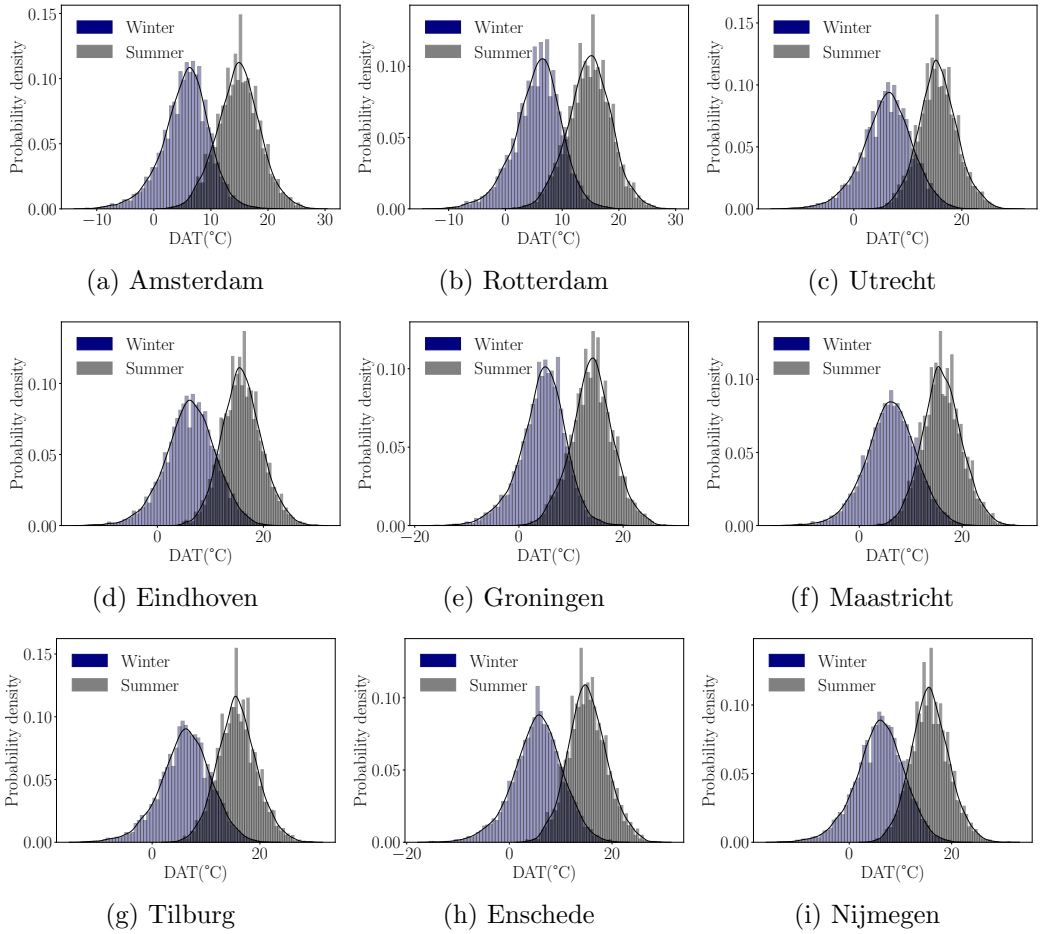

(a) Amsterdam (b) Rotterdam (c) Utrecht

(d) Eindhoven (e) Groningen (f) Maastricht

(g) Tilburg (h) Enschede (i) Nijmegen

Figure 4.6: Standard deviation of the DATs

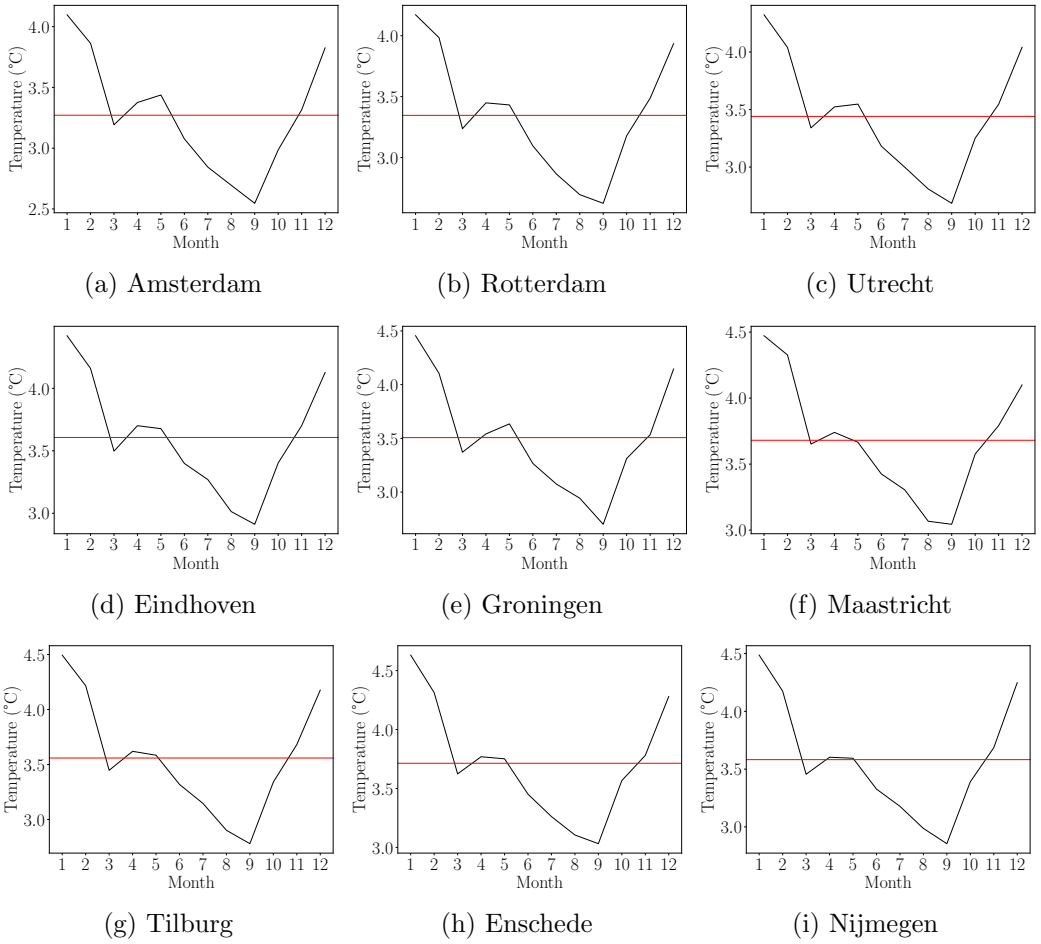

(a) Amsterdam (b) Rotterdam (c) Utrecht

(d) Eindhoven (e) Groningen (f) Maastricht

(g) Tilburg (h) Enschede (i) Nijmegen

As regards the third moment of the DATs, observing Figure 4.7, it is shown that skewness is higher in summer periods and lower in winter periods. Naturally, warm periods have their most extreme temperatures above the mean, while the coldest temperatures are below it. Nevertheless, an interesting fact that could be derived from the skewness representations in Figure 4.7 is that, in the summer, there is an increased likelihood of experiencing warmer days compared to the average, while in winter periods, there a higher probability of having days colder than the average (Bellini, 2005).

Figure 4.7: Skewness of the DATs

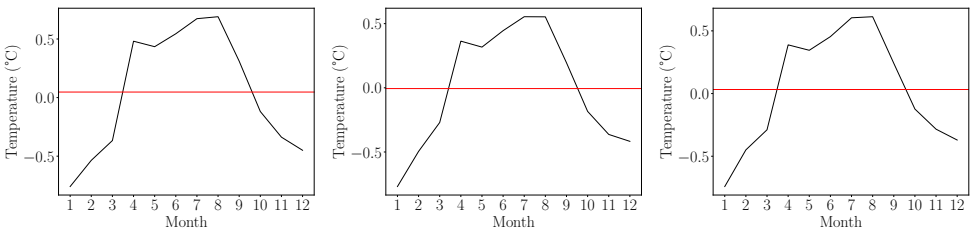

(a) Amsterdam (b) Rotterdam (c) Utrecht

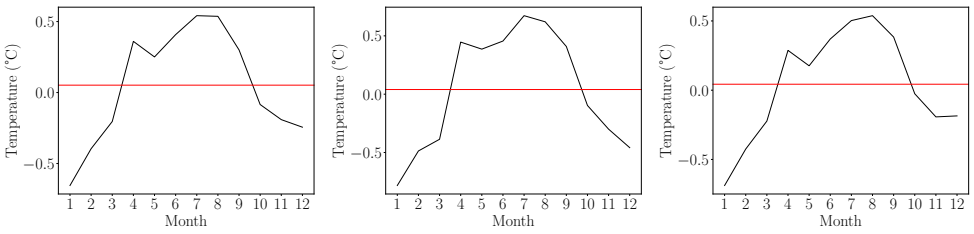

(d) Eindhoven (e) Groningen (f) Maastricht

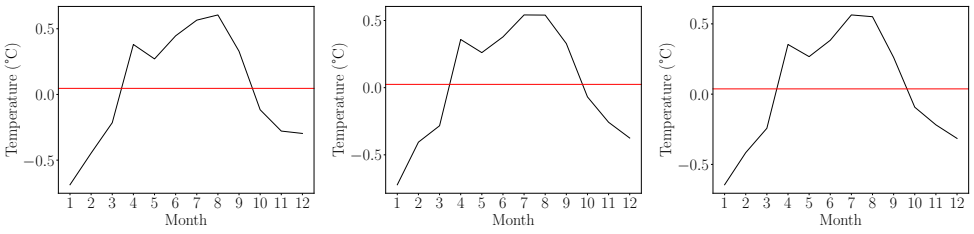

(g) Tilburg (h) Enschede (i) Nijmegen

4.3 Stationarity and Autocorrelation

When it comes to temperature data, there is an ambiguity about the stationarity of the time series. That is Table 4.1 presents the results of the augmented Dickey Fuller (ADF) test applied to models of the form

$$\Delta y_t = \alpha + \beta t + \delta y_{t-1} + \sum_{i=1}^{p-1} \delta_i \Delta y_{t-i} + \varepsilon_t \tag{4.5}$$

where y is temperature data indexed daily with t, α is a constant, β is the trend coefficient, and p is the lag order of the autoregressive process. Th regression results remain the same if a quadratic trend term is included or if $\alpha = 0$ and/or $\beta = 0$. The hypotheses of the ADF test are defined as

- $H_0 := \gamma = 0$: The time series contains a unit root and is non-stationary.
- $H_1 := \gamma < 0$: The time series is stationary of trend-stationary.

where the test statistic of the ADF tests is defined as

$$\frac{\hat{\gamma}}{\text{SE}}$$

To determine the length of p, the test algorithm includes the number of lags that minimises the Akaike Information Criterion (AIC) defined as

$$2k - \ln(L) \tag{4.6}$$

where k is the number of lagged differences included in Equation 4.5 and L is the maximized value of the likelihood function of the model. The results in Table 4.1 all show that the critical values are greater than the test statistic with p-values equal to zero, hence it indicates that it is extremely likely that the temperature data is stationary in all cities.

Table 4.1: Augmented Dickey Fuller test results summary

	ADF statistic	**Critical values**			p**-value**
		1%	**5%**	**10%**	
Amsterdam	−9.12	−2.57	−1.94	−1.62	0.000
Rotterdam	−4.26	−2.57	−1.94	−1.62	0.000
Utrecht	−4.43	−2.57	−1.94	−1.62	0.000
Eindhoven	−4.53	−2.57	−1.94	−1.62	0.000
Groningen	−4.82	−2.57	−1.94	−1.62	0.000
Maastricht	−4.63	−2.57	−1.94	−1.62	0.000
Tilburg	−4.53	−2.57	−1.94	−1.62	0.000
Enschede	4.60	−2.57	−1.94	−1.62	0.000
Nijmegen	−4.64	−2.57	−1.94	−1.62	0.000

Note: $\alpha = \beta = 0$

The results in Table 4.1 do not confirm that temperature is a stationary process *per se*, it would be counter-intuitive to assume so. That is, it is clear from the analysis so far that daily and monthly temperature have a cyclic behaviour, therefore it is highly unlikely that, even in the absence of a long-term trend, temperature is stationary. In fact, the distribution of temperature is not the same across months and seasons, hence it could be inferred that the joint distribution shifts in time. On the other hand let Amsterdam's daily average temperature be represented by a matrix of the form

$$\begin{bmatrix} a_{1,1} & a_{1,2} & \cdots & a_{1,n} \\ a_{2,1} & a_{2,2} & \cdots & a_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{m,1} & a_{m,2} & \cdots & a_{m,n} \end{bmatrix}$$

where m and n correspond are indexes of the year and day, respectively. In this case it could be argued that temperature represented in the vector $[\, a_{1,1} \;\; a_{2,1} \;\; \cdots \;\; a_{m,1} \,]^{\mathrm{T}}$ is stationary. The ambiguity described above corresponds to the phenomena of cyclostationarity, where temperature is described by multiple interleaved stationary processes. More specifically, subsequent chapters discuss temperature time series from a multivariate signal processing point of view. That is, this research is about nine time series where multi-variate information in every single one is of interest. Conditional modelling of multivariate time series helps with extracting features common to temperatures of all cities, and to price temperature derivatives as accurately as possible, one needs to fully make use of the information contained in these signals.

Concerning serial correlation of the time series, let $\{y_t\}_{t\geq 0}$ be a moving average (MA) of order p with respect to a white noise $\{w_t\}_{t\geq 0}$, then

$$y \sim MA(p) \coloneqq y_t = w_t + \beta_1 w_{t-1} + \sum_{i=2}^{p} \beta_i w_{t-i} \tag{4.7}$$

for some real numbers $\{\beta_1, \cdots, \beta_p\}$. A main property of the MA process above is that the autocorrelation function (ACF) vanishes for lags greater than p. However, Figure 4.8 shows that this is not exactly the case, the serial correlation does not disappear after a finite number of lags.

Figure 4.8: Autocorrelation of DATs

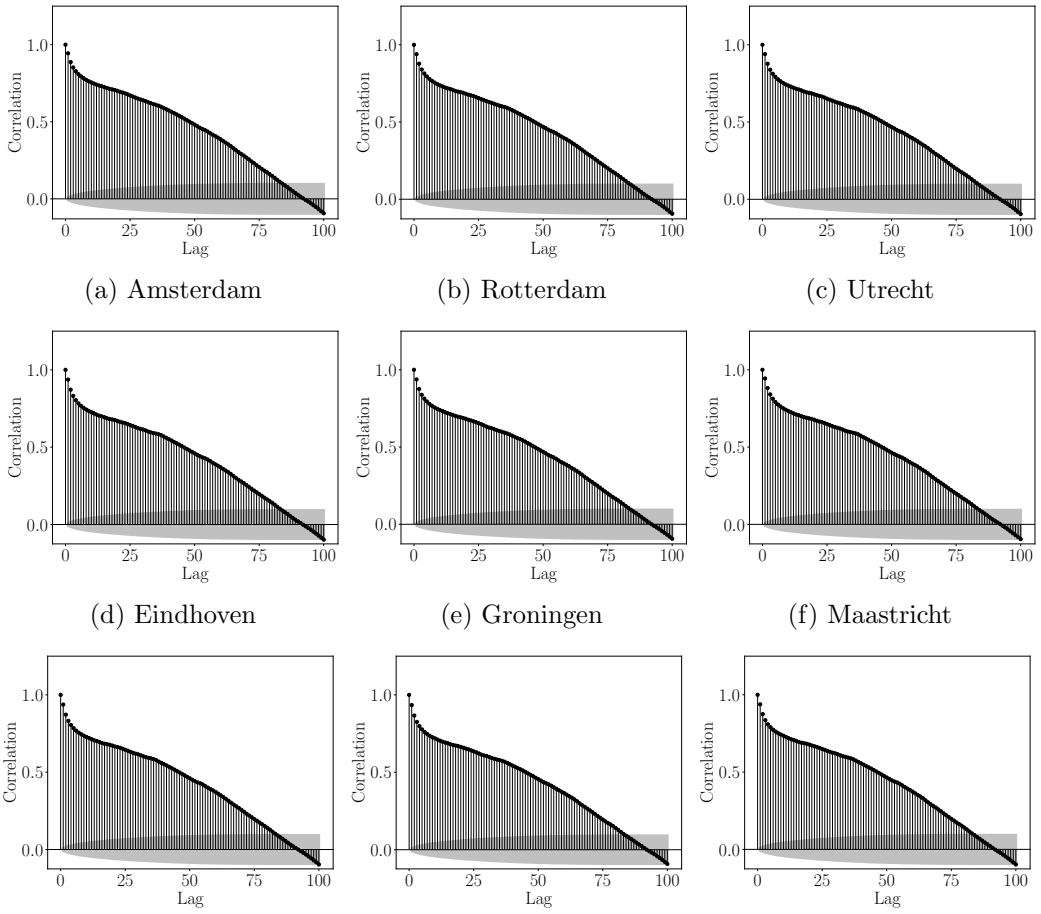

(a) Amsterdam (b) Rotterdam (c) Utrecht

(d) Eindhoven (e) Groningen (f) Maastricht

(g) Tilburg (h) Enschede (i) Nijmegen

Instead of fitting a MA model to the data, attention is shifted to autoregressive (AR) model of the form

$$y \sim AR(q) \coloneqq y_t = w_t + \gamma_1 y_{t-1} + \sum_{i=2}^{q} \gamma_i y_{t-i} \tag{4.8}$$

where $\{y_t\}_{t\geq 0}$ is an AR process of order q with respect to a white noise $\{w_t\}_{t\geq 0}$, for some real numbers $\{\gamma_1, \gamma_2, \cdots, \gamma_q\}$. In this case the partial-autocorrelation function (PACF) is used to find the lag order of the AR model as presented in Figure 4.9 . In fact, the PACF shows only the relationship between two observations in time with relationships of intervening observations removed. Mathematically, given the time series $\{y_t\}_{t\geq 0}$, the partial autocorrelation of lag s is the correlation of y_t and y_{t+s}, with the linear dependence of y_t through y_{t+s-1} on y_{t+1} not accounted for. For

instance, this property is depicted through suppressing the effects of the first two lags in Figure 4.8, where the partial autocorrelation of the third lag is as conveyed by Figure 4.9. Essentially, already found variation is removed, before finding the next correlation.

To find the order of the model defined in Equation 4.8 the AIC criterion introduced in Equation 4.6 is used which, once again, finds a parsimonious balance between best fit and the number of lags needed.

Figure 4.9: Partial autocorrelation of DATs

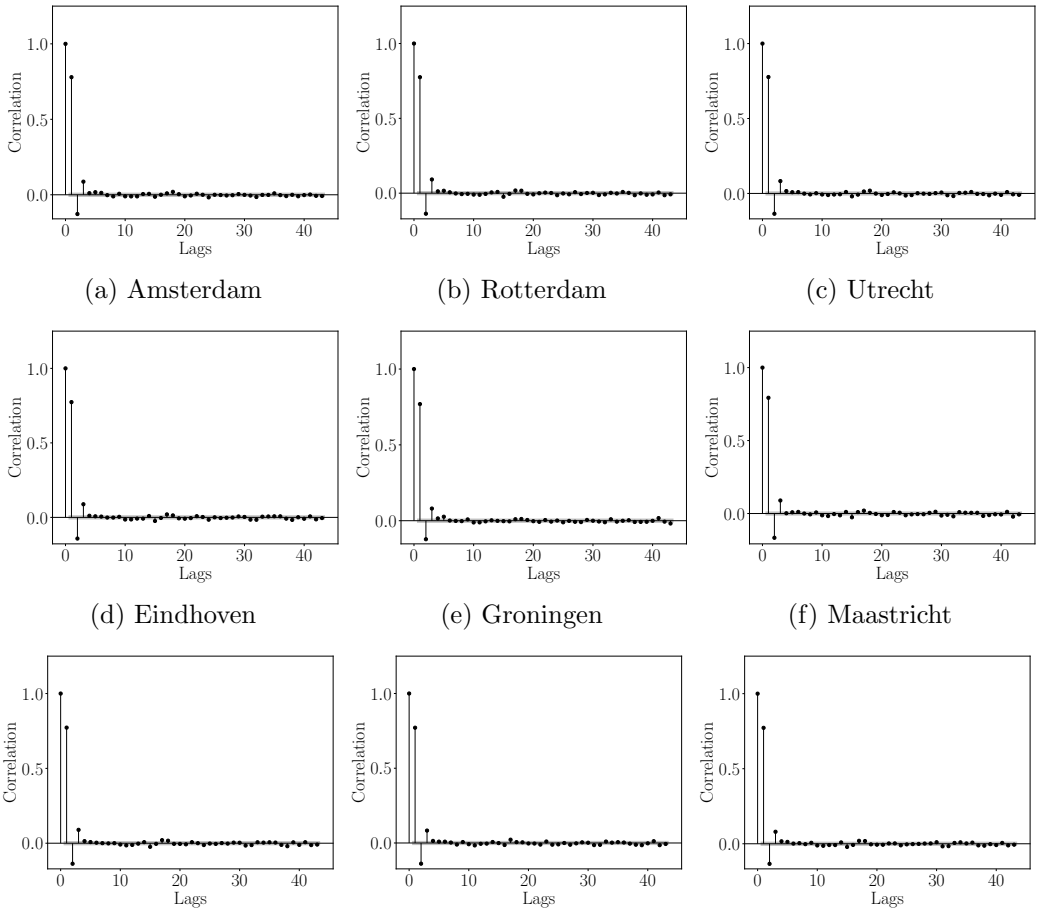

(a) Amsterdam (b) Rotterdam (c) Utrecht

(d) Eindhoven (e) Groningen (f) Maastricht

(g) Tilburg (h) Enschede (i) Nimegen

As displayed in Figure 4.10, Beyond 17 lags the AIC is approximately constant, hence it is optimal to fit an AR(7) model to the time series. The results of this regression are presented in Table A.3.

Figure 4.10: Autoregressive model order

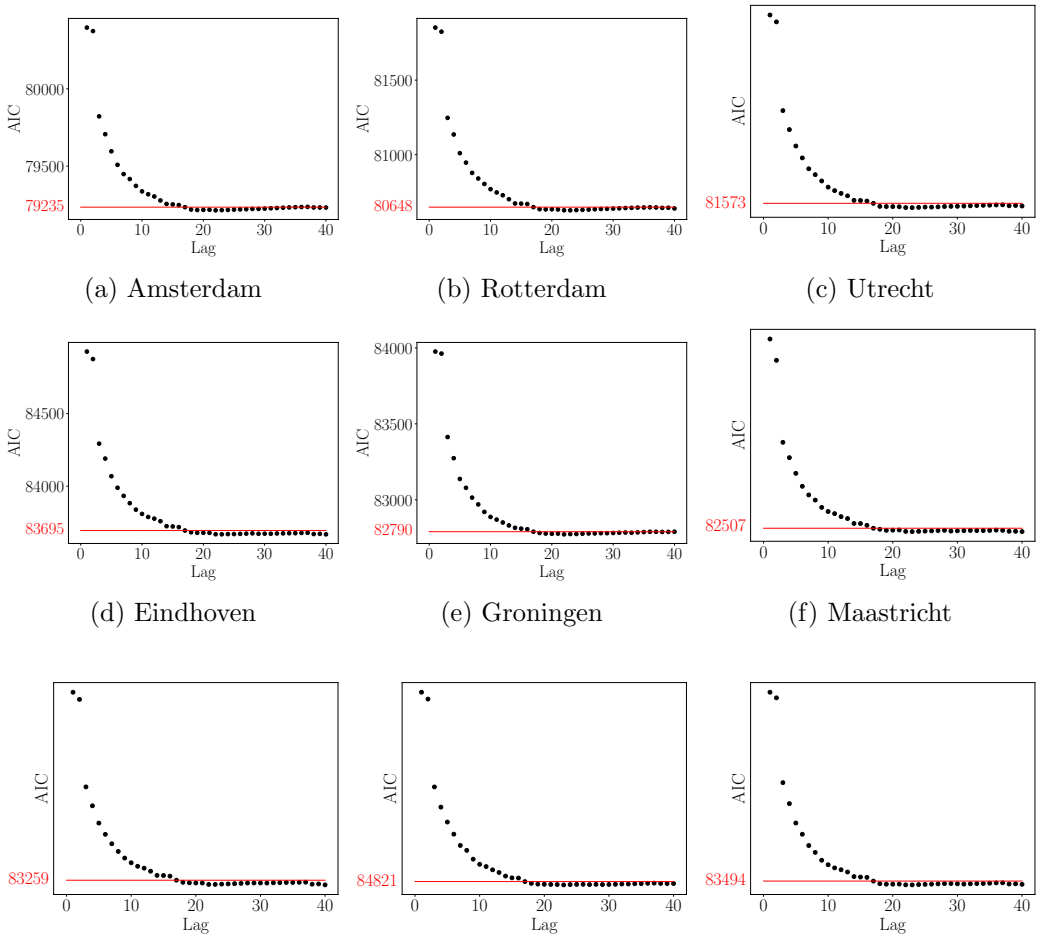

(a) Amsterdam (b) Rotterdam (c) Utrecht

(d) Eindhoven (e) Groningen (f) Maastricht

(g) Tilburg (h) Enschede (i) Nijmegen

4.4 Identifying and Removing Trends

4.4.1 Sources of Trends

Almost all temperature time series exhibit long-run trends. Hence, detrending meteorological data is a crucial step towards successfully model its dynamics. The reasons behind these trends are summarized as follows:

- *Urbanization.* Many weather stations are in, or close to, urban spaces. Urbanization results in higher neighboring and downwind temperatures due to higher absorption of solar radiation and lower cooling evaporation associated with concrete, tarmac, and buildings.

- *Anthropogenic climate change.* Increasing levels of CO_2 affect the atmospheric circulation and cause warming around the globe.

- *Predicted variability.* Observed Trends may be part of a long-timescale cycle occurring due to internal climate processes. That is, some components of

long-term variability are not random. For example, oceanic changes might be oscillatory and affect the atmosphere in a trendy fashion.

Therefore, the dynamics that affect temperature in each city may differ depending on the magnitude of the factors outlined above in each location.

4.4.2 Data detrending

Parametric trends such as quadratic and exponential trends are commonly used for detrending as they often yield satisfactory results. However, fitting non-parametric trends such as moving averages and locally estimated scatterplot smoothing (LOESS) is more suitable for datasets that go beyond 40 years such as the DAT time series shown in Figure 4.11, where it can be seen that temperatures exhibit a clear upward trend.

Figure 4.11: Temperature trend over time

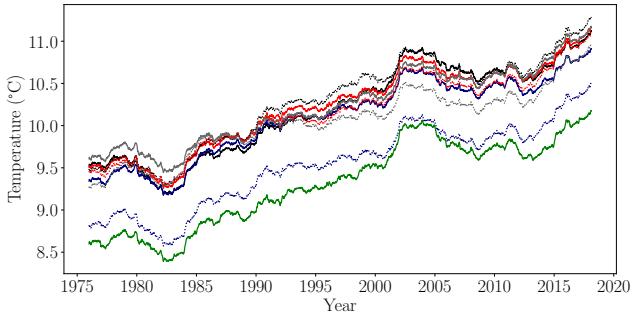

(a) Daily average temperatures' moving average

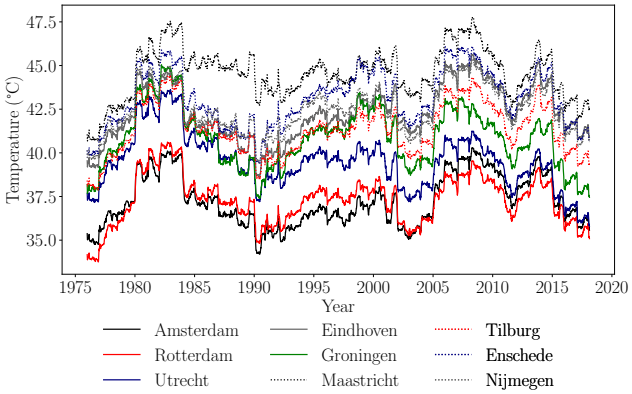

(b) Daily average temperatures' moving variance

Therefore, to decompose the cities' DAT into a trend, seasonal, and residual component the seasonal trend decomposition using loess (STL) as outlined in R. B. Cleveland et al. (1990) is best suited. The underlying algorithm of the STL method is a locally estimated scatterplot smoothing (loess), which is a non-parametric smoothing method proposed by W. S. Cleveland (1979) and developed by W. S. Cleveland and Devlin (1988).

