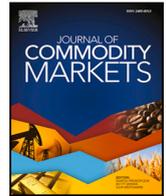


Contents lists available at [ScienceDirect](https://www.sciencedirect.com)

Journal of Commodity Markets

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

Regular article

Parametric heat wave insurance

Karl Larsson

School of Business, Unit of Statistics, Örebro University, Fakultetsgatan 1, S-70210 Örebro, Sweden

ARTICLE INFO

Keywords:

Heat waves
Insurance
Temperature
Time series
Pricing
Weather derivatives
Climate risk

ABSTRACT

This paper proposes a flexible framework for structuring and pricing parametric heat wave insurance. The framework is based on a general heat wave definition formulated in terms of an underlying temperature index. The definition can be varied in terms of the heat wave duration, intensity, measurement period and underlying index. This construction makes it straightforward to create contracts tailored to insure against heat events of many different types. A single stochastic model for the underlying index can be used to price all contracts. We consider contracts with payments that depend on the number of heat waves of a certain type occurring in the measurement period and derive the necessary pricing relations based on a general model structure encompassing several popular temperature models in the literature. An empirical case study is performed using data for Berlin where the daily maximum temperature is used as the underlying index. Model implied heat wave probabilities are consistent with historical patterns, point to high likelihoods for short duration heat events of different threshold temperatures and non-negligible risks for future heat waves of extreme temperatures and durations never before observed.

1. Introduction

With a changing climate many parts of the world are experiencing an increase in the frequency, duration and intensity of heat waves. Local governments and businesses are now in demand for insurance solutions that can help mitigate risks and provide economic compensation in the event of a costly period of extreme heat. The interest is evidenced e.g. by [Goering \(2020\)](#) who reports on an emerging market for heat wave insurance, and [Lamm et al. \(2020\)](#) who investigates the scope and potential use of such insurance relevant for California.

Prolonged periods of abnormally high temperatures are dangerous to human health and cause excess deaths and hospitalizations. In addition to increasing the pressure on the health care system heat waves can be directly damaging to a number of other sectors of society. Infrastructure for transportation such as roads, railways and runways can be made temporarily unusable leading to large costs due to delays and re-routing. Spells of hot weather often cause sharp rises in electricity demand e.g. due to increased use of cooling equipment. High temperatures will at the same time also affect power lines and reduce the efficiency of the transmission system. Nuclear plants may be forced to shut down if they use cooling water from natural sources, such as seas or rivers, that becomes too warm. These effects increase the risk of forced or uncontrolled power outages. Water supply systems are also affected when the heat drives up water demand typically under simultaneously dry conditions. Productivity in many sectors that depend on outdoor work such as agriculture, construction and tourism, is directly and adversely affected. We refer to [Smoyer-Tomic et al. \(2003\)](#) for a detailed description of these, and other examples, of the effects of heat waves.

In a recent report the World Meteorological Organization (WMO) lists the most damaging weather-related natural disasters across continents for the 50-year period from 1970 to 2019 ([WMO, 2021](#)). They estimate that heat waves accounted for 93% of the

E-mail address: karl.larsson@oru.se.

<https://doi.org/10.1016/j.jcomm.2023.100345>

Received 11 December 2022; Received in revised form 4 June 2023; Accepted 15 June 2023

Available online 6 July 2023

2405-8513/© 2023 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

weather-related human fatalities in Europe during the period. The heat wave in 2010 is the deadliest European weather event ever recorded with an estimated death toll of 55,736. Economic losses from heat waves are typically of smaller magnitude compared to that of droughts, floods and storms. According to [WMO \(2021\)](#) heat waves accounted for 4% of the total weather-related economic losses in Europe during the period 1970 to 2019. This is still a significant amount and constitutes the largest estimated loss due to heat waves of all continents. [Garcia-León et al. \(2021\)](#) perform a more detailed analysis of heat wave related costs in Europe and find that economic losses for recent heat waves in 2003, 2010, 2015 and 2018 amounted to 0.3–0.5% of GDP. They further find that the economic losses show large spatial variation, for some regions exceeding 1% of local GDP, and that heat wave related costs are projected to increase steadily for the next 40 years.

Damages from natural disasters are often insured by means of parametric insurance contracts where payments are triggered if a certain well-defined disastrous event is realized. Parametric insurance is different from indemnity type policies where the insured party need to establish the costs caused by the event and make a claim to the insurance provider. This is often difficult for many types of weather-related events including heat waves. A parametric policy will instead provide a pre-determined payment when the event occurs regardless of the actual damages. These types of policies therefore have the benefits that lengthy investigations into the actual damages are not required and that the compensation can be paid with very little delay. The parametric insurance construction is often used in so-called CAT-bonds. Many weather derivatives are also defined this way and indeed qualify as parametric insurance products, see e.g. the discussions in [Alaton et al. \(2002\)](#) and [Benth and Šaltytė-Benth \(2013\)](#). Derivative assets with payoffs contingent on temperature are among the most common weather derivatives. So called heating degree day (HDD), and cooling degree day (CDD), futures contracts for different cities are e.g. traded on the exchanges organized by the CME Group. These contracts are designed to provide insurance against increases in energy demand during cold or warm periods and are very different from the contracts proposed here that target more extreme temperatures. A thorough review of the weather derivatives market and the related literature can be found in [Benth and Šaltytė-Benth \(2013\)](#). Although the literature on temperature derivatives is well developed there is, to the best of our knowledge, no studies that specifically addresses parametric heat wave insurance.

In this paper we develop a framework for structuring and pricing contracts designed to provide insurance against heat wave events of different types. We show how to construct insurance contracts who makes payment if a suitably defined heat wave event occurs within a specified period of time. The framework is based on a general heat wave definition together with a stochastic model for the underlying temperature index. Any type of contract can then be priced using the model and standard arbitrage pricing methodology. The heat wave definition is formulated in terms of a daily temperature index and where the threshold value, duration and measurement period can all be varied. This approach has the benefit of making it possible to create contracts that insure against different types of heat waves in a flexible manner. Contract specifications can easily be tailored to different customer needs and local climates. It is therefore straightforward to employ the framework in many different insurance scenarios and geographical regions. The diverse effects of heat waves may make it difficult to establish a standard insurance policy that can simultaneously meet the needs of many different parties. Insurance industry experts therefore expect the need for tailor-made contracts in this area ([Goering, 2020](#)). This flexibility is obtained with our proposed framework. We discuss contracts designed as bonds that give payments depending on the number of heat events of a specified duration and intensity occurring in the measurement period. The prices for these contracts are expressed in terms of the conditional probabilities of the specified heat events. Modeling the time series of the underlying temperature index makes it possible to price and evaluate different heat wave contracts using the same stochastic model which is appealing from a practical perspective. We propose modeling the temperature index using a discrete time framework. In the empirical part we use a specific model similar to the one in [Šaltytė-Benth and Benth \(2012\)](#).

Our proposed framework is illustrated in an empirical case study using data for Berlin. We use a heat wave definition based on the daily maximum temperature as the underlying index. This is a common choice and used in many studies, see e.g. [Fenner et al. \(2019\)](#) who investigates several different definitions using data for Berlin. [Smoyer-Tomic et al. \(2003\)](#) find that heat wave definitions formulated in terms of maximum daily temperature efficiently describe conditions that entail economic losses for businesses and society thus making the choice relevant for insurance purposes. We use data on daily maximum temperature for a near 65 year period ranging from 1957-01-01 to 2021-08-04 to investigate the historic occurrences of heat waves and model performance. The model estimation results show that the proposed model provides an accurate representation of the historical time series of daily temperature maximums. The daily mean temperature is the common choice of index found in most empirical model evaluations in the related literature, see e.g. [Benth and Šaltytė-Benth \(2013\)](#), [Šaltytė-Benth and Benth \(2012\)](#), [Schiller et al. \(2012\)](#), [Benth et al. \(2008\)](#), [Campbell and Diebold \(2005\)](#) and [Alaton et al. \(2002\)](#). Our results show that the model works equally well for time series of daily temperature maximums. Heat waves of different duration and intensity has increased in Berlin over the sample period. The historical pattern of heat waves is consistent with a trend of more frequently occurring heat spells in line with the results in [Fenner et al. \(2019\)](#). The historical record indicate that shorter heat waves of high temperature tend to be more common than longer waves of lower intensity. Simulated heat wave probabilities for the summer months of 2022 are consistent with these patterns. The model implied probabilities point to high likelihoods for short duration heat waves, but also to non-negligible probabilities of waves of extreme temperature levels and durations never before observed.