In the case of temperature data, linear loess estimates the trend for a year y_n given other years x_n by weighted regression

$$\varphi(x) = \sum_{n=1}^{p} w_n(x) y_n \tag{4.9}$$

where nearby years have the most weight $w_n(x)$ which depends on the values of the explanatory variable x_i. Given a set $N(x)$ of the q closest x_n

$$w_x(x_n) = W\left(\frac{|x - x_n|}{\lambda_q(x)}\right) \tag{4.10}$$

where $\lambda_q(x)$ is the Euclidean distance of the q^{th} farthest x_i given by

$$\lambda_q(x) = \max_{x_i \in N(x)} |x_i - x|$$

and W is the tricube weight function

$$W(x) = \begin{cases} (1 - x^3)^3 & \text{if } 0 \leq u \leq 1 \\ 0 & \text{otherwise} \end{cases}$$

The STL uses loess in recursive procedures. An inner loop that smooths the seasonal and trend components through iteration, and an outer loop that reduces the effect of outliers. The algorithm decomposes the times series, using either an additive or multiplicative model, into a trend component T_n, seasonal component S_n, and a residual R_n

$$y_n = \hat{T}_n + \hat{S}_n + \hat{R}_n \tag{4.11}$$

$$y_n = \hat{T}_n \times \hat{S}_n \times \hat{R}_n \tag{4.12}$$

Since the amplitude of the seasonality of DATs does not change over time but rather evolves in an increasing fashion with roughly similar fluctuations, a seasonal trend decomposition using Equation 4.11 can successfully break down the time series as it is shown in Figure 4.12. The robustness factor in the STL decomposition is the number of robustness iterations of the outer loop. That is, non-robust fitting consists only of the inner loop and weighs all observations equally, while the robust estimation reweighs the data using loess where it assigns relatively small weight to observations with large outliers. However, it is outlined in R. B. Cleveland et al. (1990) that in the absence of large transient variation, robustness iteration of the outer loop can be omitted. Hence, from a parsimonious perspective, non-robust detrending of DAT is sufficient since robust and non-robust fitting exhibit a minor difference.

A more sophisticated approach is to use a multiple seasonal trend decomposition

using loess (MSTL) which is a recent extension of STL suggested by Bandara et al. (2021) where it assumed that the time series can be decomposed as

$$y_n = \hat{T}_n + \sum_{i=1}^{P} \hat{S}_n^i + \hat{R}_n \tag{4.13}$$

where i is the number of seasonal cycles in y_n. In the case of temperature data, elements such as weekly seasonality could be particularly relevant given the role of urbanization in rising temperature levels. In other words, DAT in urban areas could be significant during weekdays and less pronounced on weekends. Additionally, recurring phenomenons such as El Niño-Southern Oscillation (ENSO) which occurs at irregular intervals of two to seven years is helpful. Although the effect of El Niño on European temperatures is not entirely understood, the fact remains that multiple seasonal effects factor in when dealing DAT time series. Equation 4.14 is used to DAT data such that the seasonal cycle $\hat{S}_n^w$ represents weekly seasonality for effects such as urbanization, $\hat{S}_n^m$ monthly seasonality to account for potential factors such as monthly change in solar insolation, and, finally, $\hat{S}_n^y$ is for yearly seasonality. The results of this decomposition are shown in Figure B.4.

$$y_n = \hat{T}_n + \hat{S}_n^w + \hat{S}_n^m + \hat{S}_n^y + \hat{R}_n \tag{4.14}$$

Besides its ability to implement loess with multiple seasonal patterns, MSTL, as show in Table 4.2, also achieves a very low Root Mean Square Error (RMSE) when compared to other methods such as the Prophet approach (Taylor & Letham, 2018). Nevertheless, STL is the most accurate model as it has the lowest RMSE.

Table 4.2: Root mean square error

	RMSE		
	STL	**MSTL**	**Prophet**
Amsterdam	8.81e-16	1.74e-15	3.13
Rotterdam	8.72e-16	1.76e-15	3.22
Utrecht	8.80e-16	1.72e-15	3.33
Eindhoven	9.14e-16	1.79e-15	3.49
Groningen	8.36e-16	1.66e-15	3.39
Maastricht	9.22e-16	1.81e-15	3.54
Tilburg	8.94e-16	1.77e-15	3.45
Enschede	8.53e-16	1.68e-15	3.59
Nijmegen	8.88e-16	1.77e-15	3.46

Figure 4.12: Seasonal Trend decomposition using Loess of average temperatures

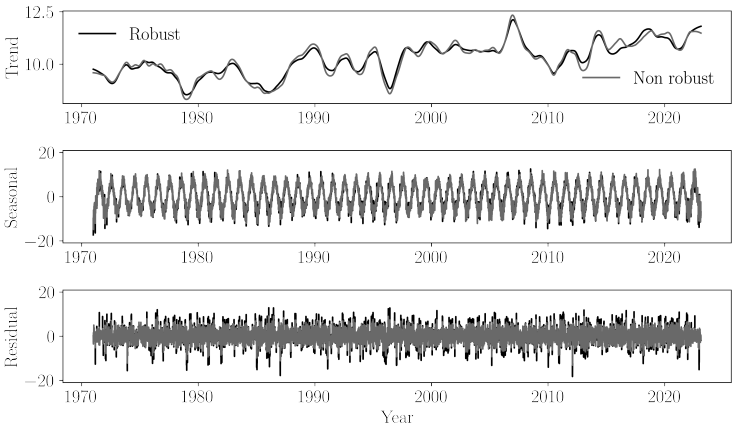

(a) Amsterdam

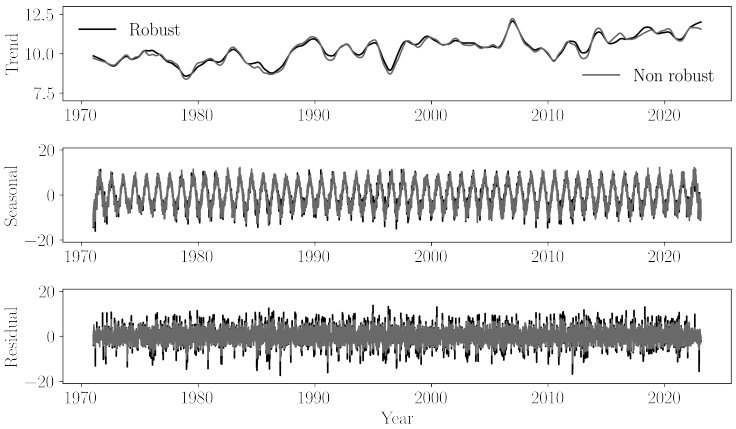

(b) Rotterdam

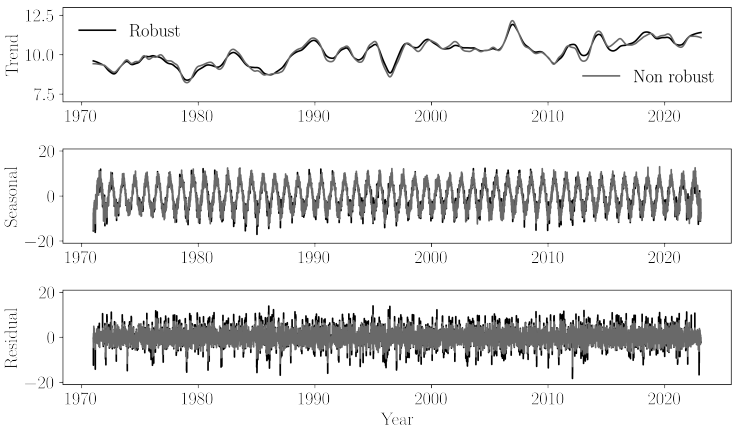

(c) Utrecht

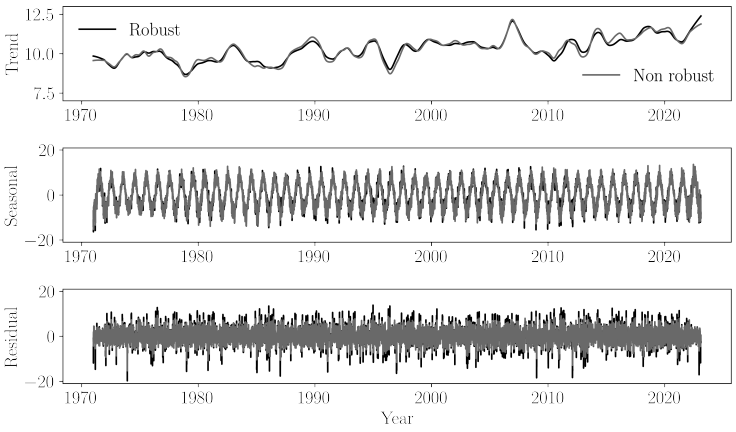

(d) Eindhoven

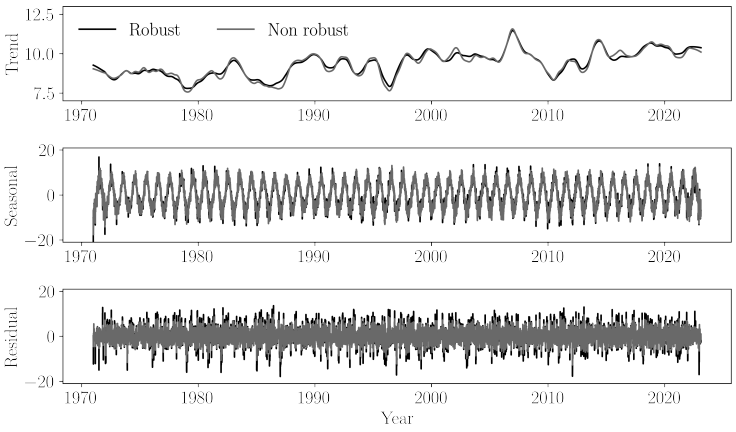

(e) Groningen

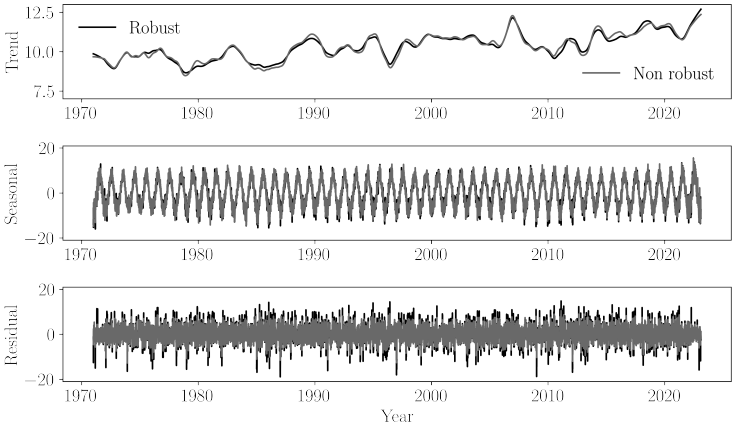

(f) Maastricht

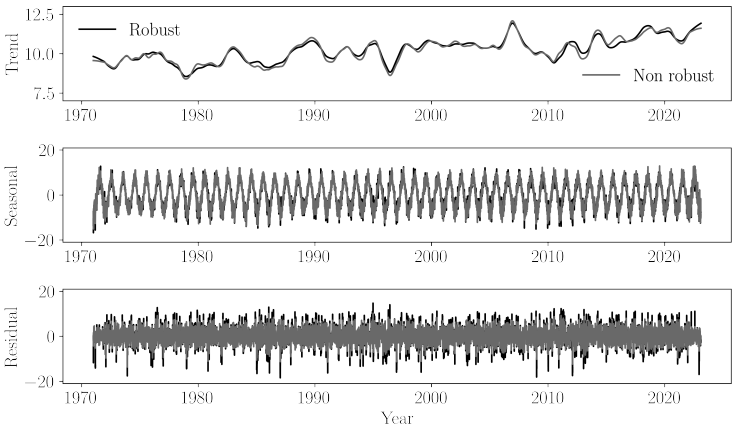

(g) Tilburg

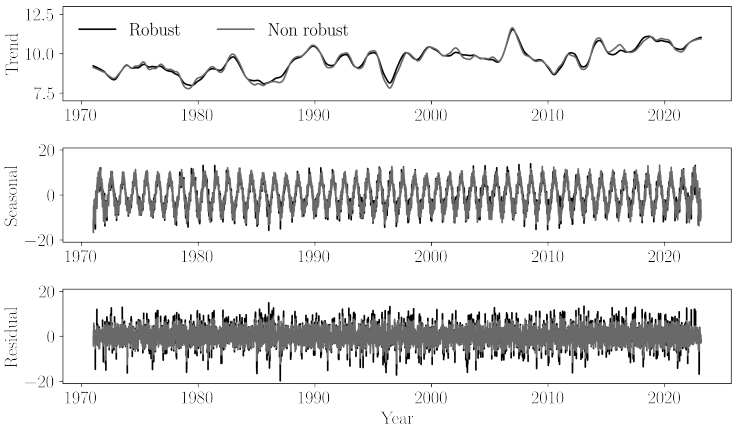

(h) Enschede

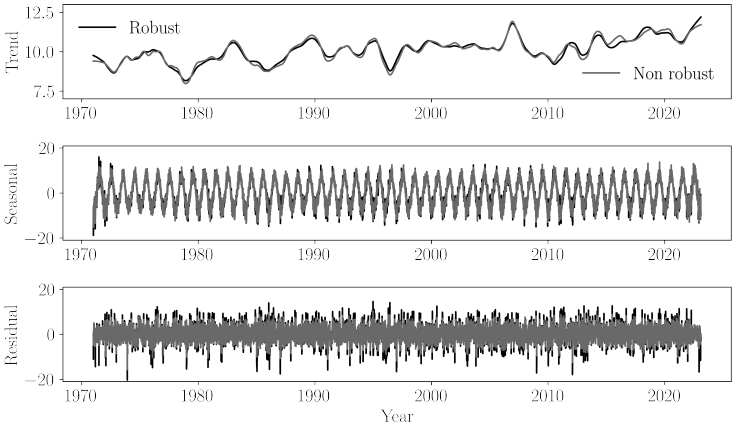

(i) Nijmegen

4.5 A Note on Residuals

It was established through Table 4.2 that STL achieved the lowest RMSE. Nonetheless, to check whether model captured all the information in temperature data, the residuals should satisfy two conditions:

- No autocorrelation
- Zero mean

Moreover, it is useful for the residuals be normally distributed and to have a constant variance across time, although it is not necessary. In fact, looking at STL residuals from the last 10 years[3] in Figure 4.13 it can be seen that residuals form a Gaussian white noise with zero mean and an approximately constant variance.

Figure 4.13: STL residuals

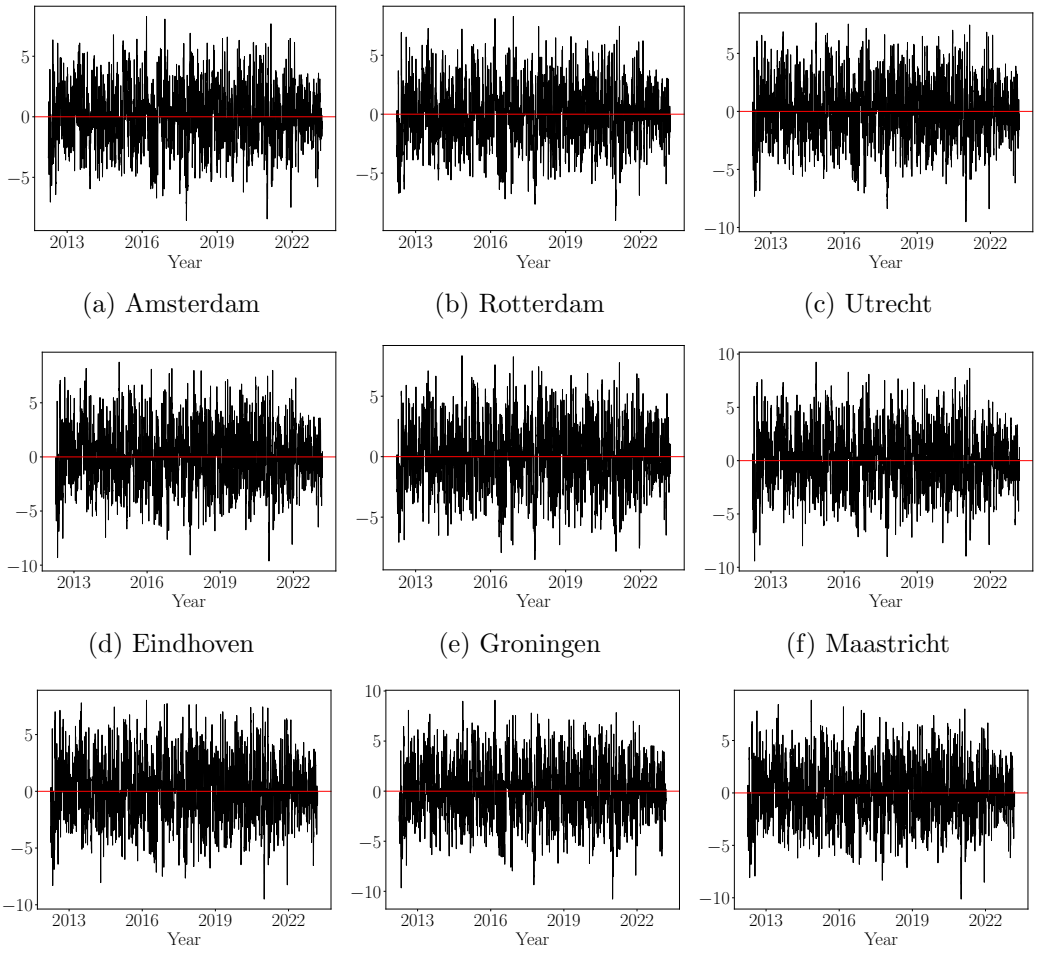

(a) Amsterdam (b) Rotterdam (c) Utrecht

(d) Eindhoven (e) Groningen (f) Maastricht

(g) Tilburg (h) Enschede (i) Nijmegen

Nevertheless, as regards the first condition, both the Ljung-Box and Box-Pierce tests confirm that there is serial correlation in the residuals. These tests verify

[3]It gets harder do graphically analyse the residuals beyond 10 years

whether or not the residuals are a white noise, and their statistics are defined as

$$Q_{LB} = T(T+2)\sum_{k=1}^{\ell}(T-K)^{-1}r_k^2$$

$$Q_{BP} = T\sum_{k=1}^{\ell}r_k^2$$

where T is the number of observations, r_k is the autocorrelation of the series at lag k, and ℓ is the number of lags. Hyndman and Athanasopoulos (2018) suggest choosing ℓ as follows:

$$\ell = \begin{cases} 10 & \text{for non-seasonal data} \\ 2m & \text{for seasonal data} \\ T/5 & \text{if } 10, 2m > T/5 \end{cases} \tag{4.15}$$

where m is the period of seasonality. Hence ℓ is chosen as

$$\ell := \min(10, T/5) = 10$$

The tests use the following hypotheses:

- H_0: The residuals are independently distributed; there is no autocorrelation.
- H_1: The residuals are not independently distributed; there is autocorrelation.

The large values of Ljung-box (LB) and Box-Pierce (BP) statistics in Table 4.3 indicate that the residuals do not come from a white noise process. Additionally, the p-values of both tests are low enough to reject the null hypothesis in favor of autocorrelation.

Table 4.3: Autocorrelation results

	LB. stat	***p*-value**	**BP. stat**	***p*-value**
Amsterdam	26714.174	0.000	26707.921	0.000
Rotterdam	25596.952	0.000	25591.061	0.000
Utrecht	25622.601	0.000	25616.732	0.000
Eindhoven	24813.059	0.000	24807.446	0.000
Groningen	25240.190	0.000	25234.351	0.000
Maastricht	26863.134	0.000	26856.993	0.000
Tilburg	24992.160	0.000	24986.469	0.000
Enschede	24582.664	0.000	24577.099	0.000

LB: Ljung-Box
BP: Box-Pierce

The results in Table 4.3 suggest that STL may not did not account for all available information. This is further confirmed by the close-to-normal distribution of the residuals shown in Figure 4.14, the distribution has a slightly longer left tail relative to the normal distribution which means that assuming normality for prediction intervals may be inaccurate.

Figure 4.14: Distribution of STL residuals

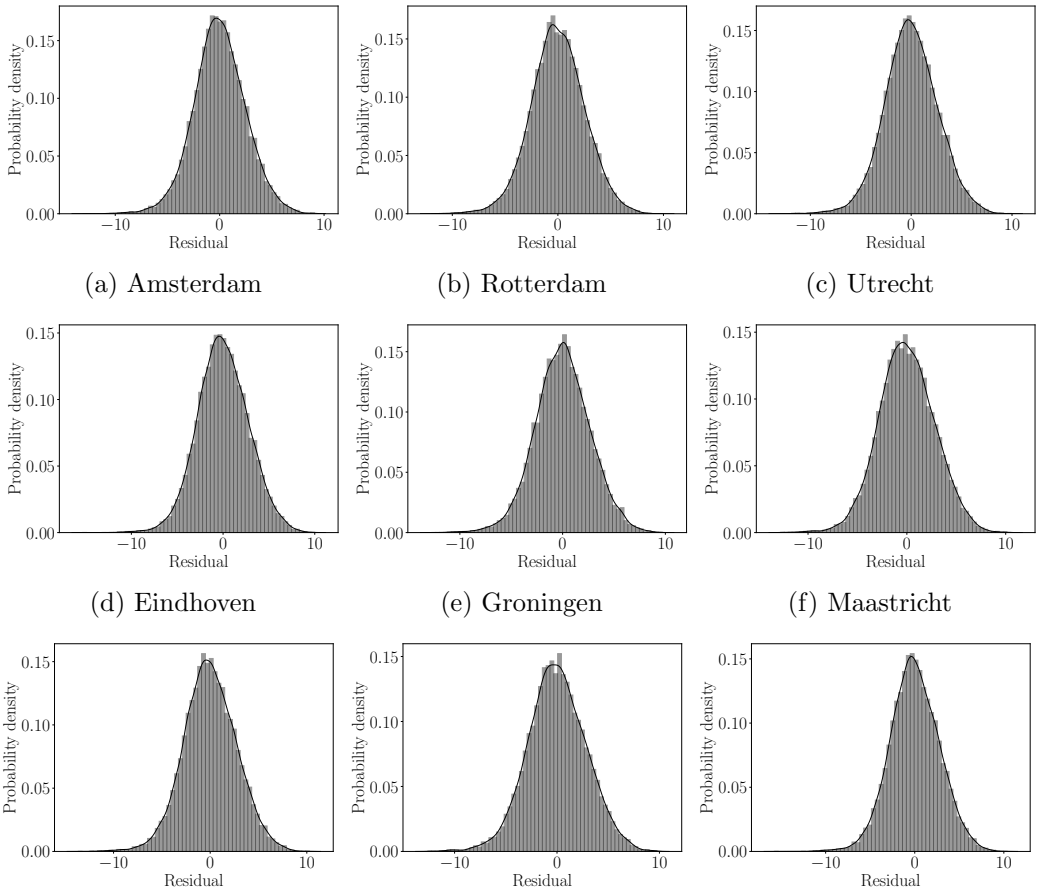

(a) Amsterdam (b) Rotterdam (c) Utrecht

(d) Eindhoven (e) Groningen (f) Maastricht

(g) Tilburg (h) Enschede (i) Nijmegen

Daily Average Temperature Modeling

In chapter 4 different models were implemented to describe DAT dynamics. Based on some assumptions about seasonality in temperature, statistical decomposition algorithms such as STL, MSTL, and Prophet generated results that align with the intuition. However, so much as non-parametric approaches allow fitting trends that are data-specific, one cannot account for all seasonalities at the expense of over-fitting the model.

5.1 Seasonal Variation

Following Cao, Wei, et al. (1999), Cao and Wei (2000), and Cao et al. (2003) it is assumed that, besides being affected by global warming and urbanization, DAT follows a yearly cycle in all cities as shown in Figure 4.1. Moreover, Figures 4.1, 4.2, and 4.6 show that DAT moves around a seasonal mean from which it cannot deviate for extended periods. Finally, DAT is more volatile during winter as depicted by Figure 4.6.

As it can be seen in Figure 4.1, the DAT has a lot of noise. Hence, it can be processed as signal through convolution in order to smooth out noisy data and extract the underlying trend, whereby the DAT time series denoted as $\widetilde{x} = \{x_t\}_{t\geq 0}$ and a rectangular pulse $\delta(t)$ as defined in Manolakis and Ingle (2011) convolve to smooth out noisy data. That is, the convolution is given by

$$(x * \delta)(t) = \int_{-\infty}^{\infty} x(\tau)\delta(t-\tau)d\tau \tag{5.1}$$

Denoting by $h(t)$ the response of the system to the impulse δ, with a width w, height $1/w$, and area $\int \delta_w(t)dt = 1$, Equation5.1 becomes

$$(x * h)(t) = \int_{-\infty}^{\infty} x(\tau)h(t-\tau)d\tau \tag{5.2}$$

where

$$h(t) = \begin{cases} \frac{1}{w}, & \text{if } |t| \leq \frac{w}{2} \\ 0, & \text{otherwise} \end{cases} \tag{5.3}$$

Since $h(t) = 0$ outside the window then

$$(x * h)(t) = \int_{t-\frac{w}{2}}^{t+\frac{w}{2}} x(\tau)h(t-\tau)d\tau \tag{5.4}$$

Substituting for $h(t)$ into Equation 5.4

$$(x * h)(t) = \int_{t-\frac{w}{2}}^{t+\frac{w}{2}} x(\tau)\frac{1}{w}dt = \frac{1}{w}\int_{t-\frac{w}{2}}^{t+\frac{w}{2}} x(\tau)dt \tag{5.5}$$

The filtered signal given by this convolution reveals that DATs have quasi-uniform peaks and troughs as shown in Figure 5.1.

Figure 5.1: Convolution of the DAT series

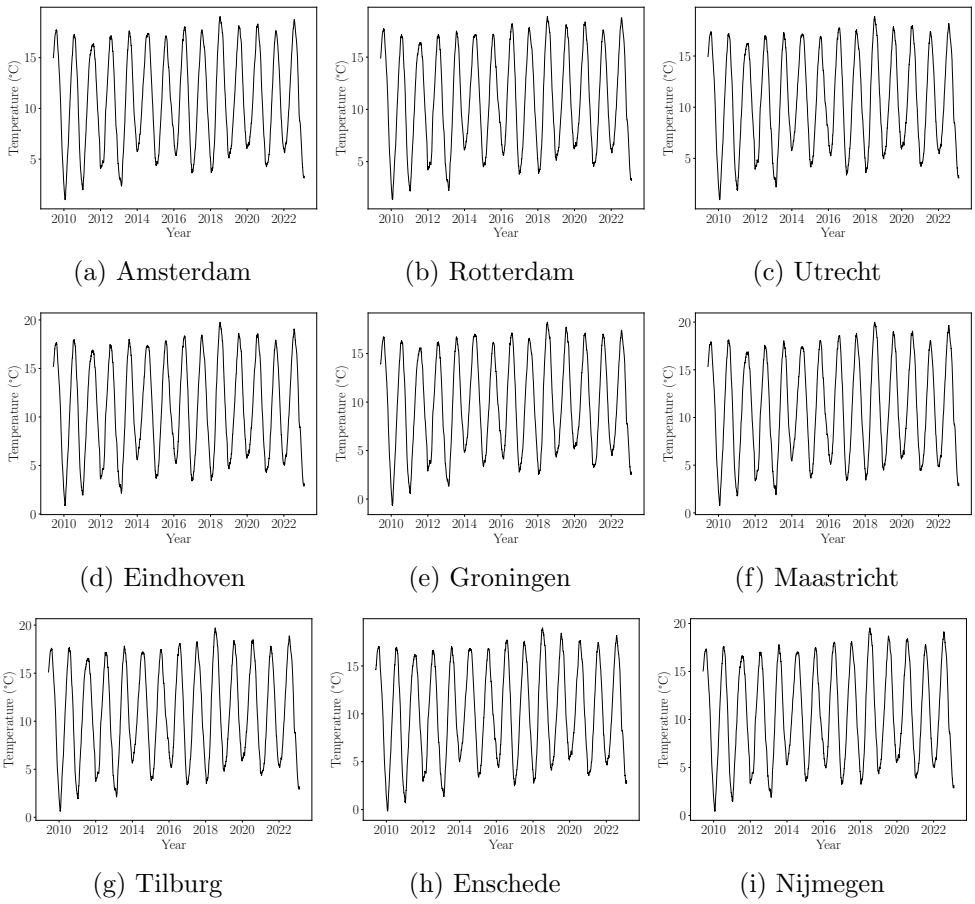

(a) Amsterdam (b) Rotterdam (c) Utrecht

(d) Eindhoven (e) Groningen (f) Maastricht

(g) Tilburg (h) Enschede (i) Nijmegen

The output of the linear time-invariant system in the section above can be represented as shifted orthogonal basis signals (Alaton et al., 2002). More explicitly, following Tindall (2006), the annual cycle or seasonal mean of DAT θ_t, $\forall t \in \mathbb{N}$, can be represented by summing the linear trend shown in Figure 4.11a and a term for the seasonal variation that depicts the periodicity of temperature as follows

$$\theta_t = T_{linear} + T_{seasonal} \tag{5.6}$$

where

$$\begin{aligned} T_{linear} &= A + Bt \\ T_{seasonal} &= \varepsilon C_0 + \sum_i C_i \sin(\omega t + \varphi) + \sum_i D_i \cos(\omega t + \zeta) \end{aligned} \tag{5.7}$$

Hence, from Equations 5.6 and 5.7 the seasonal mean can be described by a truncated Fourier series of the form

$$\theta_t = A + Bt + \sum_i C_i \sin(\omega t + \varphi) + \sum_i D_i \cos(\omega t + \zeta) \tag{5.8}$$

where the effect of the coefficient C_0 is captured in the linear trend (Tindall, 2006). The seasonal variation in Tindall (2006) was estimated using only the first order of Equation 5.8, that is

$$\theta_t = A + Bt + C \sin(\omega t + \varphi) + D \cos(\omega t + \zeta) \tag{5.9}$$

However, the framework used in Mraoua (2007), Esunge and Njong (2020), and Alaton et al. (2002) suggests that the annual cycle can be captured by the first order Fourier series of a sinusoid of the form

$$\theta_t = A + Bt + C \sin(\omega t + \varphi) \tag{5.10}$$

where $\omega \approx 2\pi/365.25$, φ is a phase shift that calibrates the sinusoid for the fact that maximum and minimum temperatures do not occur on the first of January and the first of July, respectively, and A is the amplitude of the sinusoid, i.e. the peak deviation of the function from its center position, which should be chosen alongside B, and C so that the curve fits the data.

5.2 Stochastic Representation

The time evolution of DAT is given by a stochastic differential equation (SDE) given by Equation 5.11 and whose solution is a Gaussian mean-reverting Ornstein-Uhlenbeck process

$$dT_t = \kappa(\theta_t - T_t)dt + \sigma_t dW_t, \quad T_{t_0} = T_0 > 0 \tag{5.11}$$

where the DAT is given by T_t, $\kappa \in \mathbb{R}$ is the speed of mean reversion, θ_t is the seasonal mean, σ is the daily volatility of temperature, and W_t is a Wiener process. However, to account for the mean reverting aspect of temperature in the long-run, it is important to ensure that temperature reverts back to its long-run mean θ_t

Proposition 1.

$$dT_t = \kappa(\theta_t - T_t)dt + \sigma_t dW_t \implies \lim_{t\to\infty} \mathbb{E}[T_t] = \theta_t \tag{5.12}$$

Proof 1. *Let W_t be a BM, and let $\mathcal{F}_t$ be an associated filtration, where Θ and Δ are adapted stochastic processes with respect to $\mathcal{F}_t$, then $\widetilde{T}_t$ (temperature[1].) is an Ito process of the form*

$$\widetilde{T}_T = \widetilde{T}_0 + \int_t^T \Theta_t dt + \int_t^T \Delta_t dW_t \tag{5.13}$$

And let $f = f(t, \widetilde{t})$ be a $C^2_{t,\widetilde{t}}$ function then

$$\begin{aligned} f(T, \widetilde{T}_T) = f(0, \widetilde{T}_0) &+ \int_0^T \frac{\partial f}{\partial t}(t, \widetilde{T}_t)dt + \int_0^T \frac{\partial f}{\partial \widetilde{t}}(t, \widetilde{T}_t)\Delta_t dW_t \\ &+ \int_0^T \frac{\partial f}{\partial \widetilde{t}}(t, \widetilde{T}_t)\Theta_t dt + \frac{1}{2}\int_0^T \frac{\partial^2 f}{\partial \widetilde{t}^2}(t, \widetilde{T}_t)\Delta_t^2 dt, \quad \forall T \geq 0 \end{aligned} \tag{5.14}$$

which, in differential notation, becomes

$$\begin{aligned} df(t, \widetilde{T}_t) &= \frac{\partial f}{\partial t}(t, \widetilde{T}_t)dt + \frac{\partial f}{\partial \widetilde{t}}\Delta_t dW_t + \frac{\partial f}{\partial \widetilde{t}}\Theta_t dt + \frac{1}{2}\frac{\partial^2 f}{\partial \widetilde{T}^2}\Delta_t^2 dt \\ &\iff \frac{\partial f}{\partial t}(t, \widetilde{T}_t)dt + \frac{\partial f}{\partial \widetilde{t}}d\widetilde{T}_t + \frac{1}{2}\frac{\partial^2 f}{\partial \widetilde{T}^2}(d\widetilde{T})^2 \end{aligned} \tag{5.15}$$

To derive the expected value of the temperature process[2] , let $f(t, \widetilde{t}) = e^{\kappa t}\widetilde{t}$ and substitute using Equation 5.11, hence

$$\begin{aligned} d(e^{\kappa t}\widetilde{T}_t) &= df(t, \widetilde{T}_t) \\ &= \kappa e^{\kappa t}\widetilde{T}_t dt + e^{\kappa t}\left[\kappa(\theta_t - \widetilde{T}_t)dt + \sigma_t dW_t\right] \\ &= \kappa e^{\kappa t}\theta_t dt + e^{\kappa t}\sigma_t dW_t \end{aligned} \tag{5.16}$$

and integrating over $u \in [s, t], t > s$, we get

$$\begin{aligned} \int_s^t d(e^{\kappa t}\widetilde{T}_t) &= e^{\kappa t}\widetilde{T}_t - e^{\kappa s}\widetilde{T}_s \\ &= \int_s^t \kappa e^{\kappa t}\theta_u du + e^{\kappa t}\sigma_u dW_u \end{aligned} \tag{5.17}$$

changing the the base of the Riemann integral to $d\theta_t$ and dividing through by $e^{\kappa t}$

[1] Here T_t in Equation 5.11 has been changed to $\widetilde{T}_t$ for notation convenience

[2] $e^{\kappa t}$ is similar to the integrating factor used to derive the moments of the Cox-Ingersoll-Ross model (Jafari & Abbasian, 2017)

gives

$$\begin{aligned}\widetilde{T}_t &= \widetilde{T}_s e^{-\kappa(t-s)} + \int_s^t e^{-\kappa(t-u)} d\theta_u + \int_s^t e^{-\kappa(t-s)} \sigma_u dW_u \\ &= \widetilde{T}_s e^{-\kappa(t-s)} + \theta_t e^{-\kappa(t-t)} - \theta_s e^{-\kappa(t-s)} + \int_s^t e^{-\kappa(t-s)} \sigma_u dW_u \\ &= \theta_t + \left(\widetilde{T}_s - \theta_s\right) e^{-\kappa(t-s)} + \int_s^t e^{-\kappa(t-s)} \sigma_u dW_u\end{aligned} \tag{5.18}$$

Since the expectation of the Ito integral is zero, we obtain

$$\mathbb{E}[\widetilde{T}_t] = \theta_t + \left(\widetilde{T}_s - \theta_s\right) e^{-\kappa(t-s)} \tag{5.19}$$

which goes to show that, under the settings of a normal OU process, temperature does not revert back to its long-term mean.

Proposition 2.

$$dT_t = \left(\frac{d\theta_t}{dt} + \kappa(\theta_t - T_t)\right) dt + \sigma_t dW_t \implies \lim_{t\to\infty} \mathbb{E}[T_t] = \theta_t \tag{5.20}$$

Proof 2. *Multiplying through the modified OU process by the integrating factor* $e^{\int_0^t k du}$*, we obtain*

$$e^{\int_0^t k du} d\widetilde{T}_u + e^{\int_0^t k du} d\theta_u - e^{\int_0^t k du} \kappa(\theta_u - \widetilde{T}_u) du = e^{\int_0^t k du} \sigma_u dW_u \tag{5.21}$$

Let us consider the Ito process $Z_t = e^{\int_0^t k du}(\theta_u - \widetilde{T}_u)$*, then, by the product rule, we get*

$$dZ_t = d\left[e^{\int_0^t k du}(\theta_u - \widetilde{T}_u)\right] = e^{\int_0^t k du} \sigma_u dW_u \tag{5.22}$$

Therefore, we find

$$Z_t = Z_0 - \int_0^t e^{\int_0^t \kappa du} \sigma_u dW_u \tag{5.23}$$

Substituting $Z_t = e^{\int_0^t k du}(\theta_u - \widetilde{T}_u)$ *for* $\theta_0 = \widetilde{T}_0$*, gives*

$$e^{\int_0^t k du}(\theta_t - \widetilde{T}_t) = e^{\int_0^t k du}(\theta_0 - \widetilde{T}_0) - \int_0^t e^{\int_0^t k du} \sigma_u dW_u \tag{5.24}$$

Therefore

$$\widetilde{T}_t = \theta_t + e^{-\int_0^t k du} \int_0^t e^{\int_0^t k du} \sigma_u dW_u \tag{5.25}$$

Since the expectation of the Ito integral is zero, we arrive at

$$\mathbb{E}[\widetilde{T}_t] = \theta_t \tag{5.26}$$

Hence, as Dornier and Querel (2000) showed, modifying the OU process 5.11,

where mean reversion is modeled by the time-varying term

$$\frac{d\theta_t}{dt} = B + \omega C \cos(\omega t + \varphi) \tag{5.27}$$

adjusts the drift in the SDE 5.11 so that 5.20 holds, that is

$$dT_t = \left(\frac{d\theta_t}{dt} + \kappa(\theta_t - T_t) \right) dt + \sigma_t dW_t \tag{5.28}$$

5.3 Parameter Estimation

5.3.1 Model Fitting

The parameters in Equation 5.10 can be estimated by applying the method of least squares to the series of observations. That is, if we let $A = a_1$, $B = b_2$, $C\cos(\varphi) = a_3$, and $C\sin(\varphi) = a_4$ then Equation 5.10 can be rewritten as

$$\theta_t = a_1 + a_2 t + a_3 \sin(\omega t) + a_4 \cos(\omega t) \tag{5.29}$$

Moreover it can be seen in Figure 5.2, that the models in Equation 5.29 and 5.9 fit the data in an overlapping way. Therefore, in order to avoid overfitting the model, θ_t is estimated by finding the vector $\xi = (a_1, a_2, a_3, a_4)$ such that

$$\begin{aligned} A &= a_1 \\ B &= a_2 \\ C &= \sqrt{a_3^2 + a_4^2} \\ \varphi &= \arctan\left(\frac{a_4}{a_3}\right) - \pi \end{aligned}$$

where the estimation errors are minimized using the Levenberg–Marquardt algorithm developed by Levenberg (1944) and Marquardt (1963), that is

$$\underset{\xi}{\operatorname{argmin}} \|\Theta - F(X, \xi)\|_2^2, \quad \forall t \tag{5.30}$$

where $\|.\|_2$ is the ℓ^2 norm of Θ, the vector representing the elements in Equation 5.29, and X, the data vector.

Figure 5.2: Sinusoidal representation of DAT's annual cycle

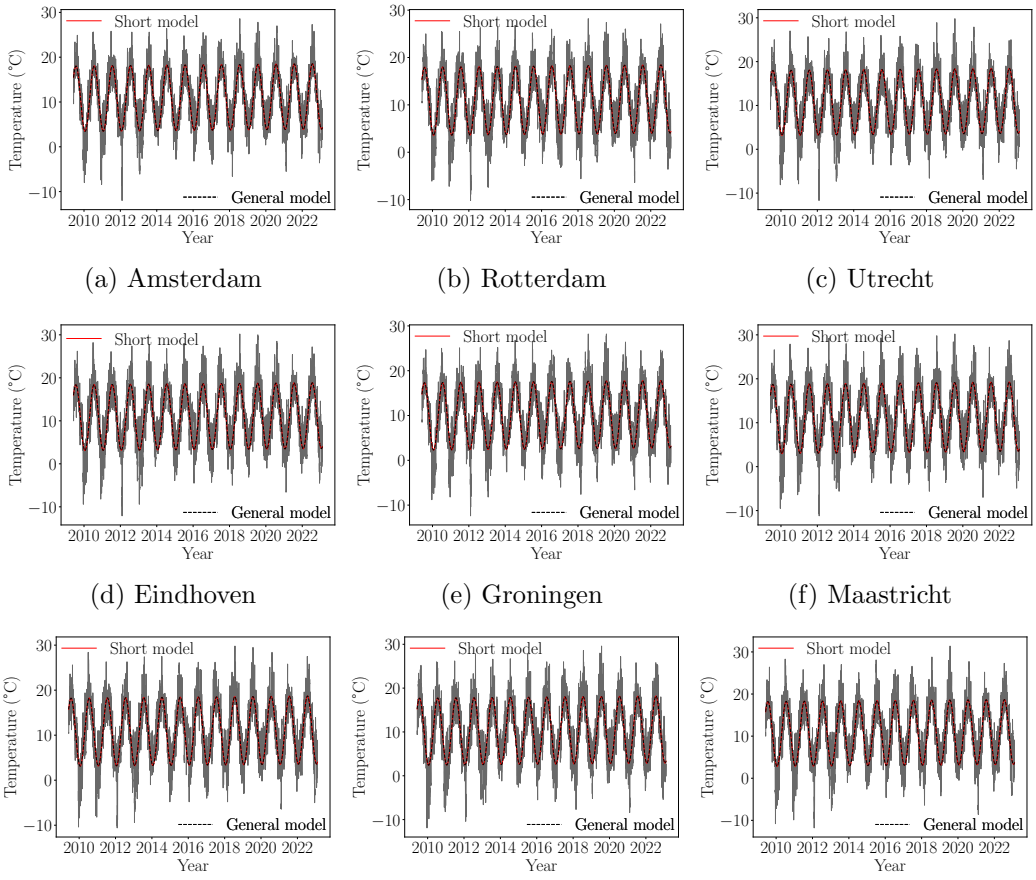

(a) Amsterdam (b) Rotterdam (c) Utrecht

(d) Eindhoven (e) Groningen (f) Maastricht

(g) Tilburg (h) Enschede (i) Nijmegen

Consequently, the fitted coefficients for every city are as follows:

$$\theta_t = \begin{cases} 9.22 + 1.07\text{e-}4t + 7.31\sin(\omega t - 1.95) & \text{in Amsterdam} \\ 9.27 + 1.05\text{e-}4t + 7.21\sin(\omega t - 1.96) & \text{in Rotterdam} \\ 9.17 + 9.89\text{e-}5t + 7.42\sin(\omega t - 1.93) & \text{in Utrecht} \\ 9.36 + 9.73\text{e-}5t + 7.67\sin(\omega t - 1.91) & \text{in Eindhoven} \\ 8.40 + 9.98\text{e-}5t + 7.54\sin(\omega t - 1.94) & \text{in Groningen} \\ 9.25 + 1.13\text{e-}4t + 7.90\sin(\omega t - 1.90) & \text{in Maastricht} \\ 9.24 + 1.01\text{e-}4t + 7.57\sin(\omega t - 1.92) & \text{in Tilburg} \\ 8.59 + 1.04\text{e-}4t + 7.66\sin(\omega t - 1.91) & \text{in Enschede} \\ 9.17 + 9.54\text{e-}5t + 7.72\sin(\omega t - 1.91) & \text{in Nijmegen} \end{cases} \tag{5.31}$$

The amplitude of these sine functions indicate that in the nine Dutch cities during summer periods, temperature is typically 7.56°C higher than in winter periods. Furthermore, the small values of B imply a weak linear trend. Nevertheless, it also shows that, depending on the city, average temperature increases by 1°C every 24

to 29 years[3].

5.3.2 Estimating the Speed of Mean Reversion

To find the parameter κ in the modified OU process proposed by Dornier and Querel (2000), Equation 5.28 can be approximated using the Euler-Maruyama scheme, that is

$$T_i \approx T_{i-1} + \int_{t_{i-1}}^{t_i} \frac{d\theta_t}{dt} dt + \int_{t_{i-1}}^{t_i} \kappa(\theta_{t_{i-1}} - T_{t_{i-1}}) dt + \int_{t_{i-1}}^{t_i} \sigma_{t_i} dB_t \tag{5.32}$$

where T_i denotes the approximation of T_{t_i} and $T_0 = T_{t0}$. Since the integrands are approximated by their values at the left-side boundary of the integration interval, we get

$$T_i \approx T_{i-1} + \theta_{t_i} - \theta_{t_{i-1}} + \kappa(\theta_{t_{i-1}} - T_{t_{i-1}})\Delta t + \sigma_{t_i} \Delta B_t \tag{5.33}$$

In Equation 5.33, The value ΔB_t is normally distributed with zero mean and variance Δt, i.e. $\Delta B_t \sim \mathcal{N}(0,\, \Delta t)$. Let $Z_t\sqrt{\Delta t}$ be a sample from this distribution where $Z_t \sim \mathcal{N}(0,\, 1)$, then Equation 5.33 becomes

$$T_i \approx T_{i-1} + \theta_{t_i} - \theta_{t_{i-1}} + \kappa(\theta_{t_{i-1}} - T_{t_{i-1}})\Delta t + \sigma_{t_i} Z_t \sqrt{\Delta t} \tag{5.34}$$

by setting $\Delta t = 1$, we have that

$$T_i \approx T_{i-1} + \theta_{t_i} - \theta_{t_{i-1}} + \kappa(\theta_{t_{i-1}} - T_{t_{i-1}}) + \sigma_{t_i} Z_t \tag{5.35}$$

by letting

$$\Lambda_j = T_j - \theta_j$$

Equation 5.35 simplifies to

$$\Lambda_{t_i} = \Lambda_{t_{i-1}} - \kappa \Lambda_{t_{i-1}} + \sigma_{t_i} Z_t \tag{5.36}$$

Equation 5.36 can be modeled as an AR(1) process of the detrended and deseasonalized temperatures with a zero constant such that

$$\Lambda_{t_i} = \gamma \Lambda_{t_{i-1}} + e_{t_i} \tag{5.37}$$

where $\kappa = 1 - \gamma$

The model in Equation 5.37 can be intuitively apprehended by inspecting the residuals of the model fitting in the last subsection. A closer look at Figure 5.3 shows that there is significant autocorrelation in the first lag term, therefore it is no coincidence that the model is estimated as an AR(1) process.

[3] t is a daily index, hence the increase is approximately estimated as $B \times 365.25 \times years$

Figure 5.3: Partial autocorrelation of the AR model lags

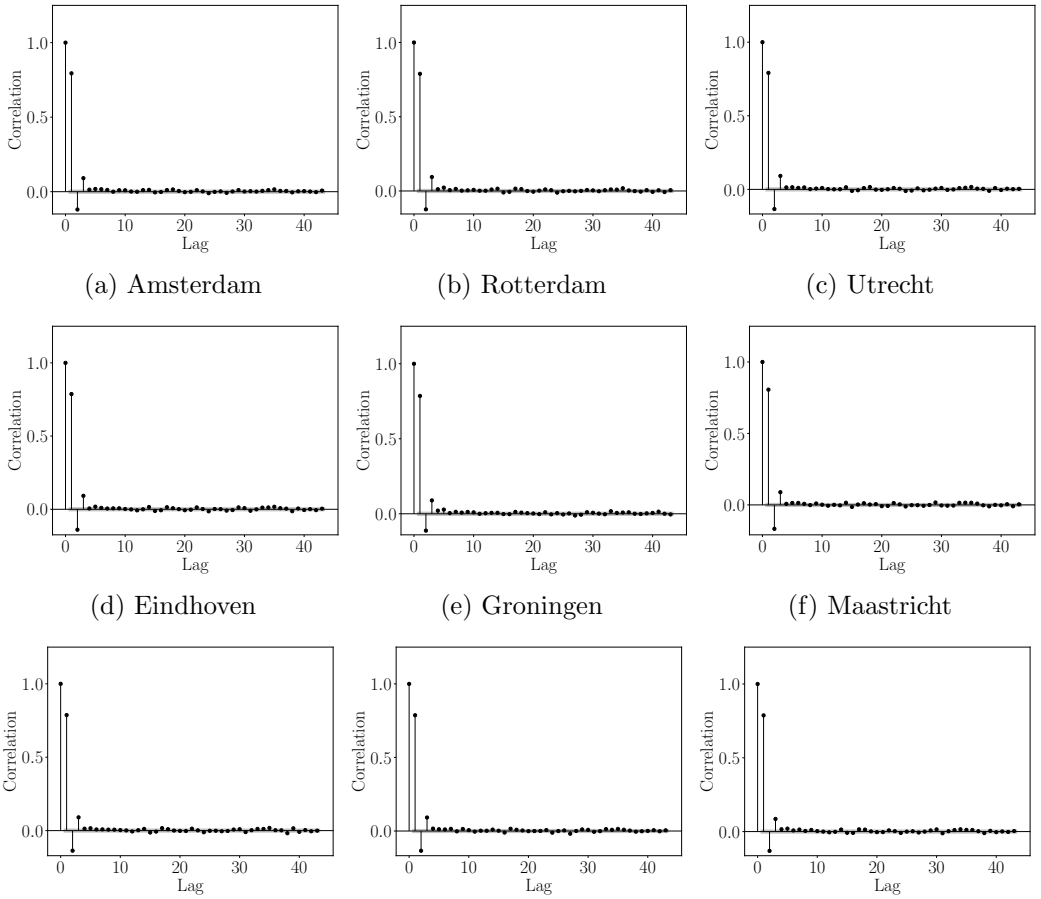

(a) Amsterdam (b) Rotterdam (c) Utrecht

(d) Eindhoven (e) Groningen (f) Maastricht

(g) Tilburg (h) Enschede (i) Nijmegen

Using the results from the previous subsection, the results of the AR regression 5.37 are provided in Table 5.1.

Table 5.1: Autoregressive model results

City:		Amsterdam	**No. Observations:**			19034
Dep. Variable:		Residual	**Log Likelihood**			-39509.286
Model:		AR(1)	**AIC**			79022.573
Method:		Conditional MLE	**BIC**			79038.280
	Coef	**Std err**	**z**	$p > \|z\|$	**[0.025**	**0.975]**
residual.L1	0.7942	0.004	180.427	0.000	0.786	0.803
City:		Rotterdam	**No. Observations:**			19034
Dep. Variable:		Residual	**Log Likelihood**			-40238.849
Model:		AR(1)	**AIC**			80481.699
Method:		Conditional MLE	**BIC**			80497.407
	Coef	**Std err**	**z**	$p > \|z\|$	**[0.025**	**0.975]**
residual.L1	0.7891	0.004	177.303	0.000	0.780	0.798

City:		Utrecht	**No. Observations:**			19034
Dep. Variable:		Residual	**Log Likelihood**			-40707.873
Model:		AR(1)	**AIC**			81419.746
Method:		Conditional MLE	**BIC**			81435.454
	Coef	**Std err**	**z**	$p > \|z\|$	**[0.025**	**0.975]**
residual.L1	0.7913	0.004	178.623	0.000	0.783	0.800

City:		Eindhoven	**No. Observations:**			19034
Dep. Variable:		Residual	**Log Likelihood**			-41773.694
Model:		AR(1)	**AIC**			83551.389
Method:		Conditional MLE	**BIC**			83567.097
	Coef	**Std err**	**z**	$p > \|z\|$	**[0.025**	**0.975]**
residual.L1	0.7869	0.004	175.973	0.000	0.778	0.796

City:		Groningen	**No. Observations:**			19034
Dep. Variable:		Residual	**Log Likelihood**			-41297.792
Model:		AR(1)	**AIC**			82599.583
Method:		Conditional MLE	**BIC**			82615.291
	Coef	**Std err**	**z**	$p > \|z\|$	**[0.025**	**0.975]**
residual.L1	0.7856	0.004	175.345	0.000	0.777	0.794

City:		Maastricht	**No. Observations:**			19034
Dep. Variable:		Residual	**Log Likelihood**			-41262.055
Model:		AR(1)	**AIC**			82528.111
Method:		Conditional MLE	**BIC**			82543.819
	Coef	**Std err**	**z**	$p > \|z\|$	**[0.025**	**0.975]**
residual.L1	0.8065	0.004	188.224	0.000	0.798	0.815

City:		Tilburg	**No. Observations:**			19034
Dep. Variable:		Residual	**Log Likelihood**			-41554.043
Model:		AR(1)	**AIC**			83112.086
Method:		Conditional MLE	**BIC**			83127.794
	Coef	**Std err**	**z**	$p > \|z\|$	**[0.025**	**0.975]**
residual.L1	0.7870	0.004	176.049	0.000	0.778	0.796

City:		Enschede	**No. Observations:**			19034
Dep. Variable:		Residual	**Log Likelihood**			-42347.546
Model:		AR(1)	**AIC**			84699.092
Method:		Conditional MLE	**BIC**			84714.800
	Coef	**Std err**	**z**	$p > \|z\|$	**[0.025**	**0.975]**
residual.L1	0.7868	0.004	175.924	0.000	0.778	0.796

City:		Nijmegen	**No. Observations:**			19034
Dep. Variable:		Residual	**Log Likelihood**			-41650.776
Model:		AR(1)	**AIC**			83305.553
Method:		Conditional MLE	**BIC**			83321.261
	Coef	**Std err**	**z**	$p > \|z\|$	**[0.025**	**0.975]**
residual.L1	0.7874	0.004	176.271	0.000	0.779	0.796

BIC: Bayesian information criterion

In the regression Table 5.1, the coefficient, i.e γ, values are used to estimate the speed of mean reversion for every city as

$$\kappa = \begin{cases} 0.206 & \text{in Amsterdam} \\ 0.211 & \text{in Rotterdam} \\ 0.209 & \text{in Utrecht} \\ 0.213 & \text{in Eindhoven} \\ 0.214 & \text{in Groningen} \\ 0.194 & \text{in Maastricht} \\ 0.213 & \text{in Tilburg} \\ 0.213 & \text{in Enschede} \\ 0.213 & \text{in Nijmegen} \end{cases} \tag{5.38}$$

5.3.3 Estimating the Volatility

The last step in modelling the underlying consists in completing the modified OU process 5.28 through estimating σ. A quick look at Figures 5.4 and 5.5 show that the variance is not constant. In other words temperature is more volatile towards the beginning/end of the years.