The paper is organized as follows. Section 2 introduces the heat wave definition, different types of contracts and general pricing results. The proposed model for the underlying temperature index is presented in Section 3. Section 4 contains empirical results on historical heat waves, model estimation and model implied heat wave probabilities for the Berlin data. Section 5 concludes.

2. Parametric heat wave insurance contracts

2.1. Heat wave definition

There are many heat wave definitions in the existing literature. Some definitions serve the purpose of describing conditions that are either harmful to human health or damaging to infrastructure, while others are used to detect deviations from historical observations e.g. in studies of climate change. However, most definitions states that a heat wave is a period of consecutive days where some suitably chosen temperature index exceeds a certain threshold value, see e.g. Fenner et al. (2019), Chen and Li (2017), Smith et al. (2013) and Robinson (2001).

We define a heat wave as the event

$$A(n, a, \tau, Z) = \{Z_t \geq a \text{ for at least } n \text{ consecutive days } t \text{ where } t \in \tau\} \quad (1)$$

where Z_t denotes the value of the temperature index at time t , n is the duration, a is the threshold value and $\tau = [\tau_a, \tau_b]$ is the measurement period. This is a general definition that encompasses a wide variety of different types of heat waves. We have included a measurement period τ that starts at time τ_a and ends at time τ_b . Including a measurement period in the heat wave definition is not standard but is introduced for the purpose of constructing well-defined financial contracts. It also highlights the stochastic nature of heat wave events and that the randomness comes from the stochastic behavior of the temperature index Z during the period τ .

The heat wave event $A(n, a, \tau, Z)$ in (1) illuminates the four main components used to specify contracts based on heat wave events. First, a temperature index Z must be chosen. A natural choice is to let Z be the daily maximum temperature at a clearly specified measurement station. This is the perhaps most common index but other choices are possible e.g. the daily mean temperature or a combination of different indexes, see Smith et al. (2013), Robinson (2001) and Fenner et al. (2019). In the empirical part we will take Z to be the daily maximum temperature. Smoyer-Tomic et al. (2003) argue that for heat wave definitions aimed at describing costly effects on infrastructure a criteria based on daily temperature maximum is sufficient which further motivates this choice. We also briefly consider the index $\tilde{Z} = \max(Z_1, Z_2, Z_3)$, where Z_1 , Z_2 and Z_3 are the daily maximum temperatures at three nearby locations, as a way of incorporating the spatial dimension.

A threshold a must be chosen to define what is meant by a sufficiently extreme temperature. This will depend on the chosen index Z . The threshold can be set in absolute (static) or relative (dynamic) terms. An absolute threshold would be to simply set a value for a that, by some rational, defines an extreme temperature level. A relative way of determining the threshold a is to set it in relation to historical conditions. An example is to let a be the 90th percentile of the temperature index distribution for some reference period. The relative approach captures the notion that a heat wave is a period when it is warmer than usual which may happen also in winter time and does not necessarily correspond to a period of damaging or hazardous temperatures. Relative heat wave definitions are useful e.g. in studies of climate change.

The duration n specifies the number of consecutive days during which the threshold is exceeded. Common choices of n are 3 to 5 days, but heat waves of moderate threshold temperatures can extend for far longer periods.

Finally, a measurement period τ must be specified that determines for which time period heat wave occurrences are recorded. A measurement period is not directly relevant for the definition of a heat wave event but is important for defining contracts based on (1). In Europe heat waves usually strike in the summer months June, July and August and these are natural choices for τ . This may differ from other parts of the world. One may naturally also consider longer periods e.g. the entire summer period June-August.

Once the temperature index Z is chosen the components a , n and τ can be varied. The duration n and threshold a can be set so as to define more or less prolonged and intense heat wave events. This is appealing from a contract perspective and makes it easy to create contracts that provide insurance for different more or less extreme heat waves and for different periods of the year.

2.2. Heat wave contracts

Let $\Gamma(n, a, \tau, Z)$ record the number of distinct heat events of type $A(n, a, \tau, Z)$ in the measurement period τ . It is natural to only allow $\Gamma(n, a, \tau, Z)$ to count distinct heat waves that occur in disjoint non-overlapping time periods. A heat wave of duration $n = 6$ is for example not allowed to be counted as two waves of duration $n = 3$. The most basic heat wave contract we consider has the payoff

$$\Phi(T, K, n, a, \tau, Z) = K \times 1\{\Gamma(n, a, \tau, Z) \geq 1\} \quad (2)$$

where $1\{B\}$ denotes the indicator function of an event B , K is the amount of money being paid if the event $\{\Gamma(n, a, \tau, Z) \geq 1\}$ occurs, and T is the payment, or maturity, date. The payment is made after the end of the measurement period when all uncertainty is resolved. We therefore set $T = \tau_b$ for simplicity. In (2) the amount K is paid in the event that at least one heat wave of type $A(n, a, \tau, Z)$ occurs within the measurement period which defines a natural and basic type of parametric insurance.

The payoff (2) makes the same payment regardless of whether there is one or several heat waves of a given type $A(n, a, \tau, Z)$ in the measurement period. This is natural since in many cases the likelihood of more than one heat wave occurring in the measurement period will be low. However, other contracts can easily be constructed. One possibility is to define a contract that makes the payment contingent on the exact number of heat waves occurring within the measurement period. Such contracts will have the payoffs

$$\phi_m(T, K_m, n, a, \tau, Z) = K_m 1\{\Gamma(n, a, \tau, Z) = m\}, \quad m = 0, \dots, N(n, \tau) \quad (3)$$

where the index m is the number of heat waves of duration n and threshold a that occurs in the period τ . The payment amount K_m depends on m as well. In a given measurement period there is a maximum number of possible separate heat waves of length n and hence m takes values in a set of integers $\{1, \dots, N(n, \tau)\}$ where $N(n, \tau)$ depends both on the number of days in the measurement period and the duration. Each index m defines a separate contract based on (3) for $m = 1, \dots, N(n, \tau)$. Other possible contract variations are possible but we will focus attention on the contracts defined by the payoff Φ in (2).

The heat wave definition (1) allows for flexible structuring of different contracts. This is important for being able to tailor the contracts to various customer requests and to be relevant in geographical regions with different climates. It also allows for the construction of contract portfolios. One may e.g. combine contracts that pay out different amounts depending on the duration n or the threshold temperature a . For example, suppose the threshold is set to $a = 30^\circ$ and two contracts of type (2) are defined for durations $n = 3$ and $n = 5$. This gives two different event definitions $A_1 = A(3, 30^\circ, \tau, Z)$ and $A_2 = A(5, 30^\circ, \tau, Z)$ with corresponding payments $\Phi(T, K_1, 3, 30^\circ, \tau, Z)$ and $\Phi(T, K_2, 5, 30^\circ, \tau, Z)$ with potentially different notional payments K_1 and K_2 . The portfolio payoff at maturity T is K_1 if $Z_t \geq 30^\circ$ for at least 3 consecutive days, and $K_1 + K_2$ if $Z_t \geq 30^\circ$ for at least 5 consecutive days, and it is zero if there is no heat wave occurring in the measurement period.