Figure 5.4: DAT dynamics model fit

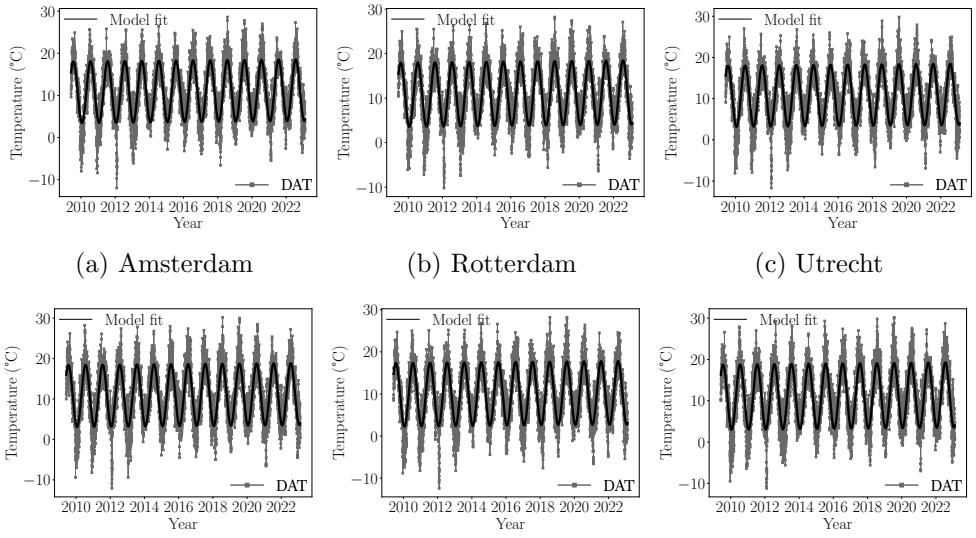

(a) Amsterdam (b) Rotterdam (c) Utrecht

(d) Eindhoven (e) Groningen (f) Maastricht

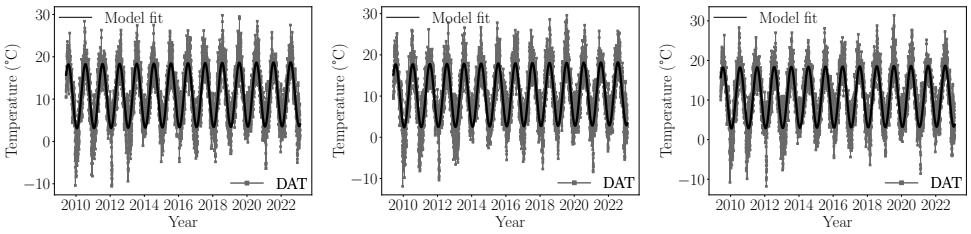

(g) Tilburg (h) Enschede (i) Nijmegen

Figure 5.5: Standard deviation of the DAT process

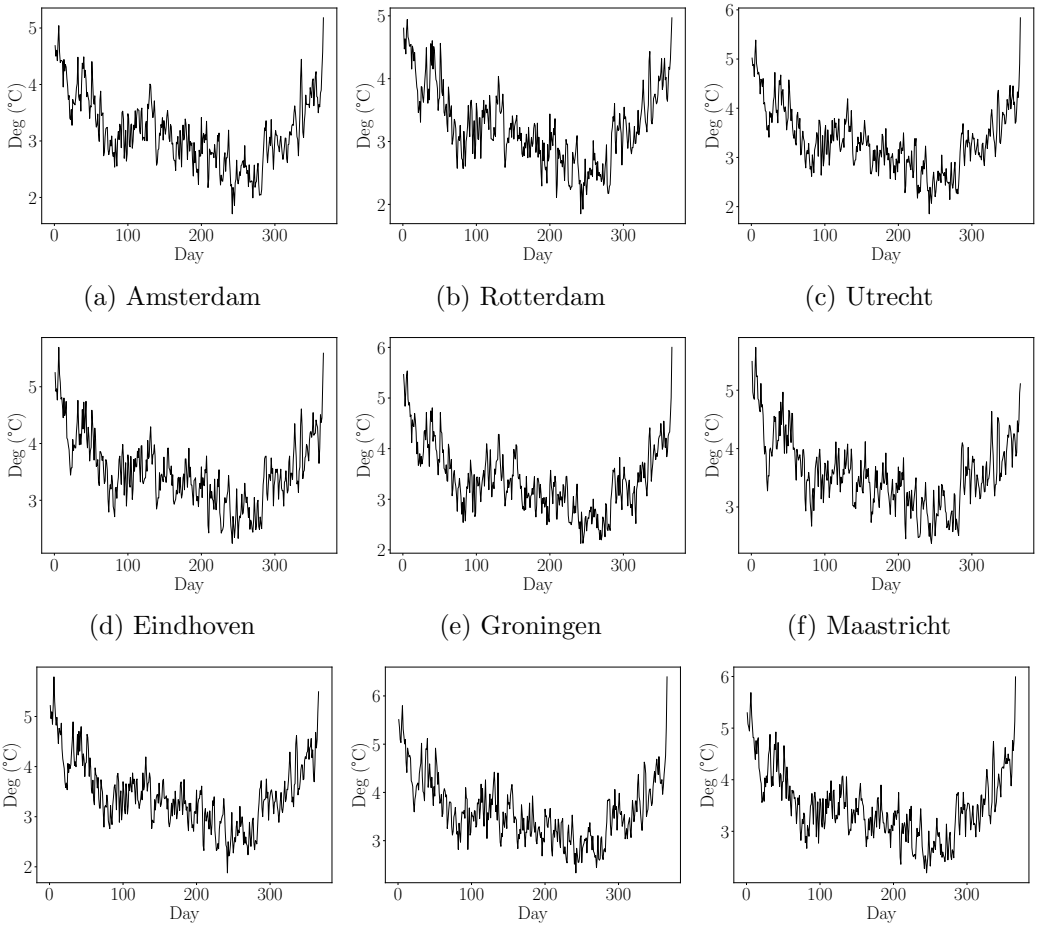

(a) Amsterdam (b) Rotterdam (c) Utrecht

(d) Eindhoven (e) Groningen (f) Maastricht

(g) Tilburg (h) Enschede (i) Nijmegen

5.3.3.1 Parametric Polynomial Regression

The volatility of the process can be estimated by fitting a curve to the data points. Multiple approaches have been taken in this regard. For instance, Tindall (2006) used a polynomial with 4 degree of freedom to Sydney Airport temperature volatility. Polynomial regression is a useful tool as it allows fitting[4] a p-degree polynomial to the volatility of the process. However, the choice of the degree p, i.e.the features, of the polynomial requires a parsimonious approach as to not over-

[4]Fitting is carried out using ordinary least squares

fit or under-fit the data. That is, looking at Figure 5.6 it can be seen that the second and third degree polynomials don't capture the high and low points of the data accurately, whereas the sixth and seventh degree polynomials produce absurd fitting as outlined by Carmona (2014). On the other hand, the AIC score of each polynomial are presented in Table 5.2 wherein it is suggested that the polynomial with the degree $p = 5$ fits the data best and hence is a good estimator and predictor of volatility.

Table 5.2: AIC results of the regeression polynomials

	$p=2$	$p=3$	$p=4$	$p=5$	$p=6$	$p=7$
Amsterdam	367.05	327.91	251.85	235.40	810.62	1148.49
Rotterdam	324.73	285.58	241.97	221.01	815.55	1159.63
Utrecht	379.23	347.68	280.71	260.72	860.32	1190.76
Eindhoven	379.52	355.21	282.83	263.19	873.72	1202.51
Groningen	428.94	404.40	310.97	296.34	889.64	1213.93
Maastricht	346.54	334.14	303.69	290.48	878.65	1207.93
Tilburg	386.70	360.19	291.07	272.73	877.85	1215.20
Enschede	395.25	375.77	315.49	304.94	910.98	1235.44
Nijmegen	392.71	369.04	283.42	272.23	882.31	1214.73

Figure 5.6: Polynomial volatility curves

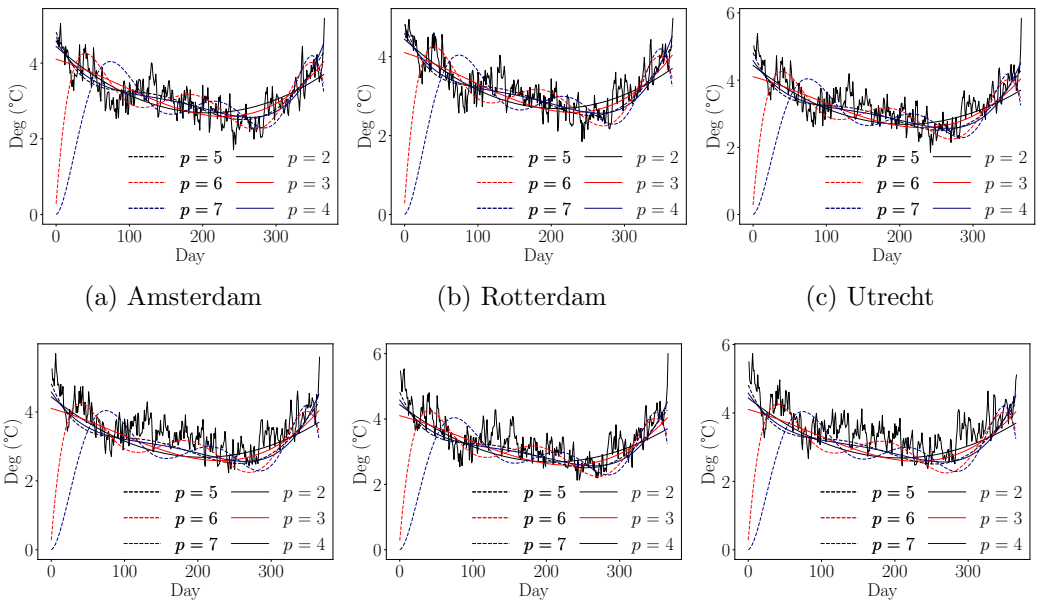

(a) Amsterdam (b) Rotterdam (c) Utrecht

(d) Eindhoven (e) Groningen (f) Maastricht

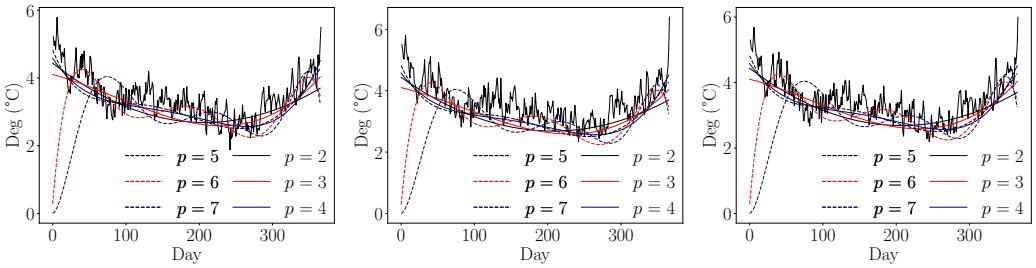

(g) Tilburg (h) Enschede (i) Nijmegen

5.3.3.2 Non-Parametric Basis Spline Interpolation

It was mentioned in the last subsection that the sixth and seventh degree polynomials over-fit the data, in fact one of the drawbacks of polynomial regression is that it tends to have a very high variance at the boundaries and hence produce inaccuracies (Hastie et al., 2009). Moreover, the polynomial regression generates a single polynomial that describes the entire volatility dataset. On the other hand, Spline regression can be used to yield a piecewise continuous function composed of multiple polynomials.

To achieve this, Basis-splines[5] (B-splines) can be used. A B-spline of order k uses several Bézier curves joined on end and is defined by

$$B(t) = \sum_{i=0}^{n} \beta_i N_{i,k}(t) \tag{5.39}$$

where β_i are de Boor points and $N_{i,k}(t)$ are basis functions defined using the following Cox-de Boor recursion formula

$$\begin{aligned} N_{i,0}(t) &= \begin{cases} 1 & \text{it}_i < t < t_{i+1} \\ 0 & \text{otherwise} \end{cases} \\ N_{i,j}(t) &= \frac{t - t_i}{t_{i+j} - t_i} N_{i,j-1}(t) + \frac{t_{i+j+1} - t}{t_{i+j+1} - t_{i+1}} N_{i+1,j-1}(t) \end{aligned} \tag{5.40}$$

where $t = \{t_i | i \in \mathbb{Z}\}$ is a sequence of non-decreasing real numbers called knot sequence[6]. In fact, constructing a smoother spline fit requires specifying the number of knots for the target data. That is, it can be seen from Table 5.3 that although a higher number of knots produces a lower residual sum of squares (RSS) data could still be subject to over-fitting as conveyed by the 50 knots B-spline in Figure 5.7.

[5]Basis-spline is a curve approximation method
[6]Knots are joints of polynomial segments

Figure 5.7: B-spline volatility curve fit

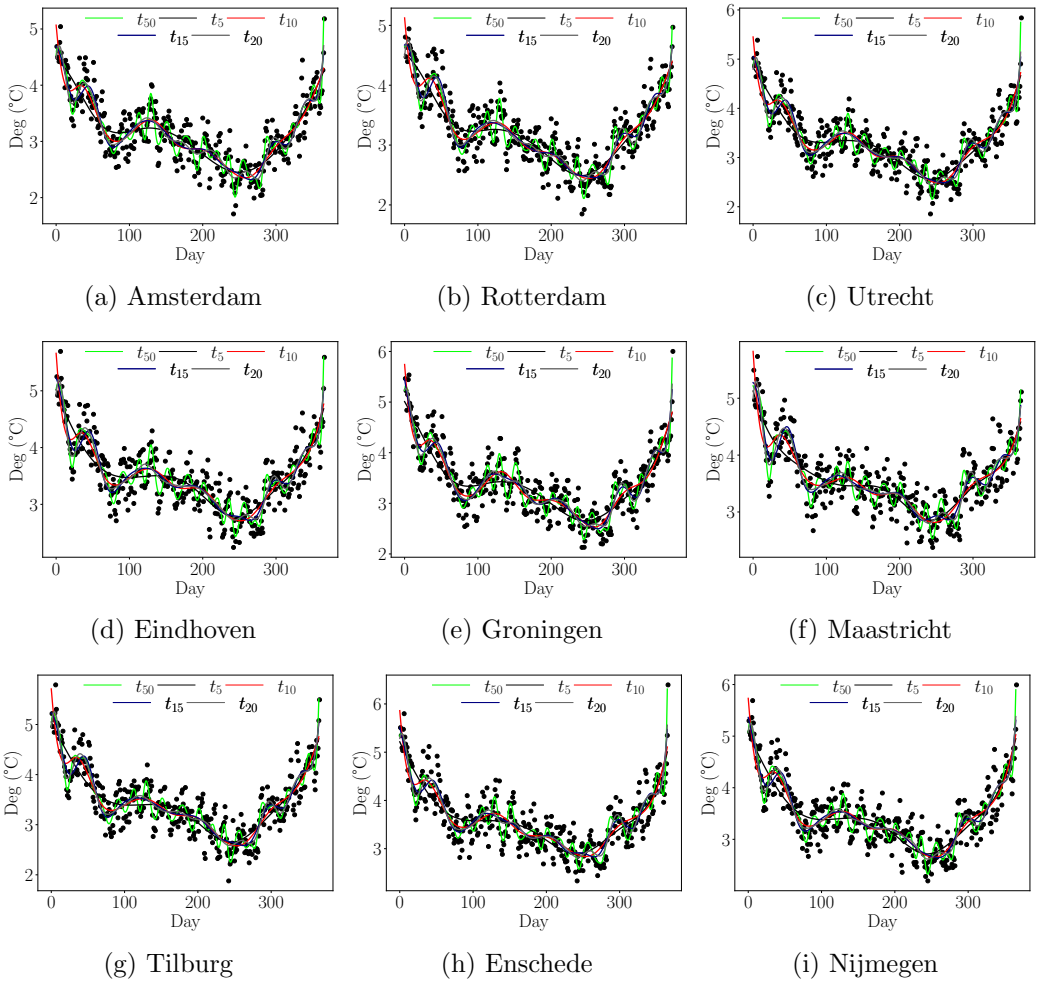

(a) Amsterdam (b) Rotterdam (c) Utrecht

(d) Eindhoven (e) Groningen (f) Maastricht

(g) Tilburg (h) Enschede (i) Nijmegen

Table 5.3: B-splines fitting residual sum of squares

	t_5	t_{10}	t_{15}	t_{20}	t_{50}
Amsterdam	37.364	32.140	29.202	27.410	18.184
Rotterdam	35.442	30.705	27.152	25.696	18.094
Utrecht	39.750	33.652	30.46	28.551	19.476
Eindhoven	40.702	35.186	31.938	29.977	21.123
Groningen	44.560	36.237	32.373	30.512	20.396
Maastricht	43.653	37.721	32.747	30.612	22.338
Tilburg	41.130	35.418	31.560	29.589	21.100
Enschede	45.750	39.908	36.305	34.075	23.538
Nijmegen	41.218	35.081	32.616	30.601	21.390

5.3.3.3 Step Function Approximation

Another method to estimate volatility is done with piecewise constant function. More specifically, Alaton et al. (2002) use piecewise construction of constant monthly

volatilities. In fact, the framework of Alaton et al. (2002) assumes that the quadratic variation $\sigma^2 \in \mathbb{R}_+$ remains approximately constant throughout the month but varies across months. Hence, the volatility is a piecewise constant function defined by

$$\hat{\sigma_\mu}^2 = \frac{1}{N_\mu} + \sum_{i=0}^{N_\mu - 1} (T_{i+1} - T_i)^2 \tag{5.41}$$

Where μ is a monthly index, N_μ is the number of days in the month, and T_i are the DATs of the month $\forall i \in [1, \cdots, N_\mu]$. As it can be seen in Figure 5.8 , the monthly approximation of variation provides a good measure of volatility of the temperature process.

Figure 5.8: Piecewise constant function volatility fit

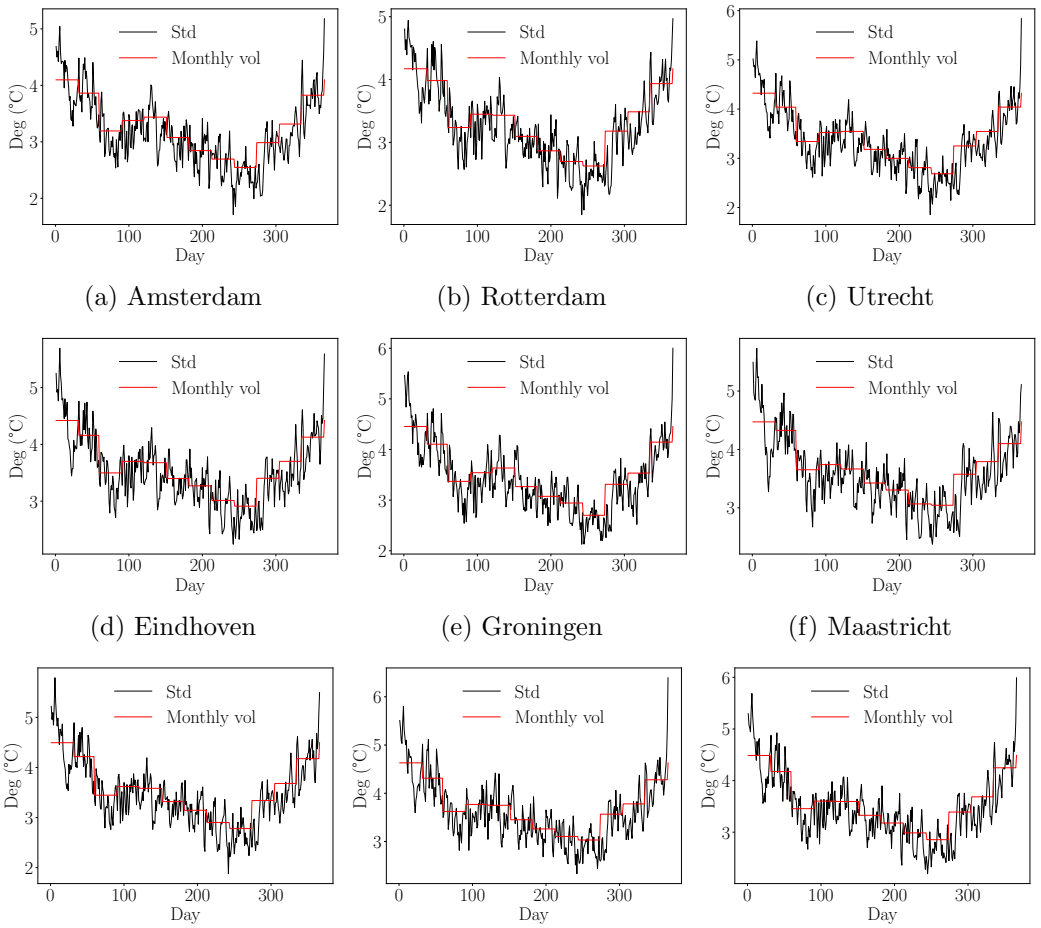

(a) Amsterdam (b) Rotterdam (c) Utrecht

(d) Eindhoven (e) Groningen (f) Maastricht

(g) Tilburg (h) Enschede (i) Nijmegen

5.3.3.4 Stochastic Model of Volatility

A different way to estimate volatility is to model it as a separate OU process. This approach was introduced by Bhowan (2003) where, following the same assumption

of Alaton et al. (2002), the volatility process is of the form

$$d\sigma_\mu = \kappa_\sigma(\widetilde{\sigma} - \sigma_\mu)d\mu + \xi_\sigma dW_\mu \tag{5.42}$$

where, once again, μ is a monthly index, $\widetilde{\sigma}$ is the long-term volatility trend, κ_σ is the speed of mean reversion, and W_μ is a Wiener process. In Equation 5.42, $\widetilde{\sigma}$ is taken to be constant as presented in Figure 5.9[7].

Figure 5.9: Monthly volatility of DAT

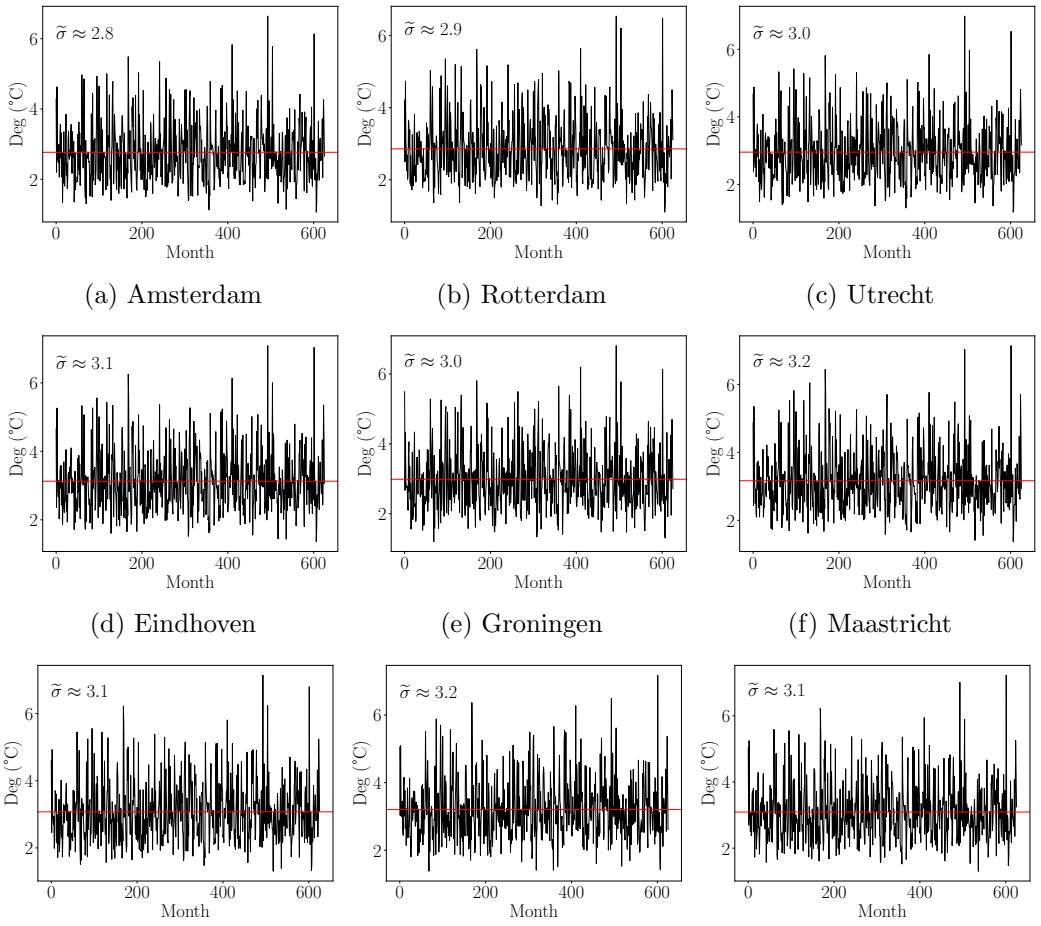

(a) Amsterdam (b) Rotterdam (c) Utrecht

(d) Eindhoven (e) Groningen (f) Maastricht

(g) Tilburg (h) Enschede (i) Nijmegen

Consequently, estimating Equation 5.42, consists in estimating ξ_σ and κ_σ, where the former can be estimated as the quadratic variation of σ_μ. That is,

$$\hat{\xi_\sigma}^2 = \frac{1}{N} + \sum_{i=0}^{N-1}(\sigma_{i+1} - \sigma_i)^2 \tag{5.43}$$

where N is the number of observations, and σ_i are the monthly volatilities $\forall i \in$

[7]The difference between this figure and figure 4.6 is due to the grouping of the data. In the latter data is grouped by both the year and the month, while in the former was grouped on a monthly basis only.

$[1, \cdots, N]$. The calculation gives

$$\hat{\xi_\sigma} = \begin{cases} 0.828 & \text{in Amsterdam} \\ 0.842 & \text{in Rotterdam} \\ 0.858 & \text{in Utrecht} \\ 0.868 & \text{in Eindhoven} \\ 0.870 & \text{in Groningen} \\ 0.871 & \text{in Maastricht} \\ 0.890 & \text{in Tilburg} \\ 0.882 & \text{in Enschede} \\ 0.878 & \text{in Nijmegen} \end{cases}$$

In the original paper, Bhowan (2003) uses a martingale estimation function to estimate κ_σ, based on the work of Bibby and Sørensen (1995), whereby SDEs of the form

$$dX_t = b(X_t, \theta)dt + \sigma(X_t, \theta)dW_t; \quad X_0 = x_0$$

allow for an unbiased estimator of θ given by

$$G_n(\theta) = \sum_{i=1}^{n} \frac{\dot{b}(X_{(i-1)\Delta}, \theta)}{\sigma^2(X_{(i-1)\Delta}, \theta)}(X_{i\Delta} - \mathbb{E}[X_i \mid X_{i-1}])$$

where $\dot{b} \equiv \frac{\partial b}{\partial \theta}$.

Nevertheless, since it was show in Figure 5.9 that $\mathbb{E}[\sigma_\mu] = \widetilde{\sigma}$ and that $\widetilde{\sigma}$ is constant across time, then the process given by Equation 5.42 can be modeled as another modified OU process of the form

$$d\sigma_\mu = \left(\frac{d\widetilde{\sigma}}{d\mu} + \kappa_\sigma(\widetilde{\sigma} - \sigma_\mu)\right) d\mu + \xi_\sigma dW_\mu \tag{5.44}$$

following the same approach in Subsection 5.3.2, the rate of speed reversion parameter κ_σ can be estimated by autoregressing on the residual of the volatility parameter, that is

$$\Lambda_{t_i} = \gamma_\sigma \Lambda_{t_{i-1}} + e_{t_i} \tag{5.45}$$

where $\kappa_\sigma = 1 - \gamma_\sigma$

Table 5.4: Autoregressive model results

City:		Amsterdam	**No. Observations:**			19034
Dep. Variable:		Standard deviation	**Log Likelihood**			-914.466
Model:		AR(1)	**AIC**			1832.931
Method:		Conditional MLE	**BIC**			1841.807
	Coef	**Std err**	**z**	$p > \|z\|$	**[0.025**	**0.975]**
std.L1	0.9311	0.014	64.319	0.000	0.903	0.959
City:		Rotterdam	**No. Observations:**			19034
Dep. Variable:		Standard deviation	**Log Likelihood**			-918.242
Model:		AR(1)	**AIC**			1840.484
Method:		Conditional MLE	**BIC**			1849.360
	Coef	**Std err**	**z**	$p > \|z\|$	**[0.025**	**0.975]**
std.L1	0.9349	0.014	66.213	0.000	0.907	0.963
City:		Utrecht	**No. Observations:**			19034
Dep. Variable:		Standard deviation	**Log Likelihood**			-940.370
Model:		AR(1)	**AIC**			1884.740
Method:		Conditional MLE	**BIC**			1893.616
	Coef	**Std err**	**z**	$p > \|z\|$	**[0.025**	**0.975]**
std.L1	0.9342	0.014	65.976	0.000	0.906	0.962
City:		Eindhoven	**No. Observations:**			19034
Dep. Variable:		Standard deviation	**Log Likelihood**			-955.426
Model:		AR(1)	**AIC**			1914.852
Method:		Conditional MLE	**BIC**			1923.727
	Coef	**Std err**	**z**	$p > \|z\|$	**[0.025**	**0.975]**
std.L1	0.9381	0.014	68.223	0.000	0.911	0.965
City:		Groningen	**No. Observations:**			19034
Dep. Variable:		Standard deviation	**Log Likelihood**			-960.194
Model:		AR(1)	**AIC**			1924.389
Method:		Conditional MLE	**BIC**			1933.264
	Coef	**Std err**	**z**	$p > \|z\|$	**[0.025**	**0.975]**
std.L1	0.9304	0.014	64.373	0.000	0.902	0.959
City:		Maastricht	**No. Observations:**			19034
Dep. Variable:		Standard deviation	**Log Likelihood**			-959.196
Model:		AR(1)	**AIC**			1922.392
Method:		Conditional MLE	**BIC**			1931.267
	Coef	**Std err**	**z**	$p > \|z\|$	**[0.025**	**0.975]**
std.L1	0.9385	0.014	68.622	0.000	0.912	0.965
City:		Tilburg	**No. Observations:**			19034
Dep. Variable:		Standard deviation	**Log Likelihood**			-956.825
Model:		AR(1)	**AIC**			1917.650
Method:		Conditional MLE	**BIC**			1926.526
	Coef	**Std err**	**z**	$p > \|z\|$	**[0.025**	**0.975]**
std.L1	0.9359	0.014	66.927	0.000	0.909	0.963

City:		Enschede	**No. Observations:**			19034
Dep. Variable:		Standard deviation	**Log Likelihood**			-962.227
Model:		AR(1)	**AIC**			1928.454
Method:		Conditional MLE	**BIC**			1937.330
	Coef	**Std err**	**z**	$p > \|z\|$	**[0.025**	**0.975]**
std.L1	0.9392	0.014	69.154	0.000	0.913	0.966
City:		Nijmegen	**No. Observations:**			19034
Dep. Variable:		Standard deviation	**Log Likelihood**			-953.283
Model:		AR(1)	**AIC**			1910.567
Method:		Conditional MLE	**BIC**			1919.442
	Coef	**Std err**	**z**	$p > \|z\|$	**[0.025**	**0.975]**
std.L1	0.9370	0.014	67.709	0.000	0.910	0.964

Consequently, the volatility dynamics in this framework are given by

$$d\sigma_\mu = \begin{cases} 0.93 + (2.80 - \sigma_\mu)d_\mu + 0.83dW_\mu & \text{in Amsterdam} \\ 0.93 + (2.90 - \sigma_\mu)d_\mu + 0.84dW_\mu & \text{in Rotterdam} \\ 0.93 + (3.00 - \sigma_\mu)d_\mu + 0.86dW_\mu & \text{in Utrecht} \\ 0.94 + (3.10 - \sigma_\mu)d_\mu + 0.83dW_\mu & \text{in Eindhoven} \\ 0.93 + (3.00 - \sigma_\mu)d_\mu + 0.87dW_\mu & \text{in Groningen} \\ 0.94 + (3.20 - \sigma_\mu)d_\mu + 0.87dW_\mu & \text{in Maastricht} \\ 0.94 + (3.10 - \sigma_\mu)d_\mu + 0.89dW_\mu & \text{in Tilburg} \\ 0.94 + (3.20 - \sigma_\mu)d_\mu + 0.88dW_\mu & \text{in Enschede} \\ 0.94 + (3.10 - \sigma_\mu)d_\mu + 0.88dW_\mu & \text{in Nijmegen} \end{cases}$$

5.3.3.5 Fourier Series Model of Volatility

A more advanced approach consists in estimating volatility as a truncated Fourier series. Following the work of Benth and Benth (2007), and Benth et al. (2007) who argue that this method allows for more flexibility in calculations and captures the observed stylized facts of temperatures, volatility can be represented by

$$\sigma_t^2 = \beta_0 + \sum_{i=1}^{I} \beta_i \sin(i\omega t) + \sum_{j=1}^{J} \alpha_j \cos(j\omega t) \tag{5.46}$$

where $\omega \approx 2\pi/365.25$. In Esunge and Njong (2020), the constant β was replaced by a linear term such that $\beta_0 = V + Ut$, to capture the higher volatility in the winter periods compared to the summer periods. Nevertheless, upon investigating the rolling volatility of the DATs it was found, as Figure 5.10 points out, that there is no linear trend to account for. Hence, Equation 5.46 is left unchanged.

Figure 5.10: Long term volatility of the DAT

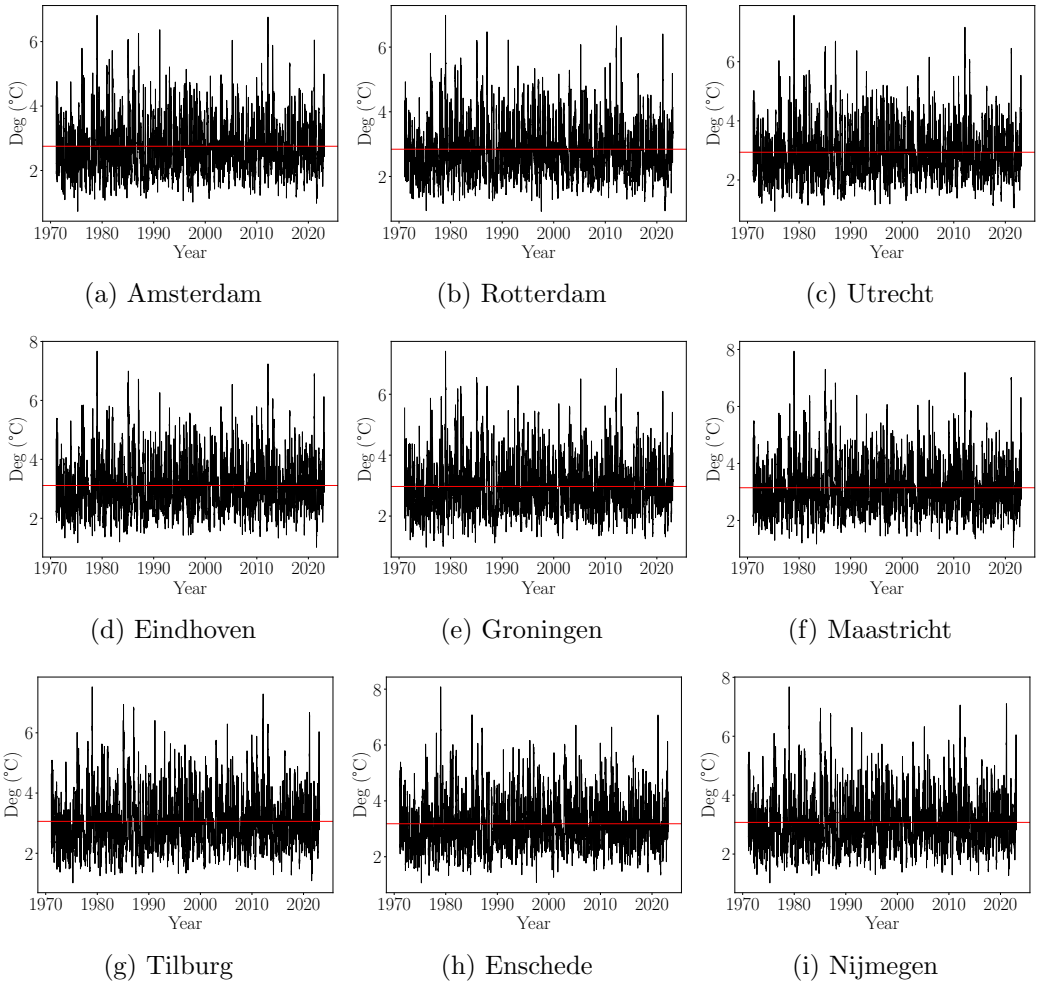

(a) Amsterdam (b) Rotterdam (c) Utrecht

(d) Eindhoven (e) Groningen (f) Maastricht

(g) Tilburg (h) Enschede (i) Nijmegen

While working on Swedish data, Benth and Benth (2007) and Benth et al. (2007) arbitrarily use a 4^{th} order Fourier series, i.e. setting $I = J = 4$ in Equation 5.46. Nevertheless, to chose the optimal model, this thesis fits fits Fourier series with orders from 1 to 15 and inspects the RSS of each fit as depicted in Figure 5.11. It can be seen that the RSS drops until order 5 with a minimal increase in performance increase at order 10. Hence, It is concluded that Dutch data should be fitted using a Fourier series of order 10, the result of this fit is presented in Figure 5.12.

Figure 5.11: Residual sum of squares of Fourier orders

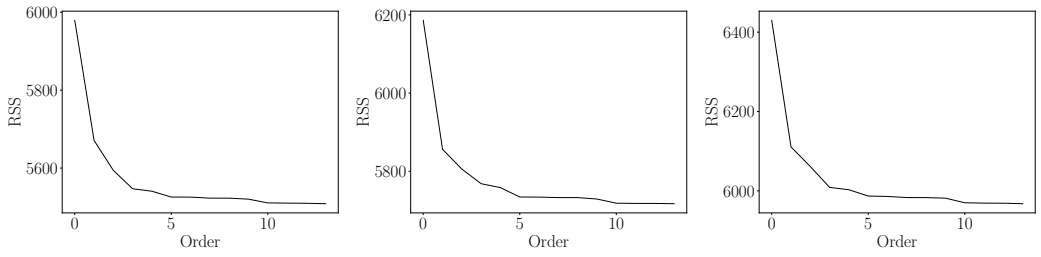

(a) Amsterdam (b) Rotterdam (c) Utrecht

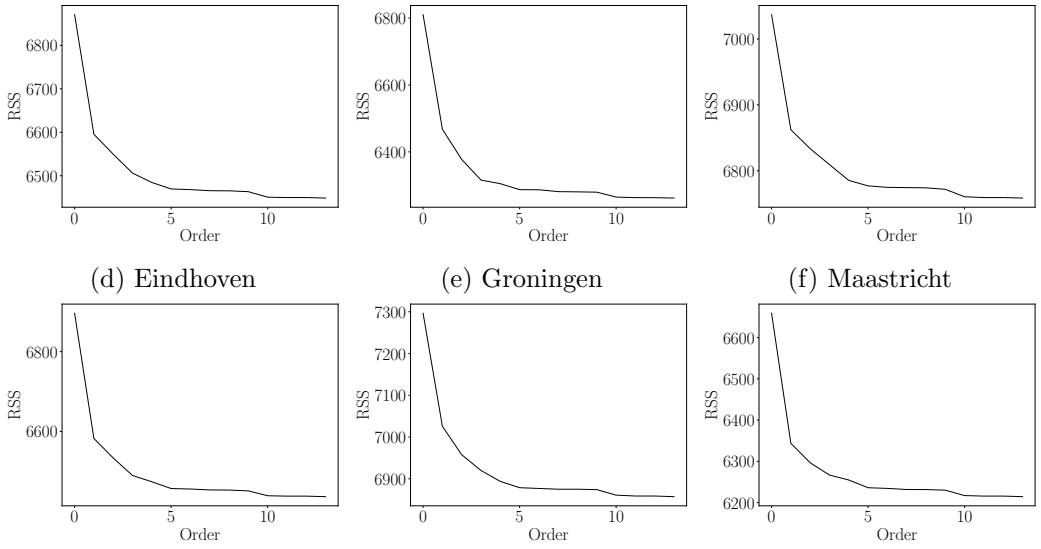

(d) Eindhoven (e) Groningen (f) Maastricht

(g) Tilburg (h) Enschede (i) Nijmegen

Figure 5.12: Fourier series volatility fit

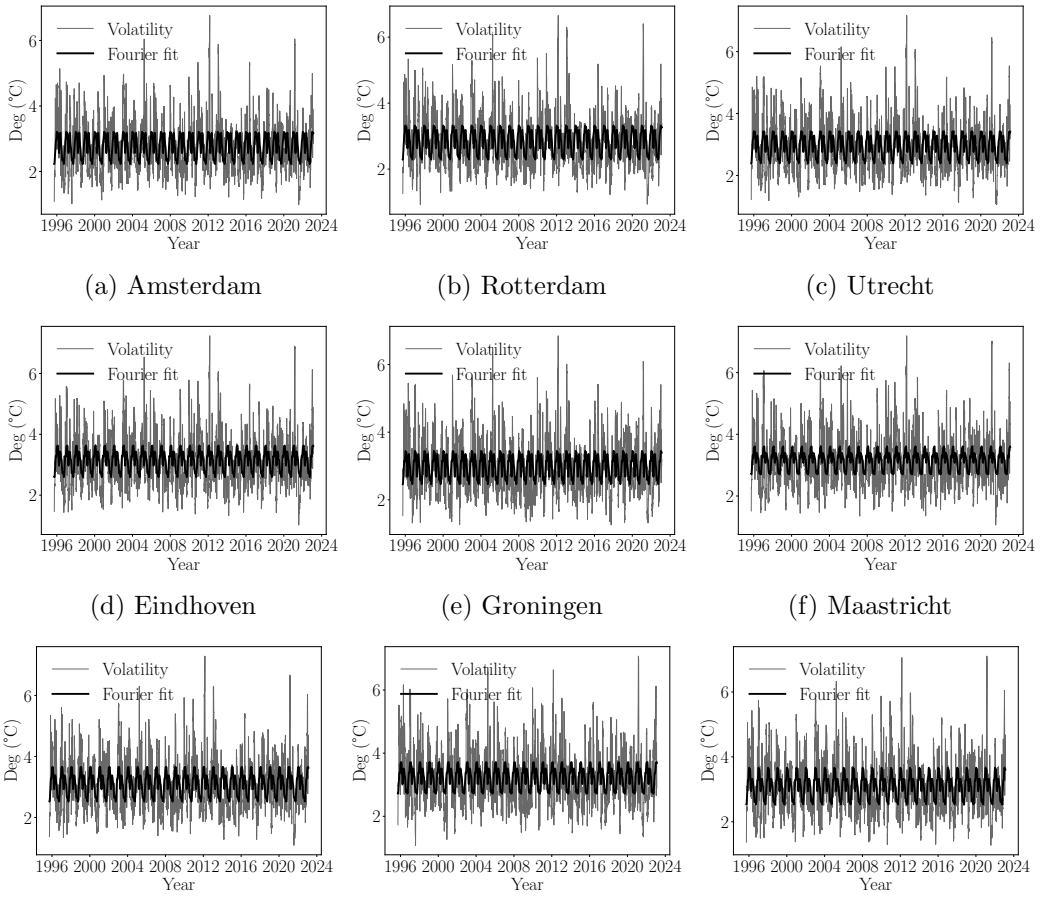

(a) Amsterdam (b) Rotterdam (c) Utrecht

(d) Eindhoven (e) Groningen (f) Maastricht

(g) Tilburg (h) Enschede (i) Nijmegen

Daily Average Temperature Simulation Under the $\mathbb{P}$-Measure

With the DATs dynamics on hand, the last step prior to pricing temperature derivatives for the Netherlands is to simulate DATs paths under the real world $\mathbb{P}$-measure using those dynamics. This approach does not assume the existence of a risk free rate or an arbitrage free setting in the market. Hence it is *backwards looking*.

6.1 Underlying Dynamics

In the last chapter a thorough modeling of the DATs was carried out. In summary, the modified OU process dynamics are as follows:

$$dT_t = \begin{cases} (d\theta_t/dt + 0.206(\theta_t - T_t))\,dt + \sigma_t dW_t & \text{in Amsterdam} \\ (d\theta_t/dt + 0.211(\theta_t - T_t))\,dt + \sigma_t dW_t & \text{in Rotterdam} \\ (d\theta_t/dt + 0.209(\theta_t - T_t))\,dt + \sigma_t dW_t & \text{in Utrecht} \\ (d\theta_t/dt + 0.213(\theta_t - T_t))\,dt + \sigma_t dW_t & \text{in Eindhoven} \\ (d\theta_t/dt + 0.214(\theta_t - T_t))\,dt + \sigma_t dW_t & \text{in Groningen} \\ (d\theta_t/dt + 0.194(\theta_t - T_t))\,dt + \sigma_t dW_t & \text{in Maastricht} \\ (d\theta_t/dt + 0.213(\theta_t - T_t))\,dt + \sigma_t dW_t & \text{in Tilburg} \\ (d\theta_t/dt + 0.213(\theta_t - T_t))\,dt + \sigma_t dW_t & \text{in Enschede} \\ (d\theta_t/dt + 0.213(\theta_t - T_t))\,dt + \sigma_t dW_t & \text{in Nijmegen} \end{cases}$$

where

$$\theta_t = \begin{cases} 9.22 + 1.07\text{e-}4t + 7.31\sin(\omega t - 1.95) & \text{in Amsterdam} \\ 9.27 + 1.05\text{e-}4t + 7.21\sin(\omega t - 1.96) & \text{in Rotterdam} \\ 9.17 + 9.89\text{e-}5t + 7.42\sin(\omega t - 1.93) & \text{in Utrecht} \\ 9.36 + 9.73\text{e-}5t + 7.67\sin(\omega t - 1.91) & \text{in Eindhoven} \\ 8.40 + 9.98\text{e-}5t + 7.54\sin(\omega t - 1.94) & \text{in Groningen} \\ 9.25 + 1.13\text{e-}4t + 7.90\sin(\omega t - 1.90) & \text{in Maastricht} \\ 9.24 + 1.01\text{e-}4t + 7.57\sin(\omega t - 1.92) & \text{in Tilburg} \\ 8.59 + 1.04\text{e-}4t + 7.66\sin(\omega t - 1.91) & \text{in Enschede} \\ 9.17 + 9.54\text{e-}5t + 7.72\sin(\omega t - 1.91) & \text{in Nijmegen} \end{cases}$$

6.2 Monte Carlo Simulation

In order to run a Monte Carlo simulation of the future temperatures in the Netherlands, we use the Ornstein-Uhlenbeck process defined in Equation5.28 which is approximated as

$$T_{i+1} \approx T_i + \frac{d\theta_t}{dt} + \kappa(\theta_{t_{i-1}} - T_{t_{i-1}}) + \sigma_{t_i} Z_t \tag{6.1}$$

where

$$\frac{d\theta_t}{dt} = \begin{cases} 1.07\text{e-}4 + 7.31(\omega)\cos(\omega t - 1.95) & \text{in Amsterdam} \\ 1.05\text{e-}4 + 7.21(\omega)\cos(\omega t - 1.96) & \text{in Rotterdam} \\ 9.89\text{e-}5 + 7.42(\omega)\cos(\omega t - 1.93) & \text{in Utrecht} \\ 9.73\text{e-}5 + 7.67(\omega)\cos(\omega t - 1.91) & \text{in Eindhoven} \\ 9.98\text{e-}5 + 7.54(\omega)\cos(\omega t - 1.94) & \text{in Groningen} \\ 1.13\text{e-}4 + 7.90(\omega)\cos(\omega t - 1.90) & \text{in Maastricht} \\ 1.01\text{e-}4 + 7.57(\omega)\cos(\omega t - 1.92) & \text{in Tilburg} \\ 1.04\text{e-}4 + 7.66(\omega)\cos(\omega t - 1.91) & \text{in Enschede} \\ 9.54\text{e-}5 + 7.72(\omega)\cos(\omega t - 1.91) & \text{in Nijmegen} \end{cases}$$

Although it is not clear which estimator from subsection 5.3.3 is best to approximate volatility, we believe that a spline with 20 knots performs better than advanced approaches such as Fourier series approximation. That is, looking at Figure 6.1 it is clear that the method suggested by Benth and Benth (2007) and Benth et al. (2007) underfits Dutch temperature data.

Figure 6.1: Comparison of Fourier series and B-spline volatility fit

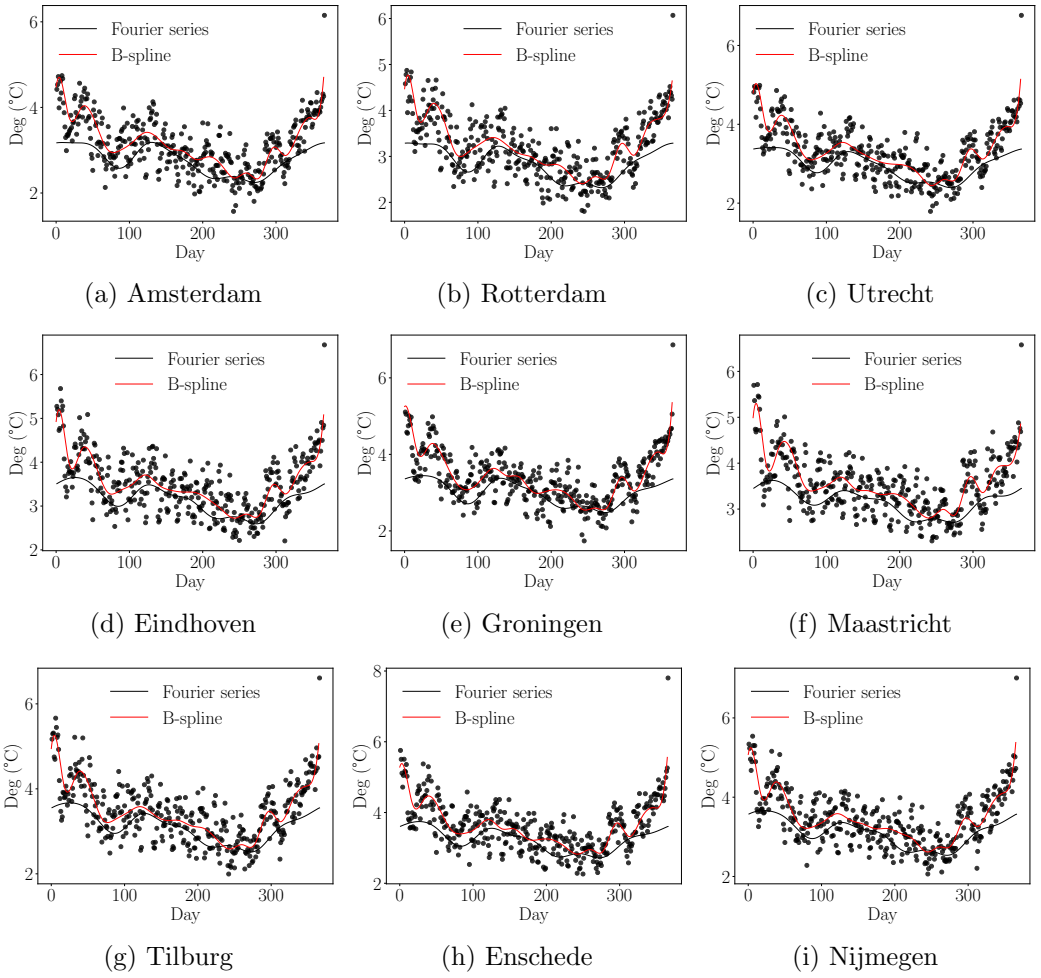

(a) Amsterdam (b) Rotterdam (c) Utrecht

(d) Eindhoven (e) Groningen (f) Maastricht

(g) Tilburg (h) Enschede (i) Nijmegen

Based on these dynamics, 10.000 scenarios of DATs between the 1st of January 2024 and the 31st of December 2028 are simulated. The results of these simulations are shown in Figure 6.2, where the x-axis is the date $i \; \forall i \in [2024/01/01, \cdots, 2028/12/31]$, the y-axis represents the different iterations from 0 to 10.000, and the z-axis corresponds to a given DAT value, hence a position along the z-axis indicates the DAT for a given simulation.

Figure 6.2: Temperature simulations from 2024 to 2029

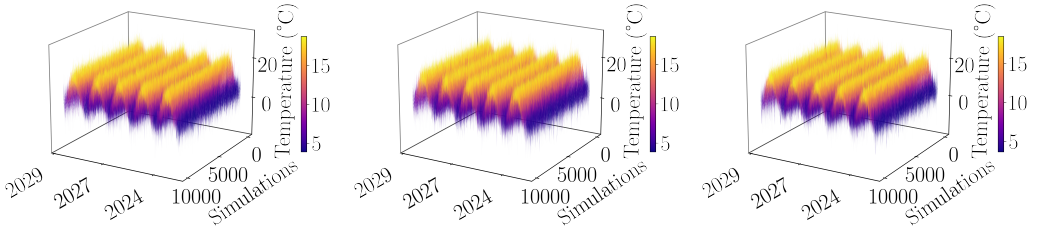

(a) Amsterdam (b) Rotterdam (c) Utrecht

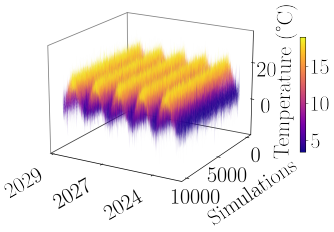

(d) Eindhoven

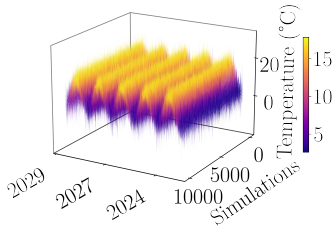

(e) Groningen

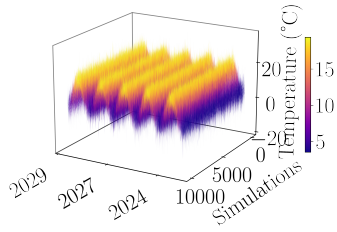

(f) Maastricht

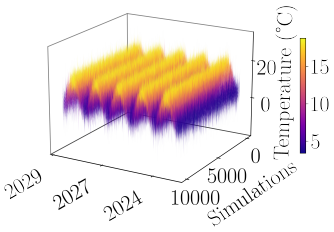

(g) Tilburg

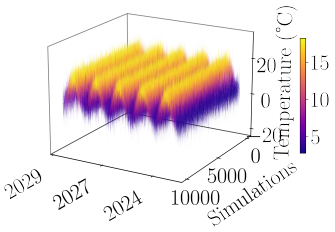

(h) Enschede

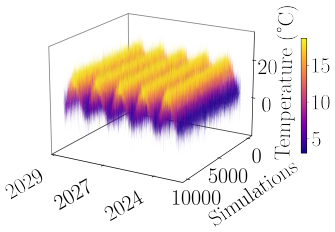

(i) Nijmegen

Putting a Price on Dutch Temperature Under the $\mathbb{Q}$-Measure

As mentioned in the beginning of this thesis, standardized European temperature derivatives are written on the cumulative average temperature (CAT) index, or the heating degree day (HDD) index. The aim of this chapter is price HDD and CAT options for the nine Dutch cities under the risk neutral $\mathbb{Q}$-measure.