2.3. Contract valuation

The contracts introduced in Section 2.2 are a form of temperature derivatives where the temperature index Z serves as the underlying asset. To price the contracts with the payoff (2) and (3) we assume that the temperature index Z is defined on a probability space Ω equipped with a probability measure \mathbb{P} and a filtration \mathcal{F}_t . From arbitrage theory the value at time t of a contract that pays the random amount $\Phi(T)$ at maturity T is given by the discounted expected value of $\Phi(T)$ under an equivalent probability measure \mathbb{Q} usually referred to as the risk-neutral measure. The fact that the temperature index Z is not a traded asset makes the market incomplete and therefore the measure \mathbb{Q} is not unique and cannot be determined from arbitrage theory alone. These are a standard result in asset pricing, see e.g. Björk (2019) for a general treatment and Benth and Šaltytė-Benth (2013), Benth et al. (2008) and Schiller et al. (2012) for discussions specifically for temperature derivatives. Whether there is a difference between the probability measures \mathbb{P} and \mathbb{Q} depends on if there is a non-zero market price of risk (MPR) associated with temperature uncertainty. If temperature risk is not priced then $\mathbb{Q} = \mathbb{P}$. Schiller et al. (2012) compare a number of pricing models for temperature derivatives and argue that the market price of risk should be zero since the temperature at a certain location can be assumed independent from economy-wide activity. They hence assume that $\mathbb{Q} = \mathbb{P}$. This is a standard assumption in other studies of temperature derivatives, see e.g. Alaton et al. (2002), Benth and Šaltytė-Benth (2005), Campbell and Diebold (2005), Benth et al. (2008), Ahčan (2012), Schiller et al. (2012) and Šaltytė-Benth and Benth (2012). In Benth et al. (2011) the authors challenge this assumption. In order to give a complete framework we show how a risk neutral measure \mathbb{Q} can be explicitly constructed in the discrete time setting in Section 3.2.

Assuming for now that a risk neutral measure \mathbb{Q} has been established, the price of the basic contract with payment (2) is given by

$$V(t, T) = e^{-r(T-t)} KE^{\mathbb{Q}} [1\{\Gamma(n, a, \tau, Z) \geq 1\} | \mathcal{F}_t] \tag{4}$$

where r is the continuously compounded interest rate which we assume constant. Since the expectation of the indicator function of a certain event is the probability of the event we have

$$V(t, T) = e^{-r(T-t)} K \xi^{\mathbb{Q}}(t, n, a, \tau, Z) \tag{5}$$

where

$$\xi^{\mathbb{Q}}(t, n, a, \tau) = \mathbb{Q}(\Gamma(n, a, \tau, Z) \geq 1 | \mathcal{F}_t). \tag{6}$$

denotes the heat wave probability relevant for the contract in question. Since the time T payoff $K1\{\Gamma(n, a, \tau, Z) \geq 1\}$ is a non-linear function of the temperature index Z the contract is in fact a type of digital, or binary, option. Here we are assuming that the buyer of the contract pays the amount $V(t, T)$ at t to receive the payoff $K1\{\Gamma(n, a, \tau, Z) \geq 1\}$ at T . We will refer to this way of scheduling the payments as a bond structure since it resembles the cashflow scheme of a zero-coupon bond. The value $V(t, T)$ will be therefore be referred to as the bond price although it is a form of option. Indeed, this is the typical basic structure used in a parametric catastrophe (CAT) bonds that also have non-linear payoffs with prices determined by the probability of the event triggering payment.

We remark that it is common that derivative contracts are structured as forwards. Exchange traded temperature derivatives such as heating degree and cooling degree futures are e.g. structured this way. In a forward contract the parties instead agree to exchange payments at contract maturity T . In this case the buyer will both pay the fixed amount H and receive the stochastic amount $K1\{\Gamma(n, a, \tau, Z) \geq 1\}$ from the seller at the time of maturity T . The amount H is set so that the value of the transaction at the time the contract is entered is zero to both parties. If the contract with payoff (2) is structured this way the value at time t is in general given by

$$V_F(t, T) = e^{-r(T-t)} E^{\mathbb{Q}} [K1\{\Gamma(n, a, \tau, Z) \geq 1\} - H | \mathcal{F}_t]$$

Assuming the contract is entered into at time t it follows that H should be set according to $H = F(t, T)$ where $F(t, T) = KE^{\mathbb{Q}} [1\{\Gamma(n, a, \tau, Z) \geq 1\} | \mathcal{F}_t]$ such that $V_F(t, T) = 0$. The quantity $F(t, T)$ is referred to as the forward price. Using the heat wave probability (6) it simplifies into the expression $F(t, T) = K \xi^{\mathbb{Q}}(t, n, a, \tau, Z)$. Hence, the essential quantity that require calculation is still the heat wave probability in (6) no matter how the cashflows are structured in time.

The contracts defined by the payoffs in (3) are handled the same way. In the case of a bond structure, the price is given by $V_m(t, T) = e^{-r(T-t)} K_m \xi_m^Q(t, n, a, \tau)$ where the contract associated heat wave probabilities are $\xi_m^Q(t, n, a, \tau) = \mathbb{Q}(\Gamma(n, a, \tau, Z) = m | \mathcal{F}_t)$, $m = 0, \dots, N(n, \tau)$.

3. Stochastic temperature modeling

To be able to calculate heat wave probabilities and contract prices a stochastic model for the temperature index Z is needed. Stylized features of temperature data are seasonality in both the mean and variance and some degree of low order autocorrelation. We do not propose a new model here and instead present a single useful model based on the works in Benth and Šaltytė-Benth (2013), Šaltytė-Benth and Benth (2012), Benth et al. (2011), Campbell and Diebold (2005), and Alaton et al. (2002). Our study is different from the previous literature since in the empirical section we take the temperature index to be the daily maximum temperature. The mentioned studies exclusively use some version of daily mean temperature. Our empirical results show that the model proposed here can be used successfully also for a time series of daily maximums.

The literature on derivative asset pricing often assumes a continuous time diffusion model for the underlying index or asset. We instead choose to work in a discrete time setting. There are benefits and drawbacks with both approaches. Derivative pricing models formulated as traditional time series models in discrete time have recently become more popular and important advances in this area have been made by e.g. Heston and Nandi (2000), Christoffersen et al. (2012), Chorro et al. (2012), Guégan et al. (2013) and Ornthanalai (2014). In the time series framework it is straightforward to include, and estimate, both seasonal and auto-regressive properties in the conditional mean. This is also possible in continuous time using so-called continuous time autoregressive (CAR) processes, see e.g. Benth and Šaltytė-Benth (2013) and Benth et al. (2008). However, these models introduce latent processes that require special treatment and are therefore more difficult to work with.