7.1 Solving the Ornstein Uhlenbeck SDE

Consider the DAT modified OU process given by the following SDE

$$dT_t = \left(\frac{d\theta_t}{dt} + \kappa(\theta_t - T_t) \right) dt + \sigma_t dB_t \tag{7.1}$$

If we set $\widetilde{T} = T_t - \theta_t$, then, following the proof to proposition 1, an explicit solution is given by

$$T_t = \theta_t + e^{-\kappa(t-s)}(T_s - \theta_s) + \int_s^t \sigma_u e^{-\kappa(t-u)} dB_u \tag{7.2}$$

Therefore, the following is also true

$$\mathbb{E}^{\mathbb{P}}[T_t|\mathcal{F}_s] = \theta_t + e^{-\kappa(t-s)}(T_s - \theta_s)$$
$$\mathbb{V}ar^{\mathbb{P}}[T_t|\mathcal{F}_s] = \int_s^t \sigma_u^2 e^{-2\kappa(t-u)} dB_u$$

This approach assumes that the driving noise process follows a Brownian motion, some authors such as Benth and Šaltytė-Benth (2005) price temperature derivatives where a Lévy motion is the driving noise process.

The weather derivatives market is an incomplete market with infinite equivalent martingale measures (Alexandridis & Zapranis, 2012). Moreover, since temperature is non-tradable, the market price of risk should be accounted for in the model in order to obtain unique prices for the contracts. In this thesis the pricing is carried out under the parametric risk neutral probability $\mathbb{Q}^{\lambda}$, where λ is a constant market price of risk.

Let $W = (W_t, t \geq 0)$ be a $\mathbb{Q}$-standard Wiener process, then the price process

can be written as

$$dT_t = \left(\frac{d\theta_t}{dt} + \kappa(\theta_t - T_t) - \lambda\sigma_t \right) dt + \sigma_t dW_t^{\mathbb{Q}} \tag{7.3}$$

Then, similarly to Equation 7.1, it holds that

$$T_t = \theta_t + e^{-\kappa(t-s)}(T_s - \theta_s) - \lambda \int_s^t \sigma_u e^{-\kappa(t-u)} du + \int_s^t \sigma_u e^{-\kappa(t-u)} dB_u \tag{7.4}$$

Therefore

$$\mathbb{E}^{\mathbb{Q}}[T_t|\mathcal{F}_s] = \mathbb{E}^{\mathbb{P}}[T_t|\mathcal{F}_s] - \lambda \int_s^t \sigma_u e^{-\kappa(t-u)} du \tag{7.5}$$

$$\mathbb{V}ar^{\mathbb{Q}}[T_t|\mathcal{F}_s] = \int_s^t \sigma_u^2 e^{-2\kappa(t-u)} dB_u \tag{7.6}$$

Additionally, as long as volatility $\sigma_i \forall i \in [s,t]$ is constant throughout the interval $[s,t]$, the expectation, variance, and covariance are given by

$$\mathbb{E}^{\mathbb{Q}}[T_t|\mathcal{F}_s] = \mathbb{E}^{\mathbb{P}}[T_t|\mathcal{F}_s] - \frac{\lambda\sigma_i}{\kappa}(1 - e^{-\kappa(t-s)}) \tag{7.7}$$

$$\mathbb{V}ar^{\mathbb{Q}}[T_t|\mathcal{F}_s] = \frac{\sigma_i^2}{2\kappa}(1 - e^{-2\kappa(t-s)}) \tag{7.8}$$

$$\mathbb{C}ov^{\mathbb{Q}}[T_t, T_u|\mathcal{F}_s] = e^{-\kappa(u-t)}\mathbb{V}ar^{\mathbb{Q}}[T_t|\mathcal{F}_s] \tag{7.9}$$

where $0 \leq s \leq t \leq u$

7.2 Pricing Dutch Heating Degree Day Options

7.2.1 Alaton Approximation

The payoff of a HDD index is given by

$$\phi = \alpha(HDD_n - K)^+ \tag{7.10}$$

where α is the tick size, K is the strike, and

$$HDD_n = \sum_{n=1}^{N} max(0, T_{ref} - DAT_n) \tag{7.11}$$

Since the reference temperature of the weather derivatives contracts traded in the CME is 18°C, Equation 7.11 becomes

$$HDD_n = \sum_{n=1}^{N} max(0, 18 - DAT_n) \tag{7.12}$$

The underlying temperature process is normally distributed. Nevertheless, under this assumption, finding a pricing formula is challenging because of the maximum function. On the other hand, Alaton et al. (2002) provide an approximation whereby under the risk neutral $\mathbb{Q}$-measure conditional on in formation at time s

$$T_t \sim \mathcal{N}(\mu_t, \sigma_t^2)$$

where $\mu_t = \mathbb{E}^{\mathbb{Q}}[T_t|\mathcal{F}_s]$ and $\sigma_t^2 = \mathbb{V}ar^{\mathbb{Q}}[T_t|\mathcal{F}_s]$. Moreover, to the observation made by Alaton et al. (2002) when working on Swedish data, it was found that winter periods in the Netherlands satisfy the following

$$\mathbb{P}[max(0, 18 - DAT_n) = 0] \approx 0$$

Hence, for Dutch HDD contracts, we have that

$$HDD_n = 18n - \sum_{i=1}^{n} T_{t_i}; \tag{7.13}$$

It is clear that the distribution of Equation 7.13 is easier to determine without the maximum function. In fact, since T_{t_i} come from a normally distributed OU process, it holds that $[T_{t_1}, T_{t_2}, \cdots, T_{t_n}]$ is also normally distributed. Therefore, HDD_n is also Gaussian with mean

$$\mathbb{E}^{\mathbb{Q}}[HDD_n|\mathcal{F}_t] = \mathbb{E}^{\mathbb{Q}}\left[18n - \sum_{i=1}^{n} T_{t_i}|\mathcal{F}_t\right] = 18n - \sum_{i=1}^{n} T_{t_i}\mathbb{E}^{\mathbb{Q}}[T_{t_i}|\mathcal{F}_t] \tag{7.14}$$

and variance

$$\mathbb{V}ar^{\mathbb{Q}}[HDD_n|\mathcal{F}_t] = \sum_{i=1}^{n} T_{t_i}\mathbb{V}ar^{\mathbb{Q}}[T_{t_i}|\mathcal{F}_t] + 2\sum\sum_{i<j} \mathbb{C}ov[T_{t_i}, T_{t_j}|\mathcal{F}_t] \tag{7.15}$$

According to the fundamental theorem of asset pricing, the absence of arbitrage, i.e. under an equivalent martingale measure, implies the existence of a state price density which in turn implies the existence of risk neutral probabilities, therefore the price of a HDD call option at time $t \leq t_1$ is

$$\begin{aligned} C_{HDD}(t) &= e^{-r(t_n-t)}\mathbb{E}^{\mathbb{Q}}[\alpha \max\{HDD_n - K, 0\}|\mathcal{F}_t] \\ &= e^{-r(t_n-t)} \int_K^{\infty} (x-k) f_{HDD_n}(x) dx \\ &= e^{-r(t_n-t)} \left((\mu_n - K)\Phi(-\xi_n) + \frac{\sigma_n}{\sqrt{2\pi}} e^{-\frac{\xi_n^2}{2}} \right) \end{aligned} \tag{7.16}$$

where $e^{-r(t_n-t)}$ is a stochastic discount factor, K is the strike, $\xi_n = (K - \mu_n)/\sigma_n$, and Φ is the cumulative distribution function of the standard normal distribution.

Similarly, the price of a HDD put option at time $t \leq t_1$ is

$$\begin{aligned} P_{HDD}(t) &= e^{-r(t_n-t)}\mathbb{E}^{\mathbb{Q}}[\alpha \max\{K - HDD_n, 0\}|\mathcal{F}_t] \\ &= e^{-r(t_n-t)} \int_0^K (K-x) f_{HDD_n}(x) dx \\ &= e^{-r(t_n-t)} \left[(K-\mu_n)\left(\Phi(\xi_n) - \Phi\left(-\frac{\mu_n}{\sigma_n}\right)\right) + \frac{\sigma_n}{\sqrt{2\pi}}\left(e^{-\frac{\xi_n^2}{2}} - e^{\frac{1}{2}(\frac{\mu_n}{\sigma_n})^2}\right)\right] \end{aligned} \tag{7.17}$$

7.2.1.1 Model Calibration to the Market

Before calculating prices for HDD calls and puts under the $\mathbb{Q}$-measure, the model needs to be calibrated to the dutch market conditions.

7.2.1.1.1 Risk Free Rate

Since this work focuses on the Dutch weather market, it was decided to use the Netherlands 10-year constant maturity treasury bond yield as proxy for the risk free rate[1]. The Netherlands is one of the few European countries that have a stable credit rating of AAA according to S&P Global Ratings. Moreover, as it can be seen in Figure 7.1, the Netherlands has a one of the lowest long-term CDS spreads which is indicates a strong credit-worthiness of the Netherlands.

Figure 7.1: Credit default swaps spread

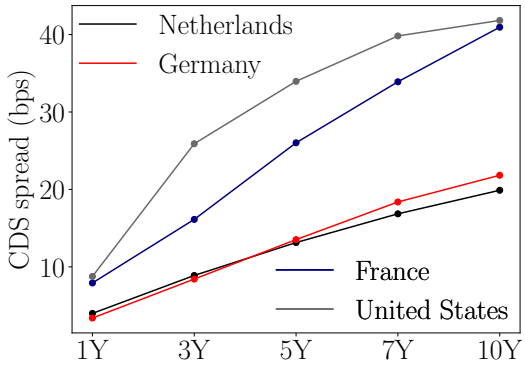

Source: FactSet.

7.2.1.2 Market Price of Risk

The market price of risk λ represents the risk preferences of market participants and their willingness to expose themselves to risk. Previous studies commonly assumed a market price of risk of zero, whereas recent research indicates the opposite. The approach commonly used in the literature is based on the work of Alaton et al.

[1] According to FactSet the yield is equal to 2.75

(2002), who suggest inferring the market price of risk from market data. That is, finding the value of λ that matches theoretical model prices with observed market prices. Although they assume, that λ is constant, subsequent research indicates the opposite. For instance, Bellini (2005), compares the theoretical future prices to the prices observed in the market under the assumption of a Lévy process, and examines the time dependence of the market price of risk and finds a relationship between λ_t and its lag. Additionally, Härdle and Cabrera (2012) estimate the implied market price of risk for German CAT derivatives in Berlin. They find a non-zero value of λ with an increasing seasonal structure as the expiration date of the temperature future increases.

In this thesis, we break away from the framework in Alaton et al. (2002) and assume a time varying λ, i.e $\lambda = \lambda_t$. Since the market price of risk can be inferred from market data then λ_t can be estimated using a root finding algorithm. That is, focusing on puts the market price of risk for a given option is the λ-value for which

$$V_p(t_0, HDD, K, T, r, \lambda) = V_p^{mkt}(K, T) \tag{7.18}$$

where $V_p(t_0, HDD, K, T, r, \lambda)$ is the predicted put price, $V_p^{mkt}(K, T)$ is the observed market price of the put, and $t_0 = 0$.
To find λ_t the Newton Raphson root-finding iteration can be used such that

$$f(\lambda) \coloneqq V_p^{mkt}(K, T) - V_p(t_0, HDD, K, T, r, \lambda) = 0 \tag{7.19}$$

Hence, since λ_t is not constant, given an initial guess λ_t^0 and $\frac{df(\lambda)}{d\lambda}$ the next approximations of λ_t^i are found as

$$\lambda_t^{i+1} = \lambda_t^i - \frac{f(\lambda_t^i)}{f'(\lambda_t^i)}; \quad \forall i \geq 0 \tag{7.20}$$

Using the CME's market prices for Amsterdam's November-March seasonal strip weather HDDs which are listed in Table 7.1, it was found, after 2000 iterations that $\lambda = -319.76$ for option 1 and $\lambda = -266.29$ for option 2.

Table 7.1: Specifications of two Amsterdam weather product

	Weather station	**Index**	**Type**	**Period**	**Price**	**Strike**
Option 1	Schipol Airport	HDD	Put	NOV-MAR 2024	€35	1700 HDDs
Option2	Schipol Airport	HDD	Put	NOV-MAR 2024	€65	1800 HDDs

The negative market prices of risk indicate that investors are willing accept a lower compensation for bearing risk compared to what is predicted by the model. That is, they are willing to take on more risk for less compensation than expected.

This risk seeking behavior might be driven by factors such as optimistic market sentiment, overconfidence, or the belief that the potential rewards outweigh the risks. Additionally, since the temperature derivatives market is incomplete, this phenomenon could be a symptom of the inefficiency of the Dutch temperature derivatives market. In other words, a negative market price of risk implies that the market prices of the securities might not accurately reflect the level of risk. In that case, early market participants can take advantage of the arbitrage opportunities that the inefficiency provides and exploit mispricings. The real underlying reasons go beyond the scope of this thesis. Nevertheless, it is important for further research to investigate this phenomenon.

The difference between the two values of λ stems from an interesting observation about the relationship between the market price of risk and either market prices or strike levels K. As it can be seen in Figure 7.2, λ increases as the market price of the put goes up. Nevertheless, there is a linear negative relationship between λ and the strike level. This can be due to the fact that investors are willing to take more risky positions for a higher payoff.

Figure 7.2: The behavior of option 1's λ

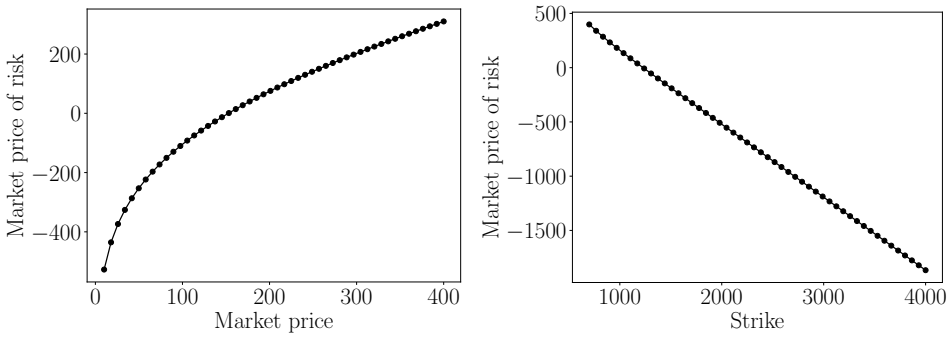

(a) Lambda values vs market prices (b) Lambda values vs strikes

7.2.2 Results

Since it was found that the market price of risk is not constant, and that it varies with respect to the strike level, to price the HDD calls and puts we will use the two λ values we found in the last subsection and their respective strikes K. The results are provided in Table 7.2. Although the distance between weather stations is relatively small, the different dynamics of the DAT of every city are such that the prices of the HDD options for Amsterdam differ those of other cities. It can be seen that the difference in the predicted price can go up to e 800 in the case Enschede's HDD call against Amsterdam's. Therefore, it can be said that although temperature is not as highly localized as wind from a meteorological perspective, it can be labeled so financially.

Table 7.2, alongside the results of Chapter 8, indicate that hedging temperature risk across the Netherlands using only Amsterdam's temperature is extremely risky as investors can be worse off if they overlook the inaccuracy of their hedging strategy.

Table 7.2: Predicted HDD options prices for the NOV-MAR 2024 strip

	Index	Type	Strike	λ	Predicted price
Amsterdam	HDD	call - put	1700 HDDs	-319.76	€630.13
			1800 HDDs	-266.29	€484.49
			1700 HDDs	-319.76	€35.00
			1800 HDDs	-266.29	€65.00
Rotterdam	HDD	call - put	1700 HDDs	-319.76	€635.07
			1800 HDDs	-266.29	€488.81
			1700 HDDs	-319.76	€38.76
			1800 HDDs	-266.29	€70.60
Utrecht	HDD	call - put	1700 HDDs	-319.76	€699.39
			1800 HDDs	-266.29	€547.05
			1700 HDDs	-319.76	€32.75
			1800 HDDs	-266.29	€60.51
Eindhoven	HDD	call - put	1700 HDDs	-319.76	€721.22
			1800 HDDs	-266.29	€566.43
			1700 HDDs	-319.76	€33.4
			1800 HDDs	-266.29	€61.40
Groningen	HDD	call - put	1700 HDDs	-319.76	€819.20
			1800 HDDs	-266.29	€657.68
			1700 HDDs	-319.76	€21.26
			1800 HDDs	-266.29	€41.06
Maastricht	HDD	call - put	1700 HDDs	-319.76	€730.10
			1800 HDDs	-266.29	€573.95
			1700 HDDs	-319.76	€34.37
			1800 HDDs	-266.29	€62.74
Tilburg	HDD	call - put	1700 HDDs	-319.76	€720.53
			1800 HDDs	-266.29	€565.83
			1700 HDDs	-319.76	€34.57
			1800 HDDs	-266.29	€63.04
Enschede	HDD	call - put	1700 HDDs	-319.76	€826.65
			1800 HDDs	-266.29	€663.04
			1700 HDDs	-319.76	€25.47
			1800 HDDs	-266.29	€47.82
Nijmegen	HDD	call - put	1700 HDDs	-319.76	€760.29
			1800 HDDs	-266.29	€602.50
			1700 HDDs	-319.76	€30.00
			1800 HDDs	-266.29	€55.48

7.3 Pricing Dutch Cumulative Average Temperature Options

7.3.1 Monte Carlo Simulation

Some researchers such as Esunge and Njong (2020) and Wang et al. (2015) suggest using a Monte Carlo simulation to price temperature options. According to them, in some instances where the condition that $\mathbb{P}[max(0, 18 - DAT_n) = 0] \approx 0$ almost does not hold, using the Alaton approximation could result in inaccurate prices. As an alternative, one could use the Monte Carlo method. Following the work of Wang et al. (2015), the Monte Carlo simulation generates a set of paths and calculates the payoff for each path, the price of the option is then calculated as the average of the payoffs from each path. In the case of the Netherlands, after generating 50.000 DAT scenarios between the 1^{st} of November 2030 and the 31^{st} of March 2031, it was found that $\mathbb{P}[max(0, 18 - DAT_n) = 0] = 0$, therefore, using Equations 7.16 and 7.17 should, in principle, provide accurate prices for the winter periods.

Although Alexandridis and Zapranis (2012) provide close form pricing formulas for call and put options written on CAT indices, we follow a the mWang et al. (2015) in pricing HDD options, to price CAR indices using a Monte Carlo simulation. The payout of a CAT index is given by

$$CAT_n = \sum_{n=1}^{N} DAT_n \tag{7.21}$$

and the price of a CAT call option at time t is given by

$$C_{CAT}(t) = e^{-r(t_n - t)} \frac{1}{M} \sum_{i=1}^{M} \alpha \mathbb{E}^{\mathbb{Q}}[\max\{\sum_{n=1}^{N} CAT_n - K, 0\}] \tag{7.22}$$

where r is the discount interest rate, α is the principal nominal, m is the length of the time space, and M is the number of paths generated. The price of a CAT put is therefore

$$P_{CAT}(t) = e^{-r(t_n - t)} \frac{1}{M} \sum_{i=1}^{M} \alpha \mathbb{E}^{\mathbb{Q}}[\max\{K - CAT_m, 0\}] \tag{7.23}$$

Figure 7.3 shows the result of this approach where Option premium (e) is plotted for strikes (CAT) from 100 to 200. Nevertheless, in the absence of information about the market price of risk , it was assumed that $\lambda = 0$

Figure 7.3: Rotterdam CAT options

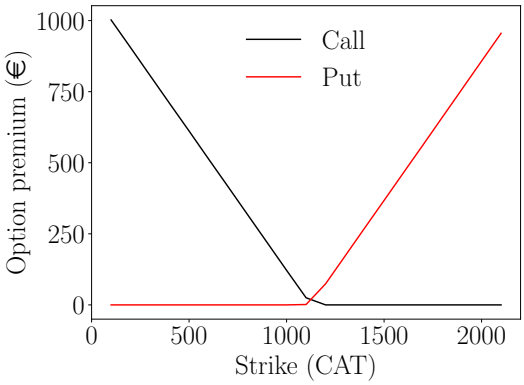

Spatial Basis Risk

J. C. Hull (2003) defines basis risk as the difference between the spot and future price. In the case of weather derivatives basis risk is defined differently. Spatial or geographical basis risk arises when there is a significant distance between the measurement site and the subject of the hedge. According to Rohrer (2004), this risk comprises also time basis risk, which results from the difference between the exposure period and the reference period of the derivative. However, since the latter can be overcome by trading weather derivatives with different maturities, this thesis will focus solely on spatial basis risk.

8.1 Weather Risk Management

Basis risk is often overlooked in weather derivative markets and receives limited attention in academic literature, typically with only brief mentions. However, some researchers such as Hess et al. (2002) have delved deeper into this aspect by examining the value of weather derivatives in developing countries, and emphasizing the significance of basis rusk in hedging crop-related risks.

In general, counterparties of hedged companies justify their offerings by relying on high correlation coefficients between the provided index station and the location being hedged.

Although Figure 8.1 demonstrates a strong correlation amond DATs, Rohrer (2004) explains that a high correlation coefficient does not eliminate the presence of basis risk. This means that even with cities located hundreds of kilometers apart, it is possible to have a high correlation (Rohrer, 2004). For instance, the coefficient correlation between Berlin and Munich on February was equal to 0.84, indicating that it is common for for Dutch cities to exhibit highly correlated temperatures. Furthermore, temperature disparities can sometimes be observed even over short distances due to the formation of micro-climates (Brockett et al., 2005; Manfredo & Richards, 2009).

Table 8.1: Correlation of Amsterdam's DAT with other cities

	Correlation coefficient
Rotterdam	0.999
Utrecht	0.999
Eindhoven	0.998
Groningen	0.999
Maastricht	0.998
Tilburg	0.999
Enschede	0.999
Nijmegen	0.998

8.2 Spatial Basis Risk Assessment

It was established in the last section that correlation might not be the best indicator of spatial basis risk. Nevertheless, the literature does not offer an accurate tool to deal with it and basis risk is accepted with no natural hedge, i.e. taking advantage from geographical diversification. Rohrer (2004) uses the distribution of the residuals of the correlation of the mean monthly temperatures of Berlin and Munich, they determine that there us a 10% probability of observing temperature in Berlin that deviates approximately 0.8 and 1.0 standard deviations below and above the expected value, respectively. Temperature derivatives are typically based on a strip or a month, hence analyzing monthly or seasonal basis risk can provide more insight. That is, observing Figure 8.1, it is seen that the mean DATs of the November-March strip of the 8 Dutch cities differ from Amsterdam's long run average during the same period. While a difference of 2°C can produce slight inaccuracies in the hedging of basis risk, Figure 8.2 indicates that difference can reach up to 7°C in the May-September strip which can lead to substantial financial losses. For instance, the mean Rotterdam DAT in the May-September strip 2022 was around 6.25 while Amsterdam's long run average during the same period is approximately 10.22, the difference of the index is about 60% which means that a hedging strategy should result in 60% less revenues.

Figure 8.1: Amsterdam's mean DAT vs other cities (Nov-Mar)

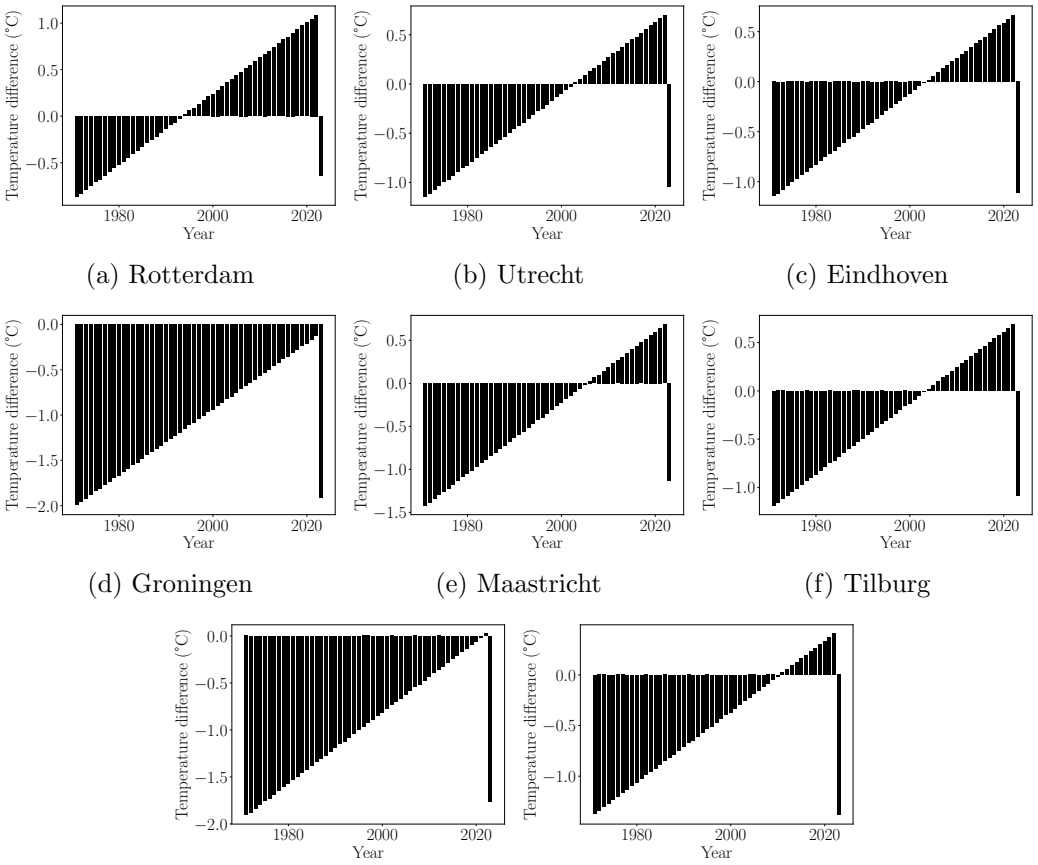

(a) Rotterdam (b) Utrecht (c) Eindhoven

(d) Groningen (e) Maastricht (f) Tilburg

(g) Enschede (h) Nijmegen

Figure 8.2: Amsterdam's mean DAT vs other cities (May-Sep strip)