3.1. Temperature models in discrete time

We consider the general time series specification for the temperature index Z_t under \mathbb{P} given by

$$Z_t = \eta_t + \sigma_t \varepsilon_t \tag{7}$$

where ε_t is an iid sequence of random variables. The processes η_t and σ_t are \mathcal{F}_{t-1} -measurable and define the conditional mean and standard deviation of Z_t . This general model structure encompasses most of the proposed discrete time temperature models found in the literature. In the empirical section we will use the following specification, which is close to the one proposed in Šaltytė-Benth and Benth (2012), and set

$$Z_t = S_t + X_t \tag{8}$$

where S_t is a deterministic function capturing trend and seasonality and X_t the stochastic process part. The seasonal function is specified as a truncated Fourier series according to

$$S_t = a_0 + a_1 t + \sum_{k=1}^{N_S} (b_{2k-1} \cos(2\pi k \omega t) + b_{2k} \sin(2\pi k \omega t)) \tag{9}$$

where $\omega = 365^{-1}$ correspond to a yearly seasonal pattern for the daily data, and N_S is truncation point for the trigonometric series. The coefficients a_0 and a_1 in (9) capture the mean level and trend respectively, and the coefficients b_k determine the seasonal pattern. The stochastic process X is the seasonal and trend adjusted temperature given by $X_t = Z_t - S_t$. We model X_t as a zero mean AR(p)-process according to

$$X_t = \sum_{i=1}^p \lambda_i X_{t-i} + u_t \tag{10}$$

where λ_i are constant coefficients. The residual u_t is specified as

$$u_t = \sigma_t \varepsilon_t \tag{11}$$

where σ_t is a time varying volatility and ε_t an i.i.d. sequence of stochastic variables. This specification fits the structure (7) with conditional mean given by

$$\eta_t = S_t + \sum_{i=1}^p \lambda_i (S_{t-i} - Z_{t-i})$$

and is often used in the temperature modeling literature, see e.g. Alaton et al. (2002), Šaltytė-Benth and Benth (2012), and Benth et al. (2008).

It is well established that temperature display seasonal heteroscedasticity. Šaltytė-Benth and Benth (2012), Campbell and Diebold (2005) and Ahčan (2012) also find ARCH-effects in temperature time-series. To allow for both sources of heteroscedasticity we model the conditional variance σ_t^2 as

$$\sigma_t^2 = \sigma_{S,t}^2 \sigma_{G,t}^2 \tag{12}$$

where $\sigma_{S,t}^2$ is given by a deterministic seasonal function and $\sigma_{G,t}^2$ is a GARCH process. For the seasonal variance part we set

$$\sigma_{S,t}^2 = c_0 + \sum_{k=1}^{N_\sigma} (c_{2k-1} \cos(2\pi k\omega t) + c_{2k} \sin(2\pi k\omega t)) \tag{13}$$

where c_k are constants, $\omega = 365^{-1}$, and N_σ is a truncation point. For the GARCH-part we set

$$\sigma_{G,t}^2 = 1 + d_1 \left(\sigma_{G,t-1}^2 - 1 \right) + d_2 \left(\frac{u_{t-1}^2}{\sigma_{S,t-1}^2} - 1 \right) \tag{14}$$

With this specification the unconditional expectation of $\sigma_{G,t}^2$ is restricted to $d_0 = (1 - d_1 - d_2)^{-1} = 1$ which is a natural requirement given the multiplicative structure. The general structure (7) accommodates many other specifications of the conditional variance. Campbell and Diebold (2005) for example use an additive specification combining seasonal and GARCH parts and Alaton et al. (2002) use a piece-wise constant seasonal variance.

Regarding the distribution of the standardized residuals $\varepsilon_t = u_t/\sigma_t$ the standard assumption in the literature is that they are $N(0, 1)$ -distributed. Benth and Šaltytė-Benth (2005) and Ahčan (2012) find evidence that the Normal Inverse Gaussian (NIG) distribution may be a better fit for some data. We will investigate both these distributional assumptions in the empirical section. This completes the model dynamics for the measure \mathbb{P} under which the time series for Z_t is observed.

3.2. Esscher transform and change of measure

As discussed in Section 2.3 a common assumption when pricing temperature derivatives is that $\mathbb{Q} = \mathbb{P}$ which is not without economic motivation as argued e.g. in Schiller et al. (2012). However, to give a more complete framework and allow for the possibility that $\mathbb{Q} \neq \mathbb{P}$ we demonstrate how an equivalent risk neutral measure \mathbb{Q} that differs from \mathbb{P} can be derived for the general discrete time model structure in Eq. (7).

Using the Esscher transform method of Gerber and Shiu (1994) it is possible to find the distribution and dynamics of Z under a different pricing measure \mathbb{Q} should that be warranted. The Esscher transform has also been used in the continuous time literature see e.g. Benth and Šaltytė-Benth (2005) and Ahčan (2012) for applications to temperature. In contrast to these authors we demonstrate how this method can be applied also in the discrete time framework which is rather different. The discrete time approach, with some variations, has been used in the financial time series literature, see e.g. Christoffersen et al. (2012), Chorro et al. (2012) and Guégan et al. (2013). Models for temperature are very different from financial asset models since stylized features include seasonality and autoregression in both the conditional mean and variance. These features are easily handled in discrete time and with the associated pricing results a consistent framework is achieved that does not require transitions between discrete and continuous time.

The Esscher transform of Z_t defined as in (7) is given by

$$M_t = \frac{e^{\theta_t Z_t}}{E^{\mathbb{P}} [e^{\theta_t Z_t} | \mathcal{F}_{t-1}]} = \frac{e^{\theta_t Z_t}}{\varphi_{t-1}^Z(\theta_t)} \tag{15}$$

where θ_t is a \mathcal{F}_{t-1} -measurable process, $\varphi_{t-1}^Z(s) = E^{\mathbb{P}} [\exp(sZ_t) | \mathcal{F}_{t-1}]$ denotes the conditional moment generating function for Z_t and $M_0 = 1$. The process θ_t is referred to as the market price of risk. Next, define the following two processes

$$L_t = \prod_{k=1}^t M_k \tag{16}$$

with $L_0 = 1$, and

$$\Lambda_{t,T} = \frac{L_T}{L_t} \tag{17}$$

Then L_t is a \mathbb{P} -martingale and $\Lambda_{t,T}$ defines a new equivalent measure \mathbb{Q} under which the relation

$$E^{\mathbb{Q}} [Y_T | \mathcal{F}_t] = E^{\mathbb{P}} [\Lambda_{t,T} Y_T | \mathcal{F}_t], \quad t \leq T,$$

holds for any \mathcal{F}_T -measurable random variable Y_T . In particular we have that the conditional moment generating function $\varphi_{t-1}^Z(s; \mathbb{Q}) = E^{\mathbb{Q}} [\exp(sZ_t) | \mathcal{F}_{t-1}]$, $s \in \mathbb{R}$, for Z_t under \mathbb{Q} is given by $\varphi_{t-1}^Z(s; \mathbb{Q}) = E^{\mathbb{P}} [\Lambda_{t-1,t} \exp(sZ_t) | \mathcal{F}_{t-1}]$. From (16) and (17) we have that $\Lambda_{t-1,t} = L_t/L_{t-1} = M_t$ and using (15) we find that

$$\varphi_{t-1}^Z(s; \mathbb{Q}) = \frac{\varphi_{t-1}^Z(s + \theta_t)}{\varphi_{t-1}^Z(\theta_t)}$$

Substituting the general model structure for Z_t from (7) we have

$$\varphi_{t-1}^Z(s) = E^{\mathbb{P}} [e^{s\eta_t + s\sigma_t \varepsilon_t} | \mathcal{F}_{t-1}] = e^{s\eta_t} \varphi_{t-1}^\varepsilon(s\sigma_t) \tag{18}$$

where $\varphi_{t-1}^\varepsilon(s) = E^{\mathbb{P}} [\exp(s\varepsilon_t) | \mathcal{F}_{t-1}]$ denotes the conditional moment generating function for ε_t under \mathbb{P} . Straightforward calculation show that we may conclude that $\varphi_{t-1}^Z(s; \mathbb{Q})$ can be expressed as

$$\varphi_{t-1}^Z(s; \mathbb{Q}) = e^{s\eta_t} \frac{\varphi_{t-1}^\varepsilon((s + \theta_t)\sigma_t)}{\varphi_{t-1}^\varepsilon(\theta_t\sigma_t)} \tag{19}$$

Table 1
Summary statistics for daily maximum temperature observations at Berlin-Tempelhof for the sample period 1957-01-01 to 2021-08-04.