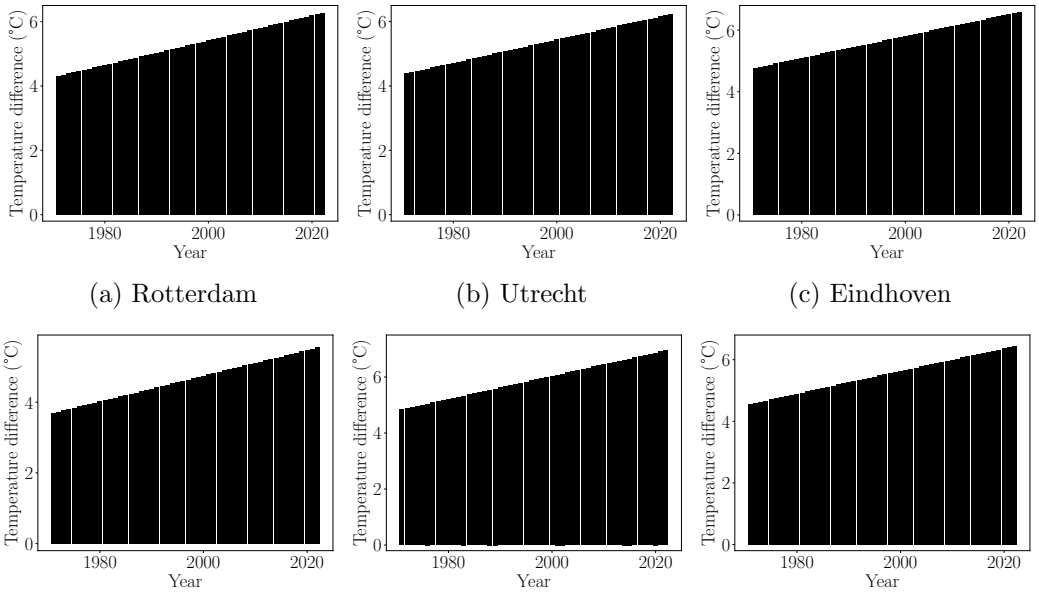

(a) Rotterdam (b) Utrecht (c) Eindhoven

(d) Groningen (e) Maastricht (f) Tilburg

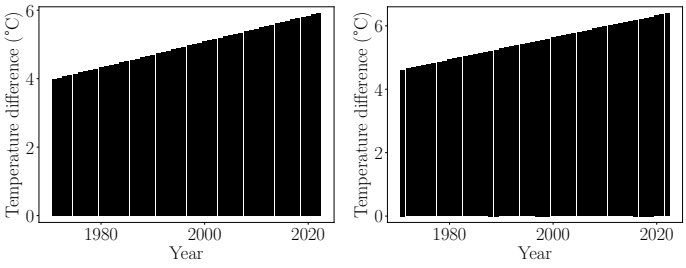

(g) Enschede (h) Nijmegen

8.3 Spatial Basis Risk Hedging

In the preceding section, we demonstrated that basis risk, which is reflected in the distance between the DAT in a reference station (Amsterdam - Schipol) and any other location in the Netherlands, can reach critical levels and induce serious financial losses. While Manfredo and Richards (2009) emphasize that the success of a hedge in the weather derivatives market relies of the behavior of the basis. When the basis is uncertain or unpredictable, it introduces challenges for hedging strategies. Moreover, basis risk also impacts the liquidity of the market due to limited market participants willing to take the risk. Unfortunately, spatial basis risk in weather derivatives remains a neglected topic. One potential solution to address this issue is the use of basis derivatives, which are based on the difference between the DATs of two stations (Alexandridis & Zapranis, 2012). For instance, suppose an energy importer wishes to hedge their profit in Rotterdam in either a winter or summer period, the investor can use the Amsterdam exchange traded derivative and the closest station to the location of the firm. That is, given ξ_n^A, Amsterdam's exchange traded temperature index (HDD or CAT), ξ_n^R, Rotterdam's temperature index, instead of trading a single temperature index, the importer can trade a bundle of indices given by

$$\xi_n = \alpha\xi_n^A + (1-\alpha)\xi_n^B \tag{8.1}$$

where ξ_n^B is a basis derivative written on the difference between ξ_n^A and ξ_n^R. However, the investor should be ready to bare the credit risk that derives from trading ξ_n^B OTC. Moreover, if the business is located far from a specific, the energy producer can hedge by using a combination of the nearby stations such that

$$\begin{aligned}\xi_n &= \alpha_0\xi_n^A + \alpha_1\xi_n^1 + \alpha_2\xi_n^2 + \cdots + \alpha_N\xi_n^N \\ &= \alpha_0\xi_n^A + \sum_{i=1}^{N}\alpha_i\xi_n^i\end{aligned} \tag{8.2}$$

where

$$0 \leq \alpha_i \leq 1, \quad \sum_{i=0}^{N}\alpha_i = 1$$

8.4 Basis Risk and Profits

Multiple authors have tried to account for temperature indices and basis risk in modeling profits from temperature-dependent businesses. Specifically, in the case of crop yields, when using weather derivatives as a hedging instrument, not only spatial basis should be considered, but also the relationship between crop yields and weather conditions. The effectiveness of the hedge depends on how closely the crop yields are linked to specific weather patterns. To account for this feature Turvey (2001) propose the following model

$$Y = AR^{\alpha}H^{\beta}$$

where Y is the crop yield, R is cumulative daily rainfall in the growing period, H is the growing degree days (GDDs)[1], and α and β are the output elasticities of rainfall and growing degree days. Other researchers such as Pardo et al. (2002) use HDDs and CDDs to model daily electricity loads in the Spanish electricity market. They demonstrate that weather and seasonality have a significant impact on electricity demand. That is, the model is given by

$$\begin{aligned} E_t =& c_1 + a_1 t + \beta_1 HDD_t + \gamma_1 CDD_t \\ &+ \sum_{i=2}^{7} \delta_{1i} W_{it} + \omega_1 H_t + \kappa_1 H_{t-1} \sum_{j=2}^{12} \lambda_{1j} M_{jt} + \varepsilon_{1t} \end{aligned}$$

where t is the time variable and the indices i and j represent the days of the week and the months of the year, excluding Monday and January, respectively. Therefore, the parameters W_{it} and M_{jt} are equal to 1 if the time variable belongs to i and j, respectively, and 0 otherwise. In addition, the parameters H_t and H_{t-1} are holiday dummies, where the latter is 1 if t is a holiday and the former is 1 if t comes after a holiday.

Consider the basic natural gas economy framework that Zweifel et al. (2017) develops, and suppose energy consumers have a demand function that reflects their knowledge about the future state of temperature represented by ξ_n such that the energy demand for a given season depends on an aggregate temperature index as given in equation 8.2

$$Q = a - p_r + \xi_n \tag{8.3}$$

where p_r is the retail price of gas sales and a is the marginal willingness to pay for the first unit of natural gas. Excluding distribution costs, the energy importer then

[1] A GDD is a weather-driven metric used in agriculture to evaluate the progress of crop growth.

has a profit of

$$\begin{aligned}\pi_i &= (p_r - p_i)Q \\ &= (p_r - p_i)(a - p_r + \xi_n)\end{aligned} \tag{8.4}$$

where p_i is the import price. Deriving equation 8.4 with respect to p_r, the optimal sales price is therefore

$$p_r^* = \frac{a + p_i + \xi_n}{2} \tag{8.5}$$

Using the fact that the quantity sold Q should be equal to the quantity imported Q_i, we get the optimal quantity imported as

$$Q_i^* = \frac{a + \xi_n - p_i}{2} \tag{8.6}$$

Hence the importer's maximum profit

$$\begin{aligned}\pi_i^* &= \left(\frac{a + p_i + \xi_n}{2} - p_i\right)\frac{a + \xi_n - p_i}{2} \\ &= \left(\frac{a + \xi_n - p_i}{2}\right)^2\end{aligned} \tag{8.7}$$

Now assume a profit maximizing natural gas producer that knows the importer's optimal demand Q_i^*, the gas producer then has a profit of

$$\begin{aligned}\pi_e &= (p_i - c(k))Q_i^* \\ &= (p_i - c(k))\frac{a + \xi_n - p_i}{2}\end{aligned} \tag{8.8}$$

where $c(k)$ is the unit cost of extracting and transporting gas which depend on capital stock k. Deriving equation 8.8 with respect to p_i, the optimal import price is therefore

$$p_i^* = \frac{a + \xi_n + c(k)}{2} \tag{8.9}$$

Using the fact that the quantity exported Q_e should be equal to the quantity imported Q_i, we get the optimal quantity exported as

$$Q_e^* = \frac{a + \xi_n - p_i*}{2} \tag{8.10}$$

$$= \frac{a + \xi_n - c(k)}{4} \tag{8.11}$$

Excluding the costs of capital, transport and distribution, the maximum profit of the producer and the importer are given by

$$\pi_e^* = \frac{1}{2}\left(\frac{a + \xi_n - c(k)}{2}\right)^2 \tag{8.12}$$

$$\pi_i^* = \frac{1}{4}\left(\frac{a + \xi_n - p_i}{2}\right)^2 \tag{8.13}$$

from equations 8.12 and 8.13 we can say that the profits of both the energy gas importer will depend on the temperature index of the given season.

Conclusion

Over the past decade, a new class of financial derivatives known as weather derivatives has emerged to allow energy and utility companies to effectively mitigate volumetric risk associated with energy sales, rather than focusing on price risk. Although the weather derivatives market, namely the temperature based options and futures market, is one of the fastest developing markets. Nonetheless, aside from a prevailing lack of awareness among investors about the advantages the market can provide, prospective investors exhibit a significant aversion to actively trade in this market. The main reason behind this is the incomplete nature of the market, which constitutes a challenge to accurately price the traded financial products and model the underlying temperature variables. In an attempt to contribute to the existing literature and the development of this market, at least regionally, this thesis provides a detailed and rigorous study of the modeling of the Dutch temperature market.

That is, the thesis carefully applies various data and signal preprocessig techniques to the DATs of nine Dutch cities and implements several approaches for identifying and and modeling trends in the data where it was found that DATs exhibit a significant seasonality. After de-noising the DATs time series, it was revealed that the data exhibits uniform peaks which allowed extracting a seasonal mean variation using a first order truncated Fourier series. This thesis allocates a substantial part to the daily modeling of temperature to capture the statistical properties and the underlying dynamics of the DATs. To do so, we use Ito calculus and specifically the Ito-Doeblin formula to justify the choice of the mean reverting Ornstein-Uhlenbeck process which we fit to the detrended and deseasonalized DATs time series and for which we estimate the parameters, such as the speed of mean reversion and the volatility, accordingly. However, since there is an absence of consensus on how to model the time temperature varying volatility in the OU process, the thesis implements different approaches to model the volatility of DATs. It was concluded that using Fourier series yields the best results from a parsimonious perspective. Although the temperature derivatives market is an incomplete and illiquid market, some authors explain that risk neutral pricing can still be performed. Hence, in chapter 8 we price temperature options under the risk-neutral measure $\mathbb{Q}$ by discounting the expected payoff of the derivative using two pricing approaches. First, assuming a normal distribution, HDD options are priced using

a closed form solution of the SDE that describes the DAT dynamics. Using this approach raised underlined the importance of the market price of risk. So far, the literature assumed that the market price of risk is equal to zero. Although some researchers attempted to successfully estimate this parameter, they have done so under unfounded premises such as the assumption that the market price of risk is constant. In this thesis, the market price of risk is approximated using a root finding algorithm which has revealed that investors might have a risk seeking attitude. Nevertheless, since it was also shown that the market price of risk is not constant in time and that it also varies with respect to the strike and market price of the underlying derivative, such behavioral inferences might not hold across the whole market. Particularly, investors avoid trading these instruments because of the presence of spatial basis risk, which was shown to reach critical levels even in small countries such as the Netherlands. Unfortunately, basis risk is not sufficiently studied in the literature and it is often approached through a brief correlation analysis which has proved to lead to erroneous conclusions.

Additionally, as many authors have mentioned, the reference temperature of 18°C was mainly developed for the US markets, but is used in European markets as well. Often, temperatures across Europe are very close to 18°C which hinders accurate pricing using closed form pricing formulas. Therefore, it can be interesting to calibrate the reference temperature to local weather conditions, which should allow more flexibility with pricing.

Bibliography

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

Alexandridis, A., & Zapranis, A. D. (2012). *Weather derivatives: Modeling and pricing weather-related risk*. Springer Science & Business Media.

Bandara, K., Hyndman, R. J., & Bergmeir, C. (2021). Mstl: A seasonal-trend decomposition algorithm for time series with multiple seasonal patterns. *arXiv preprint arXiv:2107.13462*.

Bellini, F. (2005). *The weather derivatives market: Modelling and pricing temperature* (Doctoral dissertation). Università della Svizzera italiana.

Benth, F. E., & Benth, J. š. (2007). The volatility of temperature and pricing of weather derivatives. *Quantitative Finance*, *7*(5), 553–561.

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

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

Bhowan, A. (2003). *Temperature derivatives*. School of Computational and Applied Mathematics, University of Wiwatersrand.

Bibby, B. M., & Sørensen, M. (1995). Martingale estimation functions for discretely observed diffusion processes. *Bernoulli*, 17–39.

Bob, D. (1998a). At least: A model for weather risk. *Weather risk special report, Energy and Power Risk Management*, 30–32.

Bob, D. (1998b). Black-scholes won't do. *Weather risk special report, Energy and Power Risk Management*, 8–9.

Box, G. E., Jenkins, G. M., Reinsel, G. C., & Ljung, G. M. (2015). *Time series analysis: Forecasting and control*. John Wiley & Sons.

BP. (2021). *Statistical review of world energy* (Rep. No. 70). British Petroleum.

Brockett, P. L., Cooper, W. W., Golden, L. L., & Pitaktong, U. (1994). A neural network method for obtaining an early warning of insurer insolvency. *Journal of Risk and Insurance*, 402–424.

Brockett, P. L., Wang, M., & Yang, C. (2005). Weather derivatives and weather risk management. *Risk Management and Insurance Review*, *8*(1), 127–140.

Brockwell, P. J., & Davis, R. A. (2009). *Time series: Theory and methods*. Springer science & business media.

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

Buckley, N., Hamilton, A., Harding, J., Roche, N., Ross, N., Sands, E., Skelding, R., Watford, N., & Whitlow, H. (2002). European weather derivatives. *General Insurance Convention*.

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

Cao, M., Li, A., & Wei, J. (2004). Watching the weather report. *Canadian Investment Review*, *17*(2), 27–33.

Cao, M., Li, A., & Wei, J. Z. (2003). Weather derivatives: A new class of financial instruments. *Available at SSRN 1016123*.

Cao, M., Wei, J., et al. (1999). *Pricing weather derivative: An equilibrium approach*. SSRN.

Cao, M., & Wei, J. (2000). Pricing the weather. *Risk*, 67–70.

Carmona, R. (1999). Calibrating degree day options. *3rd seminar on stochastic analysis, random field and applications. Ecole Polytechnique de Lausanne, Switzerland*.

Carmona, R. (2014). *Statistical analysis of financial data in r* (Vol. 2). Springer.

Cayrade, P. (2004). Investments in gas pipelines and liquefied natural gas infrastructure. what is the impact on the security of supply? *What is the Impact on the Security of Supply*.

Cleveland, R. B., Cleveland, W. S., McRae, J. E., & Terpenning, I. (1990). Stl: A seasonal-trend decomposition. *J. Off. Stat*, *6*(1), 3–73.

Cleveland, W. S. (1979). Robust locally weighted regression and smoothing scatterplots. *Journal of the American statistical association*, *74*(368), 829–836.

Cleveland, W. S., & Devlin, S. J. (1988). Locally weighted regression: An approach to regression analysis by local fitting. *Journal of the American statistical association*, *83*(403), 596–610.

Council of the EU & the European Council. (2023a). Energy prices and security of supply. https://www.consilium.europa.eu/en/policies/energy-prices-and-security-of-supply/

Council of the EU & the European Council. (2023b). Infographic - energy price rise since 2021. https://www.consilium.europa.eu/en/infographics/energy-prices-2021/

Cummins, J. D., & Mahul, O. (2003). Optimal insurance with divergent beliefs about insurer total default risk. *Journal of Risk and Uncertainty*, *27*, 121–138.

D'arcy, S. P., & France, V. G. (1992). Catastrophe futures: A better hedge for insurers. *Journal of Risk and Insurance*, 575–600.

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

De Cian, E., Lanzi, E., & Roson, R. (2013). Seasonal temperature variations and energy demand: A panel cointegration analysis for climate change impact assessment. *Climatic Change*, *116*, 805–825.

Dorfleitner, G., & Wimmer, M. (2010). The pricing of temperature futures at the chicago mercantile exchange. *Journal of Banking & Finance*, *34*(6), 1360–1370.

Dornier, Q., & Querel, M. (2000). Caution to the wind. energy and power risk managment. *Weather risk special report*, *13*(8), 30–32.

Dowling, P. (2013). The impact of climate change on the european energy system. *Energy Policy*, *60*, 406–417.

Engle, R. F., Granger, C. W., Rice, J., & Weiss, A. (1986). Semiparametric estimates of the relation between weather and electricity sales. *Journal of the American statistical Association*, *81*(394), 310–320.

Engle, R. F., Mustafa, C., & Rice, J. (1992). Modelling peak electricity demand. *Journal of forecasting*, *11*(3), 241–251.

Erdmann, G., & Zweifel, P. (2008). *Energieökonomik: Theorie und anwendungen*. Springer.

Esunge, J., & Njong, J. J. (2020). Weather derivatives and the market price of risk. *Journal of Stochastic Analysis*, *1*(3), 7.

European Commission. (2023). Commission proposes reform of the eu electricity market design to boost renewables, better protect consumers and enhance industrial competitiveness. https://ec.europa.eu/commission/presscorner/detail/en/IP_23_1591

Eurostat, S. E. (2022). Energy production and imports.

Field, C. B., Barros, V., Stocker, T. F., & Dahe, Q. (2012). *Managing the risks of extreme events and disasters to advance climate change adaptation: Special report of the intergovernmental panel on climate change*. Cambridge University Press.

Gabbi, G., Zanotti, G., et al. (2006). Climate variables and weather derivatives. In *Climate variables and weather derivatives*.

Geman, H., & Leonardi, M.-P. (2005). Alternative approaches to weather derivatives pricing. *Managerial Finance*, *31*(6), 46–72.

Golden, L. L., Wang, M., & Yang, C. (2007). Handling weather related risks through the financial markets: Considerations of credit risk, basis risk, and hedging. *Journal of Risk and Insurance*, *74*(2), 319–346.

Hamlet, A. F., Lee, S.-Y., Mickelson, K. E., & Elsner, M. M. (2010). Effects of projected climate change on energy supply and demand in the pacific northwest and washington state. *Climatic Change*, *102*(1-2), 103–128.

Härdle, W. K., & Cabrera, B. L. (2012). The implied market price of weather risk. *Applied Mathematical Finance*, *19*(1), 59–95.

Hastie, T., Tibshirani, R., Friedman, J. H., & Friedman, J. H. (2009). *The elements of statistical learning: Data mining, inference, and prediction* (Vol. 2). Springer.

Henley, A., & Peirson, J. (1998). Residential energy demand and the interaction of price and temperature: British experimental evidence. *Energy Economics*, *20*(2), 157–171.

Hess, U., Richter, K., & Stoppa, A. (2002). Weather risk management for agriculture and agri-business in developing countries. *Climate Risk and the Weather Market, Financial Risk Management with Weather Hedges. London: Risk Books*.

Holland, M., Amann, M., Heyes, C., Rafaj, P., Schöpp, W., Hunt, A., & Watkiss, P. (2011). The reduction in air quality impacts and associated economic benefits of mitigation policy: Summary of results from the ec rtd climatecost project.

Hull, J., & White, A. (1990). Pricing interest-rate-derivative securities. *The review of financial studies*, *3*(4), 573–592.

Hull, J. C. (2003). *Options futures and other derivatives*. Pearson.

Hyndman, R. J., & Athanasopoulos, G. (2018). *Forecasting: Principles and practice*. OTexts.

International Energy Agency. (2022a). *Netherlands climate resilience policy indicator* [Licence: Creative Commons Attribution CC BY 4.0].

International Energy Agency. (2022b). *Never too early to prepare for next winter: Europe's gas balance for 2023-2024* [Licence: Creative Commons Attribution CC BY 4.0].

Isaac, M., & Van Vuuren, D. P. (2009). Modeling global residential sector energy demand for heating and air conditioning in the context of climate change. *Energy policy*, *37*(2), 507–521.

Jafari, M., & Abbasian, S. (2017). The moments for solution of the cox-ingersoll-ross interest rate model. *Journal of Finance and Economics*, *5*(1), 34–37.

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

Jewson, S., & Caballero, R. (2003). Seasonality in the statistics of surface air temperature and the pricing of weather derivatives. *Meteorological applications*, *10*(4), 367–376.

Levenberg, K. (1944). A method for the solution of certain non-linear problems in least squares. *Quarterly of applied mathematics*, *2*(2), 164–168.

Li, X., & Sailor, D. (1995). Electricity use sensitivity to climate and climate change. *World Resource Review*, *7*(3).

Major, J. (1999). Index hedge performance: Insurer market penetration and basis risk. In *The financing of catastrophe risk* (pp. 391–432). University of Chicago Press.

Manfredo, M. R., & Richards, T. J. (2009). Hedging with weather derivatives: A role for options in reducing basis risk. *Applied Financial Economics*, *19*(2), 87–97.

Manolakis, D. G., & Ingle, V. K. (2011). *Applied digital signal processing: Theory and practice*. Cambridge university press.

Marquardt, D. W. (1963). An algorithm for least-squares estimation of nonlinear parameters. *Journal of the society for Industrial and Applied Mathematics*, *11*(2), 431–441.

Martin, R., & Peter C, K. (2014). Symfit. *J.P.B.* https://doi.org/10.5281/zenodo.1133336

Moreno, M. (2000). Riding the temp. *Weather Derivatives, FOW Special Supplement*.

Mraoua, M. (2007). Temperature stochastic modeling and weather derivatives pricing: Empirical study with moroccan data. *Afrika Statistika*, *2*(1).

Oosterlee, C. W., & Grzelak, L. A. (2019). *Mathematical modeling and computation in finance: With exercises and python and matlab computer codes*. World Scientific.

Pardo, A., Meneu, V., & Valor, E. (2002). Temperature and seasonality influences on spanish electricity load. *Energy Economics*, *24*(1), 55–70.

Peirson, J., & Henley, A. (1994). Electricity load and temperature: Issues in dynamic specification. *Energy Economics*, *16*(4), 235–243.

Pirrong, C., & Jermakyan, M. (2008). The price of power: The valuation of power and weather derivatives. *Journal of Banking & Finance*, *32*(12), 2520–2529.

Praktiknjo, A. J. (2014). Stated preferences based estimation of power interruption costs in private households: An example from germany. *Energy*, *76*, 82–90.

Ranson, M., Morris, L., & Kats-Rubin, A. (2014). Climate change and space heating energy demand: A review of the literature.

Ritter, M., Musshoff, O., & Odening, M. (2014). Minimizing geographical basis risk of weather derivatives using a multi-site rainfall model. *Computational Economics*, *44*, 67–86.

Rohrer, M. (2004). The relevance of basis risk in the weather derivatives market. *WIT Transactions on Modelling and Simulation*, *38*.

Sailor, D. J., & Muñoz, J. R. (1997). Sensitivity of electricity and natural gas consumption to climate in the usa—methodology and results for eight states. *Energy*, *22*(10), 987–998.

Shreve, S. E., et al. (2004). *Stochastic calculus for finance ii: Continuous-time models* (Vol. 11). Springer.

Stoft, S. (2002). *Power system economics: Designing markets for electricity* (Vol. 468). IEEE press Piscataway.

Taylor, S. J., & Letham, B. (2018). Forecasting at scale. *The American Statistician*, *72*(1), 37–45.

Tindall, J. (2006). Weather derivatives: Pricing and risk management applications. *Institute of Actuaries of Australia*.

Turvey, C. G. (2001). Weather derivatives for specific event risks in agriculture. *Applied Economic Perspectives and Policy*, *23*(2), 333–351.

Ukkola, A., Roderick, M., Barker, A., & Pitman, A. (2019). Exploring the stationarity of australian temperature, precipitation and pan evaporation records over the last century. *Environmental Research Letters*, *14*(12), 124035.

Valor, E., Meneu, V., & Caselles, V. (2001). Daily air temperature and electricity load in spain. *Journal of applied Meteorology*, *40*(8), 1413–1421.

Wang, Z., Li, P., Li, L., Huang, C., & Liu, M. (2015). Modeling and forecasting average temperature for weather derivative pricing. *Advances in Meteorology*, *2015*.

Yang, C. C., Brockett, P. L., & Wen, M.-M. (2009). Basis risk and hedging efficiency of weather derivatives. *The Journal of Risk Finance*, *10*(5), 517–536.

Zanotti, G., Gabbi, G., & Laboratore, D. (2003). Climate variables and weather derivatives: Gas demand, temperature and seasonality effects in the italian case. *Temperature and Seasonality Effects in the Italian Case (14/01/2003)*.

Zapranis, A., & Alexandridis, A. (2008). Modelling the temperature time-dependent speed of mean reversion in the context of weather derivatives pricing. *Applied Mathematical Finance*, *15*(4), 355–386.

Zweifel, P., Praktiknjo, A., & Erdmann, G. (2017). *Energy economics: Theory and applications*. Springer.

Tables

Table A.1: Properties of gas energy sources

		Density (kg/m^3)	**Upper heating value** (Mj/m^3)	**Lower heating value** (Mj/m^3)
Methane	CH_4	0.7175	39.819	35.883
Ethane	C_2H_6	1.3550	70.293	64.345
Propane	C_3H_8	2.0110	101.242	93.215
Butane	C_{410}	2.7080	134.061	123.810
Hydrogen	H_2	0.08988	12.745	10.783
Carbon monoxide	CO	1.25050	12.633	12.633
Nitrogen	N_2	1.2504		
Oxygen	O_2		1.4290	
Carbon dioxide	CO_2	1.9770		
Air		1.2930		
Natural gas H		0.79	∼41	∼37
Natural gas L		0.83	∼35	∼32
Biogas		1.12	∼27	∼24

Source: Zweifel et al. (2017).