Location	mean	std.dev.	Skewness	Kurtosis	max	min
Tempelhof	13.7049	8.9998	-0.0262	2.2560	38.5	-17.0

In many cases of interest the conditional distribution of Z under \mathbb{Q} can now be easily evaluated from (19). Note that if $\theta_t = 0$ for all t then $M_t = 1$ and $\mathbb{P} = \mathbb{Q}$.

If ε_t is $N(0,1)$ -distributed, for all t , then $\varphi_{t-1}^\varepsilon(s) = \exp(s^2/2)$ and we find from (19) that

$$\varphi_{t-1}^Z(s; \mathbb{Q}) = e^{s(\eta_t + \theta_t \sigma_t^2) + \frac{1}{2} s^2 \sigma_t^2}$$

showing that $Z_t | \mathcal{F}_{t-1}$ is $N(\eta_t + \theta_t \sigma_t^2, \sigma_t^2)$ -distributed under \mathbb{Q} . The time series dynamics can therefore be written

$$Z_t = \eta_t + \theta_t \sigma_t^2 + \sigma_t \varepsilon_t^\mathbb{Q}$$

where $\varepsilon_t^\mathbb{Q}$ is an iid sequence of $N(0,1)$ -distributed random variables. The change of measure preserves the model structure and the conditional distribution is still Gaussian under \mathbb{Q} . The only change is that a compensation for risk term $\theta_t \sigma_t^2$ appears in the conditional mean of Z_t . We refer to Appendix for an example of how this change of measure works when ε_t are assumed to follow a normal inverse Gaussian distribution.

The common way of pinning down the market price of risk, and hence the measure \mathbb{Q} , is to make assumptions on the functional form of θ_t which is estimated from market data on derivative assets. In our setting there is however no such data directly available.

4. Empirical case study: heat waves in Berlin

In this section we perform an empirical case study using data for Berlin. We present results for historical heat wave patterns, model estimation, and simulated heat wave probabilities. There are several reasons for choosing data for Berlin. Berlin is one of the largest metropolitan areas in Europe and heat warnings are frequently issued in Berlin during summers by the Deutscher Wetterdienst (DWD), see e.g. Matzarakis et al. (2020). Heat waves in Berlin have also been studied in e.g. Fenner et al. (2014, 2019). In addition, the DWD has an excellent service for historical temperature data.

4.1. Data and historical heat wave characteristics

We use data from the Deutscher Wetterdienst (DWD) for the weather station Tempelhof to represent Berlin urban temperature.¹ The Tempelhof weather station is located ca 2.5 km south of the city center of Berlin. Fenner et al. (2019) also use Tempelhof station as representative for studying heat waves in inner city Berlin. We use a sample of data spanning from 1957-01-01 to 2021-08-04 and consisting of 23,592 observations of daily temperature maximums. Summary statistics for the time series of daily temperature maximums during the sample period are presented in Table 1. The maximum record for Tempelhof in Table 1 is from June 30 in 2019. We start with an empirical investigation of the historical heat wave occurrences in Berlin for different variations of the general definition $A(n, a, \tau, Z)$ in (1). Note that we are then effectively assuming that the measurement period consist of our entire sample from 1957-01-01 to 2021-08-04. Table 2 reports the number of identified heat waves of durations $n = 3, \dots, 11$ and threshold temperatures $a = 30^\circ, \dots, 36^\circ$. There are 68 heat waves of duration $n = 3$ for $a = 30^\circ$ but only 2 when $a = 36^\circ$. For values of $a \geq 37^\circ$ there are no recorded heat waves of any duration although the heat event in July 2010 comes close with 2 days above 37° and one day at 36.9° . The longest observed heat wave duration is $n = 11$ observed in 1994 and 2006. In 2006 the 11 day long heat wave was observed for $a = 30^\circ$ and in 1994 for $a = 31^\circ$. The 1994 heat wave was both long and hot including three consecutive days of temperatures higher than $a = 36^\circ$. The only other 3-day heat wave with temperatures above 36° is the one mentioned in July 2010. As further seen from Table 2, increasing the threshold by 1° will quickly reduce the number of heat waves of any duration. For the shortest duration $n = 3$ the reduction in numbers from 68, at $a = 30^\circ$, to 2, at $a = 36^\circ$, takes place in an exponential fashion. Longer durations at increasingly high thresholds a are more rare and for thresholds larger than $a = 33^\circ$ there are no observed heat waves of duration $n = 5$ or longer. Increasing the duration by 1 day reduces the number of heat waves in a similar, but slightly faster, fashion compared to increasing the threshold temperature by 1° .

Fig. 1 plots the number of heat waves of threshold temperature $a = 30^\circ$ per year for different combinations of durations n during our sample. The graph indicates that the frequency of 3-day heat waves is increasing sharply towards the end of our sample. Indeed, there are several more occurrences in the late sample with 55.9% of the 3-day heat waves, at $a = 30^\circ$, taking place after year 2000 which corresponds to roughly the last third of the sample. There are only six heat waves of duration $n = 3$ exceeding the threshold $a = 34^\circ$ and four of them has occurred after 2009. It may prove relevant in some circumstances to define a temperature index that acknowledges the spatial dimension. Costs of heat waves are typically concentrated to densely populated areas and especially larger cities so for insurance purposes large areas, extending far beyond city limits, are likely to be of less interest. Nevertheless,

¹ Data from DWD can be freely downloaded from the climate data center at <https://cdc.dwd.de/portal/>.

Table 2
Number of heat waves recorded at Tempelhof station 1957-01-01 to 2021-08-04 for different durations (n) and threshold temperatures (a).

n	a						
	30°	31°	32°	33°	34°	35°	36°
3	68	36	21	8	6	3	2
4	34	14	6	2	1	0	0
5	16	10	3	0	0	0	0
6	7	4	1	0	0	0	0
7	5	2	1	0	0	0	0
8	4	1	1	0	0	0	0
9	2	1	1	0	0	0	0
10	2	1	1	0	0	0	0
11	2	1	0	0	0	0	0

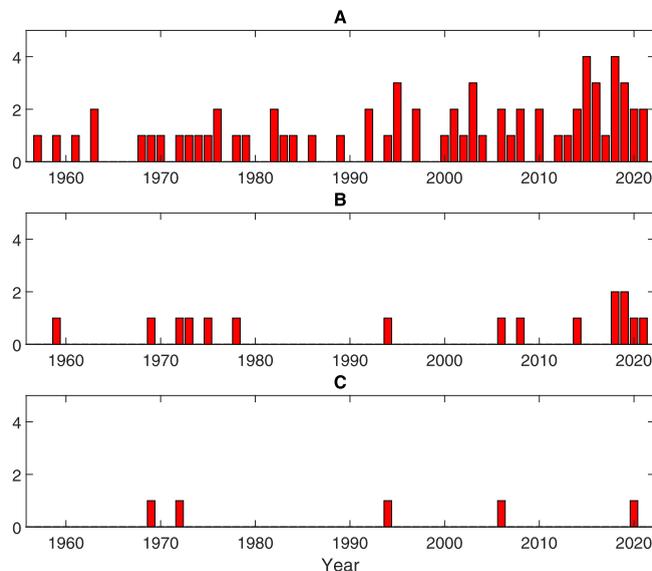


Fig. 1. Number of heat waves per year during 1957-01-01 to 2021-08-04 of threshold temperature $a = 30^\circ$ and durations $n = 3$ (panel A), $n = 5$ (panel B) and $n = 7$ (panel C).

spatial variation may prove relevant. Cities are e.g. known to be affected by the so-called Urban Heat Island (UHI) effect where infrastructure, such as buildings, pavements etc, will absorb and conserve heat thereby preventing temperatures from dropping as fast as they would otherwise have done during nighttime, see e.g. [Arnfield \(2003\)](#) and [Stewart \(2010\)](#). Due to the UHI effect warmer temperatures may arise in inner city areas compared to more peripheral locations. [Fenner et al. \(2014\)](#) provide a study of spatial temperature differences in Berlin and observe higher temperatures in the city compared to nearby rural areas.

We include data from two nearby weather stations to assess the relation between heat wave occurrences and daily maximum temperatures at nearby locations. The closest weather stations to Tempelhof with available data are the stations in Berlin-Brandenburg and Potsdam.² These stations are located in the periphery of urban Berlin. [Fenner et al. \(2019\)](#) find that Tempelhof is subject to the UHI effect while the Potsdam and Brandenburg stations are not.

One way of incorporating the spatial dimension is to use the alternative temperature index

$$\tilde{Z}_t = \min(Z_t^1, Z_t^2, Z_t^3)$$

where Z_t^i , $i = 1, 2, 3$, denote the daily temperature maximums at the stations in Tempelhof, Brandenburg and Potsdam. The index \tilde{Z}_t defined this way require that a heat event $A(n, a, \tau, Z^i)$ is observed at each station for a heat wave to be declared according to the definition $A(n, a, \tau, \tilde{Z})$. We remark that the index \tilde{Z} is observable, easily constructed, and can be directly modeled using our proposed framework.

The stations Tempelhof, Potsdam and Brandenburg share very similar characteristics. The correlations for the respective time series of daily maximum temperature are all close to unity (0.9974 between Tempelhof and Brandenburg, 0.9960 between Potsdam and Brandenburg and 0.9964 between Potsdam and Tempelhof. The high correlations indicate little room for different heat wave

² There are more weather stations in the Berlin area but either without, or with insufficient, data on daily maximum temperature.

Table 3

Number of heat waves recorded at Tempelhof, Brandenburg and Potsdam stations 1957-01-01 to 2021-08-04 using the temperature index \tilde{Z}_t for different durations (n) and threshold temperatures (a).

n	a						
	30°	31°	32°	33°	34°	35°	36°
3	59	32	17	8	5	1	1
4	26	11	5	1	1	0	0
5	15	7	2	0	0	0	0
6	7	3	1	0	0	0	0
7	4	1	0	0	0	0	0
8	4	1	0	0	0	0	0
9	2	1	0	0	0	0	0
10	2	1	0	0	0	0	0
11	2	1	0	0	0	0	0

Table 4

Parameter estimates.

Parameter	Estimate	Std.err.	Parameter	Estimate	Std.err.
a_0	12.4812	0.0576	c_0	8.0073	0.0756
a_1	1.0354×10^{-4}	4.2301×10^{-6}	c_1	-2.3474	0.1122
b_1	-10.2725	0.0407	c_2	0.8368	0.1010
b_2	-3.9973	0.0408	c_3	0.4908	0.1017
λ_1	0.8110	0.0060	c_4	-0.2523	0.1025
λ_2	-0.1034	0.0078	d_1	0.6931	0.0420
λ_3	0.0366	0.0080	d_2	0.0597	0.0067
λ_4	0.0166	0.0080			
λ_5	-0.0029	0.0080			
λ_6	0.0264	0.0063			

behavior. However, there are differences in heat wave counts especially for shorter durations. In Table 3 we present the heat wave counts when the index $\tilde{Z}_t = \min(Z_t^1, Z_t^2, Z_t^3)$ is used. Using this index the criterion for a heat event is stricter and there will be fewer heat waves. The differences are larger for short duration heat waves. There are e.g. 59 heat waves of duration $n = 3$ and $a = 30^\circ$, compared to 68 for Tempelhof considered alone. Heat wave counts for longer durations and higher thresholds show smaller differences and are more similar to the ones in Table 2 recorded at Tempelhof station. The overall patterns also show similar characteristics. We conclude that differences are concentrated to shorter and less intense heat waves. It is possible that an index such as \tilde{Z}_t , or similar, that takes spatial considerations into account is relevant in certain insurance scenarios. This depends on how costs relate to different heat waves and how they extend in the spatial dimension.

4.2. Model estimation

Following Šaltytė-Benth and Benth (2012) we employ a step-wise procedure to estimate the model presented in Section 3.1 verifying crucial properties at each step. All parameter estimates for the full sample are presented in Table 4. First, we fit the seasonality and trend function in (9) to data. The order N_S of the Fourier series in (9) is set to $N_S = 1$ thus only including the first two trigonometric terms with coefficients b_1 and b_2 . We find that increasing the order to $N_S = 2$ thus including two additional terms gives no improvement; the coefficients b_3 and b_4 are very small and the resulting curve is very close to the one obtained from setting $N_S = 1$ (not reported). The parameter a_1 determines the linear trend and is estimated to $1.0354 \times 10^{-4} \approx 0.0001$ which corresponds to a 0.38° increase in the daily temperature maximum over a 10 year period and a total increase of 2.44° over our entire sample. This result is in line with results found in other studies. Šaltytė-Benth and Benth (2012) and Benth and Šaltytė-Benth (2013) both report an estimate of a_1 rounded to 0.0001 for average daily temperature data in Stockholm and several locations in Lithuania respectively. Campbell and Diebold (2005) also report significant trend estimates for a number of U.S. cities.

Turning to estimation of the AR-parameters of the process X_t in (10) we first need to determine an appropriate AR-order p . We estimate models of orders $p = 1, \dots, 8$ and compare the Akaike (AIC) and Bayesian (BIC) information criteria across models. The BIC has a distinct minimum for $p = 6$ and the AIC ceases to decrease also at $p = 6$ (not reported). These considerations therefore also points to $p = 6$ as the appropriate order. The AR-parameters λ_i , $i = 1, \dots, 6$ in (10) are estimated using a standard least squares approach. The parameter estimates of all AR-parameters, except λ_5 , are significant at 5% and satisfy the stationarity criteria for X_t .

In Fig. 2 we plot the ACF for the residuals u_t and u_t^2 obtained from fitting the AR-part. The ACF for u_t show no signs of seasonality or serial correlation while the ACF for u_t^2 indicates a clear yearly seasonal pattern. We fit the seasonal variance part $\sigma_{S,t}^2$ by quasi-maximum-likelihood assuming the conditional distribution for u_t is $N(0, \sigma_{S,t}^2)$ distributed. We find that truncation at

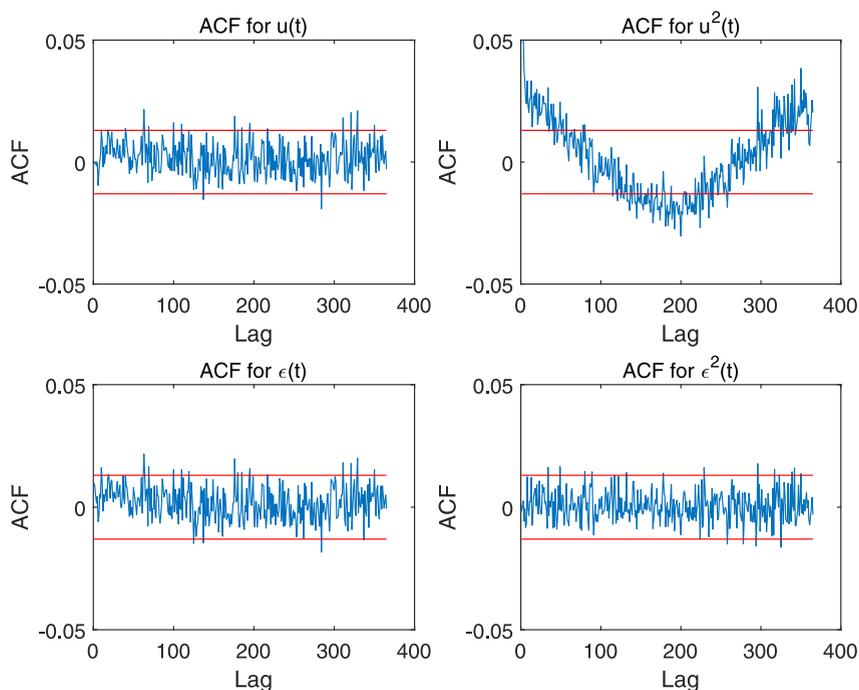


Fig. 2. Autocorrelation function (ACF) for u_t (upper left panel), u_t^2 (upper right panel), ε_t (lower left panel) and ε_t^2 (lower right panel).