Table A.2: Descriptive Statistics

	Amsterdam			Rotterdam			Utrecht		
	$\mathbf{T^{max}}$	$\mathbf{T^{min}}$	**DAT**	$\mathbf{T^{max}}$	$\mathrm{T^{min}}$	**DAT**	$\mathbf{T^{max}}$	$\mathrm{T^{min}}$	**DAT**
Count					19034				
Mean	13.93	6.51	10.22	14.07	6.43	10.25	14.19	6.01	10.10
Std	6.96	5.60	6.10	6.92	5.70	6.09	7.26	5.73	6.26
Min	−9.50	−18.80	−12.00	−9.70	−17.10	−12.60	−10.60	−18.90	−13.20
Max	36.40	21.40	28.60	37.20	20.90	28.20	37.50	22.40	29.80
Skewness	0.01	−0.29	−0.15	0.03	−0.31	−0.16	0.02	−0.31	−0.16
Kurtosis	−0.44	−0.20	−0.39	−0.41	−0.25	−0.40	−0.46	−0.19	−0.38
JB-Stat	153.74	303.58	193.18	137.62	363.16	207.14	171.41	325.69	193.99
***p*-value**	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	Eindhoven			**Groningen**			**Maastricht**		
	$\mathbf{T^{max}}$	$\mathbf{T^{min}}$	**DAT**	$\mathbf{T^{max}}$	$\mathbf{T^{min}}$	**DAT**	$\mathbf{T^{max}}$	$\mathbf{T^{min}}$	**DAT**
Mean	14.62	5.91	10.27	13.49	5.17	9.33	14.29	6.33	10.31
Std	7.61	5.84	6.49	7.45	5.76	6.36	7.73	5.97	6.67
Min	−9.80	−19.70	−13.50	−11.60	−22.00	−16.40	−12.30	−19.30	−14.90
Max	40.40	23.60	30.20	36.90	21.50	28.20	39.60	23.20	30.20
Skewness	0.03	−0.29	−0.13	0.03	−0.34	−0.16	0.03	−0.29	−0.13
Kurtosis	−0.52	−0.21	−0.42	−0.47	0.00	−0.32	−0.53	−0.24	−0.44
JB-Stat	213.63	298.44	191.72	177.01	360.10	163.16	221.7	314.85	206.36
***p*-value**	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	Nijmegen			**Tilburg**			**Enschede**		
	$\mathbf{T^{max}}$	$\mathbf{T^{min}}$	**DAT**	$\mathbf{T^{max}}$	$\mathbf{T^{min}}$	**DAT**	$\mathbf{T^{max}}$	$\mathbf{T^{min}}$	**DAT**
Mean	14.48	5.64	10.06	14.53	5.83	10.18	13.87	5.25	9.56
Std	7.67	5.83	6.51	7.46	5.86	6.41	7.72	5.92	6.55
Min	−10.30	−20.10	−14.00	−10.60	−18.30	−13.20	−13.00	−21.80	−15.60
Max	40.10	22.60	31.40	40.70	22.10	29.80	40.20	23.10	29.60
Skewness	0.03	−0.32	−0.15	0.04	−0.35	−0.16	0.02	−0.34	−0.16
Kurtosis	−0.49	−0.20	−0.40	−0.45	−0.15	−0.37	−0.51	−0.05	−0.34
JB-Stat	194.67	362.6	199.70	168.41	409.42	190.80	209.81	365.10	170.59
***p*-value**	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

JB-Stat: Jarque-Bera statistic

Table A.3: Autoregressive model results

City:		Amsterdam	**No. Observations:**			19034
Dep. Variable:		DAT	**Log Likelihood**			-39648.734
Model:		AR(17)	**AIC**			79335.468
Method:		Conditional MLE	**BIC**			79484.677
	Coef	**Std err**	**z**	$p > \|z\|$	**[0.025**	**0.975]**
Intercept	0.2619	0.030	8.812	0.000	0.204	0.320
L1	0.9426	0.007	130.035	0.000	0.928	0.957
L2	−0.1962	0.010	−19.695	0.000	−0.216	−0.177
L3	0.0928	0.010	9.217	0.000	0.073	0.112
L4	0.0079	0.010	0.782	0.434	−0.012	0.028
L5	0.0136	0.010	1.347	0.178	−0.006	0.033
L6	0.0144	0.010	1.431	0.152	−0.005	0.034
L7	0.0192	0.010	1.906	0.057	−0.001	0.039
L8	−0.0008	0.010	−0.081	0.935	−0.021	0.019
L9	0.0089	0.010	0.882	0.378	−0.011	0.029
L10	0.0139	0.010	1.376	0.169	−0.006	0.034
L11	0.0087	0.010	0.858	0.391	−0.011	0.028
L12	−0.0012	0.010	−0.121	0.904	−0.021	0.019
L13	0.0046	0.010	0.458	0.647	−0.015	0.024
L14	0.0214	0.010	2.122	0.034	0.002	0.041
L15	0.0053	0.010	0.527	0.598	−0.014	0.025
L16	−0.0089	0.010	−0.891	0.373	−0.028	0.011
L17	0.0283	0.007	3.914	0.000	0.014	0.043
City:		Rotterdam	**No. Observations:**			19034
Dep. Variable:		DAT	**Log Likelihood**			-40355.715
Model:		AR(17)	**AIC**			80749.429
Method:		Conditional MLE	**BIC**			80898.638
	Coef	**Std err**	**z**	$p > \|z\|$	**[0.025**	**0.975]**
Intercept	0.2824	0.031	9.060	0.000	0.221	0.343
L1	0.9412	0.007	129.870	0.000	0.927	0.955
L2	−0.2039	0.010	−20.482	0.000	−0.223	−0.184
L3	0.0999	0.010	9.927	0.000	0.080	0.120
L4	-8.572×10^{-5}	0.010	−0.008	0.993	−0.020	0.020
L5	0.0283	0.010	2.804	0.005	0.009	0.048
L6	0.0022	0.010	0.216	0.829	−0.018	0.022
L7	0.0189	0.010	1.872	0.061	−0.001	0.039
L8	0.0069	0.010	0.687	0.492	−0.013	0.027
L9	0.0045	0.010	0.445	0.656	−0.015	0.024
L10	0.0124	0.010	1.229	0.219	−0.007	0.032
L11	0.0076	0.010	0.755	0.450	−0.012	0.027
L12	0.0029	0.010	0.292	0.771	−0.017	0.023
L13	0.0006	0.010	0.062	0.951	−0.019	0.020
L14	0.0289	0.010	2.863	0.004	0.009	0.049
L15	0.0046	0.010	0.454	0.650	−0.015	0.024
L16	−0.0171	0.010	−1.719	0.086	−0.037	0.002
L17	0.0347	0.007	4.786	0.000	0.020	0.049

City:	Utrecht	**No. Observations:**	19034
Dep. Variable:	DAT	**Log Likelihood**	-40821.287
Model:	AR(17)	**AIC**	81680.574
Method:	Conditional MLE	**BIC**	81829.783

	Coef	**Std err**	**z**	$p > \|z\|$	**[0.025**	**0.975]**
Intercept	0.2815	0.031	9.107	0.000	0.221	0.342
L1	0.9522	0.007	131.369	0.000	0.938	0.966
L2	−0.2111	0.010	−21.084	0.000	−0.231	−0.191
L3	0.0936	0.010	9.242	0.000	0.074	0.113
L4	0.0103	0.010	1.015	0.310	−0.010	0.030
L5	0.0164	0.010	1.618	0.106	−0.003	0.036
L6	0.0067	0.010	0.662	0.508	−0.013	0.027
L7	0.0197	0.010	1.943	0.052	0.000	0.040
L8	0.0053	0.010	0.525	0.600	−0.015	0.025
L9	0.0047	0.010	0.464	0.643	−0.015	0.025
L10	0.0147	0.010	1.450	0.147	−0.005	0.035
L11	0.0071	0.010	0.704	0.481	−0.013	0.027
L12	0.0106	0.010	1.041	0.298	−0.009	0.030
L13	−0.0083	0.010	−0.819	0.413	−0.028	0.012
L14	0.0287	0.010	2.824	0.005	0.009	0.049
L15	0.0039	0.010	0.389	0.697	−0.016	0.024
L16	−0.0116	0.010	−1.163	0.245	−0.031	0.008
L17	0.0292	0.007	4.027	0.000	0.015	0.043

City:	Eindhoven	**No. Observations:**	19034
Dep. Variable:	DAT	**Log Likelihood**	-41881.556
Model:	AR(17)	**AIC**	83801.112
Method:	Conditional MLE	**BIC**	83950.321

	Coef	**Std err**	**z**	$p > \|z\|$	**[0.025**	**0.975]**
Intercept	0.2962	0.032	9.161	0.000	0.233	0.360
L1	0.9546	0.007	131.725	0.000	0.940	0.969
L2	−0.2189	0.010	−21.846	0.000	−0.239	−0.199
L3	0.1025	0.010	10.100	0.000	0.083	0.122
L4	3.549×10^{-5}	0.010	0.003	0.997	−0.020	0.020
L5	0.0189	0.010	1.861	0.063	−0.001	0.039
L6	0.0150	0.010	1.478	0.139	−0.005	0.035
L7	0.0085	0.010	0.832	0.406	−0.011	0.028
L8	0.0082	0.010	0.809	0.418	−0.012	0.028
L9	0.0129	0.010	1.265	0.206	−0.007	0.033
L10	0.0089	0.010	0.878	0.380	−0.011	0.029
L11	0.0108	0.010	1.063	0.288	−0.009	0.031
L12	0.0039	0.010	0.388	0.698	−0.016	0.024
L13	−0.0095	0.010	−0.937	0.349	−0.029	0.010
L14	0.0304	0.010	2.989	0.003	0.010	0.050
L15	0.0054	0.010	0.528	0.598	−0.015	0.025
L16	−0.0168	0.010	−1.680	0.093	−0.036	0.003
L17	0.0364	0.007	5.026	0.000	0.022	0.051

City:		Groningen	**No. Observations:**			19034
Dep. Variable:		DAT	**Log Likelihood**			-41427.248
Model:		AR(17)	**AIC**			82892.496
Method:		Conditional MLE	**BIC**			83041.705
	Coef	**Std err**	**z**	$p > \|z\|$	**[0.025**	**0.975]**
Intercept	0.2612	0.030	8.790	0.000	0.203	0.319
L1	0.9253	0.007	127.660	0.000	0.911	0.940
L2	−0.1870	0.010	−18.932	0.000	−0.206	−0.168
L3	0.0886	0.010	8.892	0.000	0.069	0.108
L4	0.0047	0.010	0.467	0.641	−0.015	0.024
L5	0.0335	0.010	3.355	0.001	0.014	0.053
L6	0.0019	0.010	0.188	0.851	−0.018	0.021
L7	0.0158	0.010	1.583	0.113	−0.004	0.035
L8	0.0037	0.010	0.373	0.709	−0.016	0.023
L9	0.0126	0.010	1.259	0.208	−0.007	0.032
L10	0.0150	0.010	1.504	0.133	−0.005	0.035
L11	0.0046	0.010	0.457	0.648	−0.015	0.024
L12	0.0047	0.010	0.470	0.638	−0.015	0.024
L13	0.0068	0.010	0.680	0.496	−0.013	0.026
L14	0.0124	0.010	1.243	0.214	−0.007	0.032
L15	0.0101	0.010	1.018	0.309	−0.009	0.030
L16	−0.0113	0.010	−1.143	0.253	−0.031	0.008
L17	0.0306	0.007	4.226	0.000	0.016	0.045

City:		Maastricht	**No. Observations:**			19034
Dep. Variable:		DAT	**Log Likelihood**			-41286.890
Model:		AR(17)	**AIC**			82611.779
Method:		Conditional MLE	**BIC**			82760.988
	Coef	**Std err**	**z**	$p > \|z\|$	**[0.025**	**0.975]**
Intercept	0.2778	0.031	9.047	0.000	0.218	0.338
L1	1.0011	0.007	138.124	0.000	0.987	1.015
L2	−0.2502	0.010	−24.391	0.000	−0.270	−0.230
L3	0.0968	0.010	9.289	0.000	0.076	0.117
L4	0.0052	0.010	0.500	0.617	−0.015	0.026
L5	0.0107	0.010	1.023	0.307	−0.010	0.031
L6	0.0147	0.010	1.410	0.159	−0.006	0.035
L7	0.0177	0.010	1.699	0.089	−0.003	0.038
L8	−0.0028	0.010	−0.266	0.791	−0.023	0.018
L9	0.0138	0.010	1.321	0.186	−0.007	0.034
L10	0.0147	0.010	1.408	0.159	−0.006	0.035
L11	−0.0033	0.010	−0.317	0.751	−0.024	0.017
L12	0.0162	0.010	1.548	0.122	−0.004	0.037
L13	−0.0164	0.010	−1.575	0.115	−0.037	0.004
L14	0.0356	0.010	3.408	0.001	0.015	0.056
L15	−0.0073	0.010	−0.696	0.486	−0.028	0.013
L16	−0.0067	0.010	−0.649	0.517	−0.027	0.013
L17	0.0333	0.007	4.596	0.000	0.019	0.047

City:		Tilburg	**No. Observations:**			19034
Dep. Variable:		DAT	**Log Likelihood**			-41665.514
Model:		AR(17)	**AIC**			83369.029
Method:		Conditional MLE	**BIC**			83518.237
	Coef	**Std err**	**z**	$p > \|z\|$	**[0.025**	**0.975]**
Intercept	0.2932	0.032	9.155	0.000	0.230	0.356
L1	0.9506	0.007	131.184	0.000	0.936	0.965
L2	−0.2127	0.010	−21.266	0.000	−0.232	−0.193
L3	0.0936	0.010	9.248	0.000	0.074	0.113
L4	0.0090	0.010	0.885	0.376	−0.011	0.029
L5	0.0188	0.010	1.853	0.064	−0.001	0.039
L6	0.0102	0.010	1.009	0.313	−0.010	0.030
L7	0.0109	0.010	1.074	0.283	−0.009	0.031
L8	0.0092	0.010	0.904	0.366	−0.011	0.029
L9	0.0115	0.010	1.136	0.256	−0.008	0.031
L10	0.0083	0.010	0.821	0.412	−0.012	0.028
L11	0.0127	0.010	1.257	0.209	−0.007	0.033
L12	0.0019	0.010	0.191	0.849	−0.018	0.022
L13	−0.0040	0.010	−0.391	0.696	−0.024	0.016
L14	0.0279	0.010	2.750	0.006	0.008	0.048
L15	0.0039	0.010	0.382	0.702	−0.016	0.024
L16	−0.0190	0.010	−1.900	0.057	−0.039	0.001
L17	0.0384	0.007	5.300	0.000	0.024	0.053

City:		Enschede	**No. Observations:**			19034
Dep. Variable:		DAT	**Log Likelihood**			-42445.384
Model:		AR(17)	**AIC**			84928.768
Method:		Conditional MLE	**BIC**			85077.977
	Coef	**Std err**	**z**	$p > \|z\|$	**[0.025**	**0.975]**
Intercept	0.2852	0.031	9.070	0.000	0.224	0.347
L1	0.9462	0.007	130.568	0.000	0.932	0.960
L2	−0.2108	0.010	−21.124	0.000	−0.230	−0.191
L3	0.0907	0.010	8.989	0.000	0.071	0.111
L4	0.0164	0.010	1.618	0.106	−0.003	0.036
L5	0.0133	0.010	1.313	0.189	−0.007	0.033
L6	0.0048	0.010	0.471	0.637	−0.015	0.025
L7	0.0273	0.010	2.695	0.007	0.007	0.047
L8	−0.0072	0.010	−0.716	0.474	−0.027	0.013
L9	0.0158	0.010	1.564	0.118	−0.004	0.036
L10	0.0168	0.010	1.657	0.098	−0.003	0.037
L11	−0.0002	0.010	−0.019	0.985	−0.020	0.020
L12	0.0093	0.010	0.919	0.358	−0.011	0.029
L13	0.0003	0.010	0.025	0.980	−0.020	0.020
L14	0.0109	0.010	1.075	0.282	−0.009	0.031
L15	0.0227	0.010	2.249	0.025	0.003	0.042
L16	−0.0220	0.010	−2.206	0.027	−0.042	−0.002
L17	0.0361	0.007	4.986	0.000	0.022	0.050

City:	Nijmegen	**No. Observations:**	19034
Dep. Variable:	DAT	**Log Likelihood**	-41783.149
Model:	AR(17)	**AIC**	83604.299
Method:	Conditional MLE	**BIC**	83753.507

	Coef	**Std err**	**z**	$p > \|\mathbf{z}\|$	**[0.025**	**0.975]**
Intercept	0.2842	0.032	9.016	0.000	0.222	0.346
L1	0.9443	0.007	130.299	0.000	0.930	0.958
L2	−0.2014	0.010	−20.208	0.000	−0.221	−0.182
L3	0.0869	0.010	8.622	0.000	0.067	0.107
L4	0.0077	0.010	0.763	0.445	−0.012	0.027
L5	0.0245	0.010	2.426	0.015	0.005	0.044
L6	0.0040	0.010	0.398	0.691	−0.016	0.024
L7	0.0197	0.010	1.950	0.051	0.000	0.039
L8	0.0022	0.010	0.217	0.828	−0.018	0.022
L9	0.0160	0.010	1.586	0.113	−0.004	0.036
L10	0.0093	0.010	0.920	0.358	−0.011	0.029
L11	0.0101	0.010	0.997	0.319	−0.010	0.030
L12	0.0059	0.010	0.581	0.561	−0.014	0.026
L13	−0.0087	0.010	−0.866	0.387	−0.029	0.011
L14	0.0252	0.010	2.498	0.012	0.005	0.045
L15	0.0101	0.010	1.004	0.316	−0.010	0.030
L16	−0.0198	0.010	−1.991	0.046	−0.039	0.000
L17	0.0361	0.007	4.982	0.000	0.022	0.050

BIC: Bayesian information criterion

Figures

Figure B.1: Gas reserves and production

(a) Strategic Ellipse countries

(b) Natural gas reserves (trillion m^3)

(c) Natural gas production (billion m^3)

Note: the maps were generated using data in Table 1.1

Figure B.2: Common payoff functions

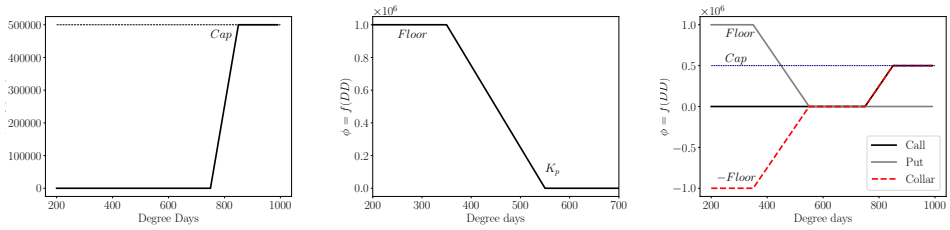

(a) Call with a cap (b) Put with a floor (c) Collar

Figure B.3: Process analysis for modeling energy demand

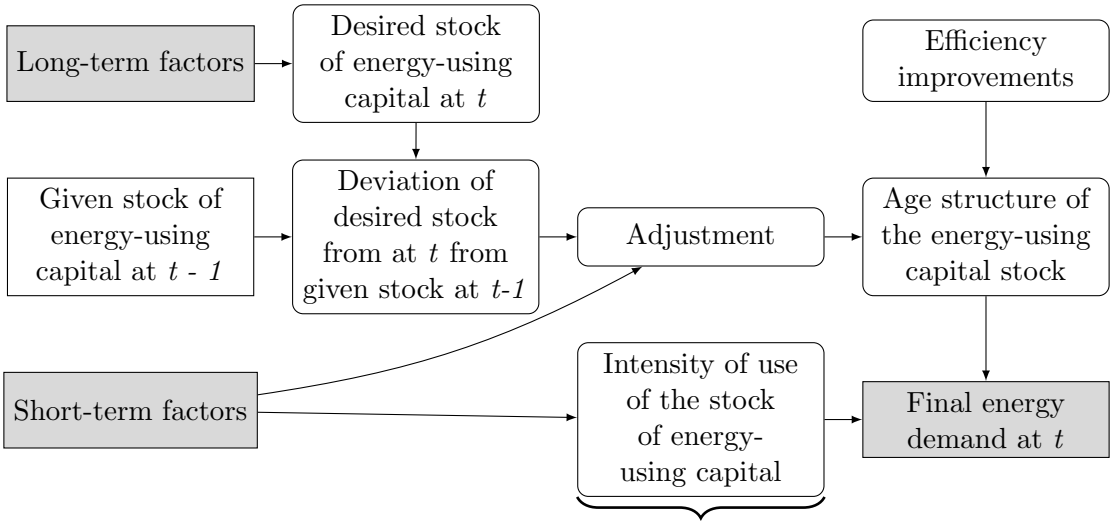

Note: With adjustments, investors verify whether the changes in factors influencing the desired stock are really long-term or just transitory

Figure B.4: Multiple Seasonal Trend decomposition using Loess of average temperatures

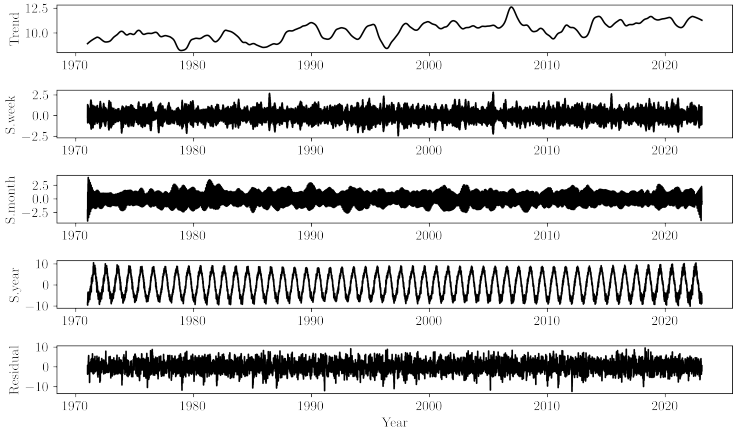

(a) Amsterdam

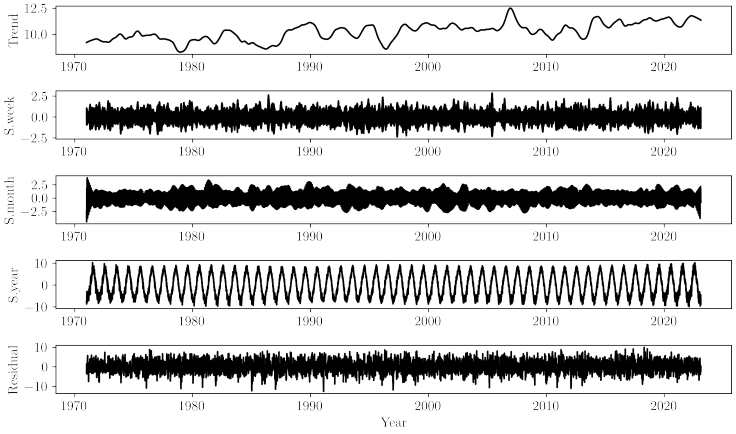

(b) Rotterdam

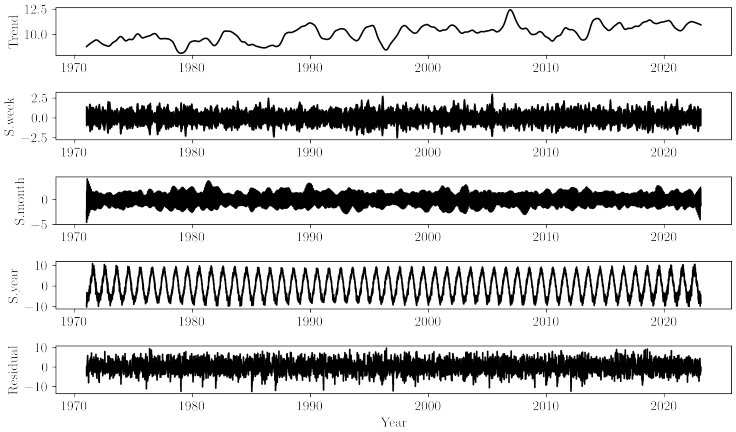

(c) Utrecht

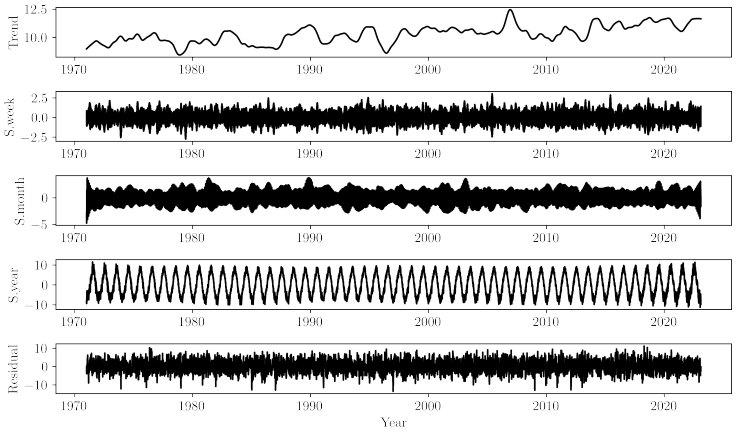

(d) Eindhoven

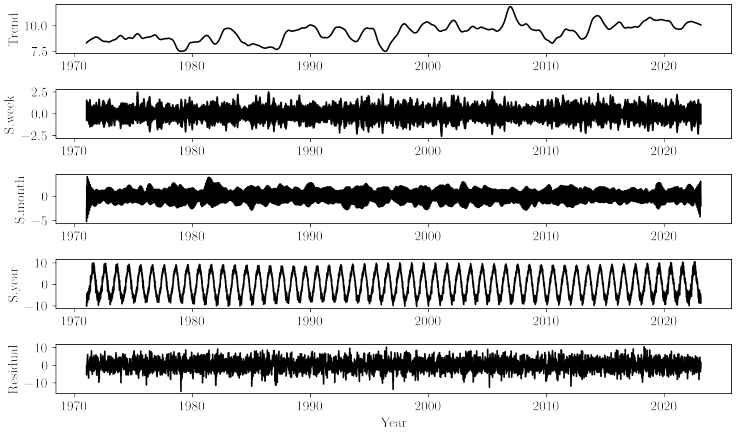

(e) Groningen

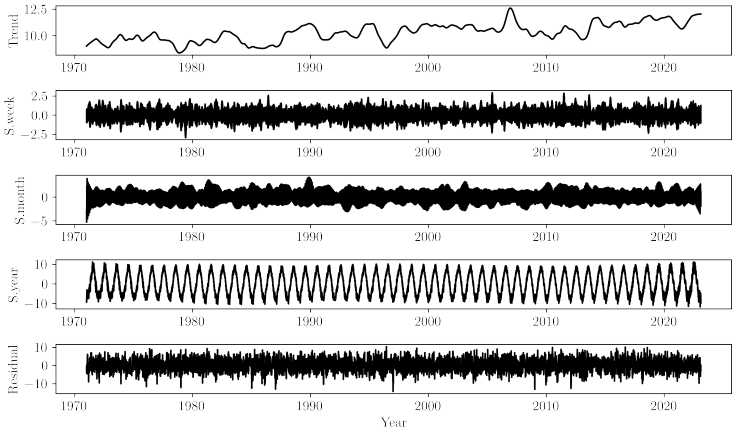

(f) Maastricht

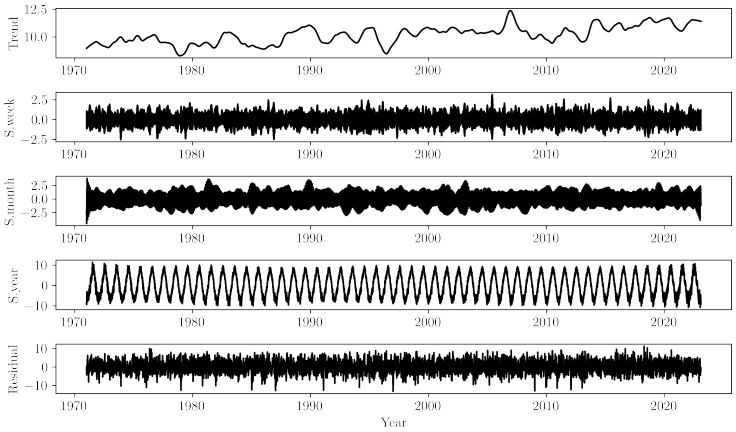

(g) Tilburg

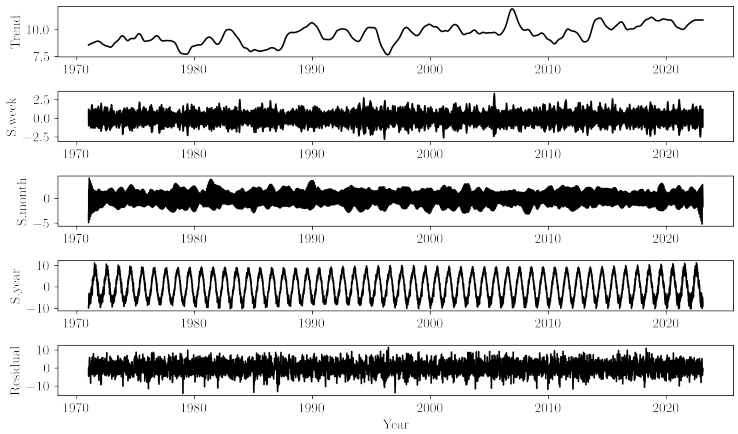

(h) Enschede

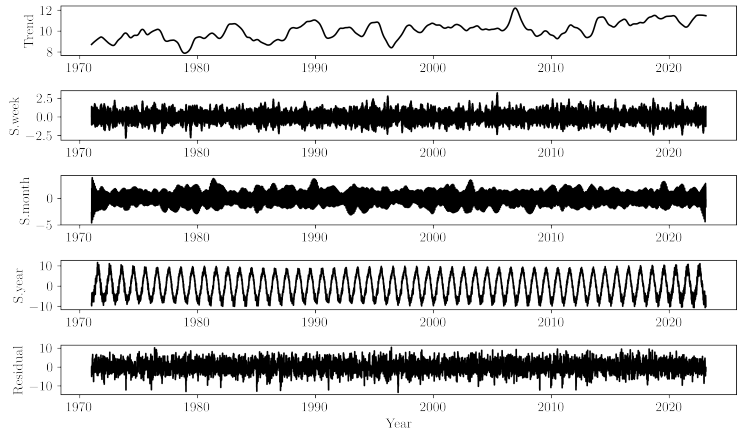

(i) Nijmegen