$N_\sigma = 2$, i.e. including parameters c_0, c_1, c_2, c_3 and c_4 , is enough to describe the seasonally varying variance.³ The estimate of c_0 is 8.0073 which can be compared to the sample variance of u_t which is 8.0115. From the estimated parameters $c_i, i = 0, 1, 2, 3, 4$, we construct $\sigma_{S,t}$ in (13) and form the partly standardized residuals $v_t = u_t/\sigma_{S,t}$. We use these to estimate the GARCH component $\sigma_{G,t}^2$ in (14). The parameters d_1 and d_2 in (14) are also fitted using a quasi-maximum-likelihood procedure assuming v_t is $N(0, \sigma_{G,t}^2)$ distributed. Parameter estimates are positive, significant and fulfill the requirements of stationarity.

A final set of standardized residuals ε_t are given by $\varepsilon_t = u_t/\sigma_t$. The standardized residuals are assumed to be an i.i.d. sequence of random variables. In Fig. 2 we also plot the ACF for ε_t and ε_t^2 . The graphs show that all seasonality and serial correlation have been removed both from ε_t and ε_t^2 .

What is left is to evaluate the distributional fit. A standard assumption is that the ε_t are iid and $N(0,1)$ -distributed which requires no further estimation. In Fig. 3 we plot the normalized histogram of the standardized residuals together with the density for a standard Gaussian $N(0,1)$ random variable. We also plot the empirical distribution function of ε_t together with the cumulative distribution function for the $N(0,1)$ distribution. The plots indicate a satisfactory fit to the standard Gaussian $N(0,1)$ distribution. The Gaussian distribution function is also found to be a good fit for time series on daily average temperature in other studies, see e.g. Šaltytė-Benth and Benth (2012) and Benth et al. (2011). We also evaluated the fit from assuming ε_t are iid but instead following a standardized NIG distribution. These results are presented in Appendix. We found no strong motivation for using the NIG distribution. The empirical skewness and kurtosis of the standardized residuals are 0.0335 and 3.4203 which are close to the theoretical values of 0 and 3 for the $N(0,1)$ -distribution. The NIG distribution fits nicely to the observed skewness and kurtosis but the difference in calculated heat wave probabilities are very close to the ones obtained from using the standard Gaussian distribution. We therefore see no compelling reason for employing the NIG distribution for the data considered here.

The essential point of using this kind of reduced form statistical model is that it gives an accurate representation of both the conditional mean and variance in data and can be used for consistent contract valuation. The specification of the variance and underlying distribution are paramount to be able to calculate prices and risks. The occurrence of heat waves depends on many atmospheric variables which is not a feature of the statistical time series approach used for pricing. Prodhomme et al. (2021) study heat wave prediction using large scale meteorological models that take many atmospheric and physical variables into account. They find that heat waves cannot be predicted with any reliability for horizons longer than two weeks. We remark that in pricing applications it is common practice to calibrate some parameters to make the model consistent with other data sources such as market prices of related assets or even expert opinion. The parameters a_0 and a_1 can e.g. be used this way to calibrate the model to reflect heat wave predictions obtained from an advanced meteorological model. This is likely to be a minor concern since the

³ Šaltytė-Benth and Benth (2012) use a different estimation approach where they aggregate data for squared residuals for each day of the year and fit the truncated Fourier series in (13) to the aggregated curve. We also estimated the seasonal variance parameters using this method but find only very small differences in parameter estimates and that the conclusions regarding the order N_σ are the same for both methods.

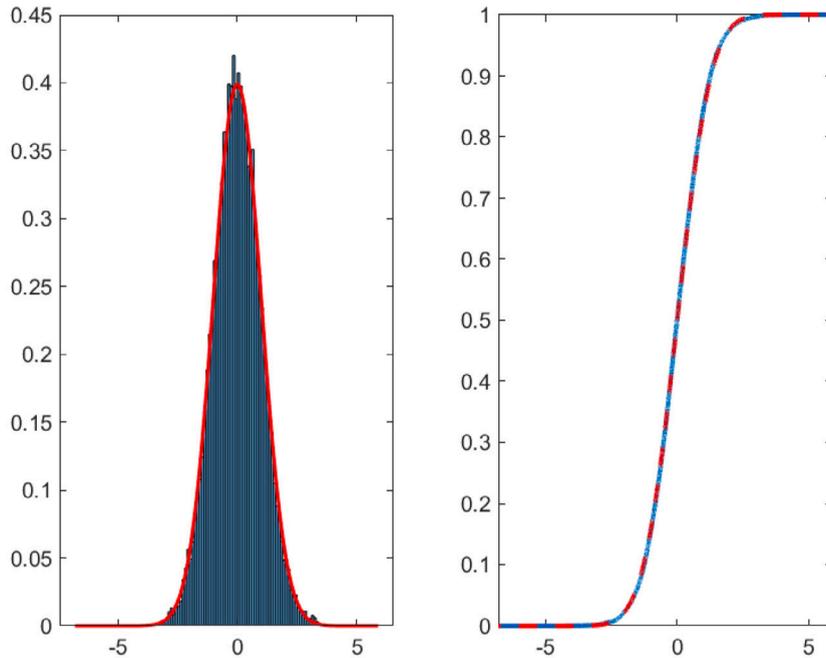


Fig. 3. Left panel: Normalized histogram of ϵ_i with probability density for the $N(0, 1)$ distribution (red curve). Right panel: empirical distribution function for ϵ_i with cumulative distribution function for the $N(0, 1)$ -distribution (red dashed curve). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

decision to enter a heat wave insurance contract should typically be made months before the start of the measurement period. At these long horizons heat waves are not predictable and probabilities of future waves are essentially determined by the expected seasonal behavior of the temperature during the measurement period which is well represented by the time series model used here.

4.3. Heat wave probabilities

We study heat wave probabilities calculated under the real world (or physical) probability measure \mathbb{P} using the estimated model for Berlin-Tempelhof station. These probabilities are of crucial interest since they will give insight into the risks of heat waves of different types actually coming to pass and that the corresponding insurance contracts will have to provide payment. They also determine the prices of the heat wave insurance contracts we have considered under the standard assumption that $\mathbb{P} = \mathbb{Q}$.

We focus on the heat wave probabilities associated with the contracts (2) and recall that these are given by

$$\xi^{\mathbb{P}}(t, n, a, \tau, Z) = \mathbb{P}(\Gamma(n, a, \tau, Z) \geq 1 | \mathcal{F}_t) \tag{20}$$

The basic scenario we consider is that of a prospective insurance buyer who contemplates buying heat wave insurance for the coming summer in Berlin 2022. The valuation date t is set to 2022-01-01 and we let the contract maturity be $T = \tau_b$ in all cases. The impact on the heat wave probabilities of the time point t is minor since the index Z_t mean reverts quickly to the seasonal function S_t . The conditional heat wave probabilities are determined by the projected behavior of Z_t during the measurement period and this is dominated by the trend and the seasonal mean and variance functions even when t is close to the start of the measurement period.

We consider four measurement periods corresponding to June ($\tau^{(1)}$), July ($\tau^{(2)}$), August ($\tau^{(3)}$) and the full summer period June to August ($\tau^{(4)}$), in 2022. In the corresponding contract specifications we always set τ_b to the end of the period plus $n - 1$ days. A June contract will e.g. then have τ_b equal to July 2 if the contract duration is $n = 3$. With this construction a heat wave of duration n that starts June 30 is still counted as belonging to the June contract measurement period.

The time series dynamics for Z_t are simulated forward to contract maturity using $M = 500,000$ daily sample paths. This provides M simulated observations $\Gamma_i(n, a, \tau, Z)$ that can be used to calculate the heat wave probabilities (20) through standard Monte Carlo integration.

Table 5 gives the heat wave probabilities $\xi^{\mathbb{P}}(n, a, \tau^{(4)}, Z)$ for $n \in \{3, 5, 7, 9, 11\}$ and $a \in \{30^\circ, 32^\circ, 34^\circ, 36^\circ, 38^\circ\}$ for the full summer period ($\tau^{(4)}$). The estimated probability of a 3-day long heat wave at $a = 30^\circ$ is as high as 85.06%. The probabilities tail off more quickly for increasing n than for increasing a which is consistent with the historical heat wave patterns observed in Table 2. At $a = 30^\circ$ the probability $\xi^{\mathbb{P}}(n, 30^\circ, \tau^{(4)}, Z)$ drops from 85.06% for $n = 3$ to 55.02% for $n = 5$ and then to 31.49% for $n = 7$. Keeping the duration fixed while increasing the threshold temperature the heat wave probabilities decay slightly slower. Several never before

Table 5

Heat wave probabilities $100 \times \xi^{\mathbb{P}}(n, a, \tau^{(4)}, Z)$ (%) for the summer measurement period ($\tau^{(4)}$) and different durations (n) and threshold temperatures (a). Probabilities are calculated by Monte Carlo simulation using 500,000 simulations. Corresponding standard errors are given in parenthesis. The valuation date is set to 2022-01-01.

n	$a = 30^\circ$	$a = 32^\circ$	$a = 34^\circ$	$a = 36^\circ$	$a = 38^\circ$
3	85.06 (0.0504)	62.73 (0.0684)	40.65 (0.0695)	24.87 (0.0611)	15.10 (0.0506)
5	55.02 (0.0704)	32.13 (0.066)	17.70 (0.054)	9.81 (0.0421)	5.59 (0.0325)
7	31.49 (0.0657)	16.03 (0.0519)	8.14 (0.0387)	4.25 (0.0285)	2.34 (0.0214)
9	17.79 (0.0541)	8.23 (0.0389)	3.94 (0.0275)	2.00 (0.0198)	1.07 (0.0146)
11	9.99 (0.0424)	4.32 (0.0288)	2.03 (0.0199)	1.01 (0.0141)	0.53 (0.0103)

Table 6

Heat wave probabilities $100 \times \xi^{\mathbb{P}}(n, a, \tau^{(1)}, Z)$ (%) for the June period ($\tau^{(1)}$) and different durations (n) and threshold temperatures (a). Probabilities are calculated by Monte Carlo simulation using 500,000 simulations. Corresponding standard errors are given in parenthesis. The valuation date is set to 2022-01-01.

n	$a = 30^\circ$	$a = 32^\circ$	$a = 34^\circ$	$a = 36^\circ$	$a = 38^\circ$
3	54.50 (0.0704)	34.28 (0.0671)	20.26 (0.0568)	11.8 (0.0456)	7.00 (0.0361)
5	29.02 (0.0642)	15.62 (0.0513)	8.32 (0.0391)	4.53 (0.0294)	2.59 (0.0225)
7	15.42 (0.0511)	7.56 (0.0374)	3.78 (0.027)	1.98 (0.0197)	1.09 (0.0147)
9	8.55 (0.0395)	3.87 (0.0273)	1.85 (0.019)	0.93 (0.0136)	0.51 (0.0100)
11	4.83 (0.0303)	2.06 (0.0201)	0.97 (0.0139)	0.48 (0.0098)	0.26 (0.0072)

Table 7

Heat wave probabilities $100 \times \xi^{\mathbb{P}}(n, a, \tau^{(2)}, Z)$ (%) for the July measurement period ($\tau^{(2)}$) and different durations (n) and threshold temperatures (a). Probabilities are calculated by Monte Carlo simulation using 500,000 simulations. Corresponding standard errors are given in parenthesis. The valuation date is set to 2022-01-01.

n	$a = 30^\circ$	$a = 32^\circ$	$a = 34^\circ$	$a = 36^\circ$	$a = 38^\circ$
3	59.15 (0.0695)	36.93 (0.0683)	21.28 (0.0579)	12.1 (0.0461)	7.00 (0.0361)
5	31.70 (0.0658)	16.61 (0.0526)	8.58 (0.0396)	4.53 (0.0294)	2.50 (0.0211)
7	16.66 (0.0527)	7.9 (0.0381)	3.83 (0.0271)	1.92 (0.0194)	1.03 (0.0143)
9	8.98 (0.0404)	3.95 (0.0275)	1.83 (0.019)	0.91 (0.0134)	0.48 (0.0097)
11	4.93 (0.0306)	2.03 (0.02)	0.92 (0.0135)	0.44 (0.0094)	0.22 (0.0067)

Table 8

Heat wave probabilities $100 \times \xi^{\mathbb{P}}(n, a, \tau^{(3)}, Z)$ (%) for the August measurement period ($\tau^{(3)}$) and different durations (n) and threshold temperatures (a). Probabilities are calculated by Monte Carlo simulation using 500,000 simulations. Corresponding standard errors are given in parenthesis. The valuation date is set to 2022-01-01.

n	$a = 30^\circ$	$a = 32^\circ$	$a = 34^\circ$	$a = 36^\circ$	$a = 38^\circ$
3	30.67 (0.0652)	16.7 (0.0527)	9.07 (0.0406)	5.08 (0.0311)	2.95 (0.0239)
5	12.98 (0.0475)	6.37 (0.0345)	3.24 (0.025)	1.76 (0.0186)	0.98 (0.0140)
7	5.85 (0.0332)	2.71 (0.0229)	1.35 (0.0163)	0.70 (0.0118)	0.39 (0.0089)
9	2.83 (0.0234)	1.26 (0.0158)	0.61 (0.011)	0.32 (0.0079)	0.17 (0.0059)
11	1.41 (0.0167)	0.62 (0.0111)	0.3 (0.0077)	0.15 (0.0055)	0.08 (0.0041)

observed heat waves have fairly high probabilities. For example, the probability of a 7-day long wave at 34° is 8.14% and for a 5-day wave at 36° it is 9.81%. We have also included a threshold temperature of $a = 38^\circ$ which is higher than ever observed for a heat wave in Berlin. The probability of a 3-day long heat wave with temperatures exceeding 38° each day is a non-negligible 15.10%. For $n = 5$ the 38° heat wave probability is 5.59% and for $n \geq 7$ it is less than 2.34%.

Tables 6–8 report probabilities $\xi^{\mathbb{P}}(n, a, \tau, Z)$ for the June, July and August contracts. An overall comment is that July have the highest probabilities and August have the lowest. For June and July the probabilities are close in magnitude while they are substantially lower for August. Heat waves occurrences are thus concentrated to the first two thirds of the summer period which is expected given the climate in Berlin. The behavior of the probabilities as functions of n and a is consistent with those found for the entire summer period. Probabilities of the more extreme heat waves are in some cases slightly higher for June than for July despite the fact that July is the warmer month with higher average temperatures. This can be explained by the fact that the seasonal variance is slightly higher in June than in July.

The results presented here give an indication of the scope for heat wave insurance with Berlin as example. As is evident from both the historical record of heat events and the model implied probabilities, short heat waves with 3–4 day durations and moderate threshold temperatures of $30^\circ - 32^\circ$ occur frequently and have high likelihoods. Long duration heat waves with $n > 5$ at 30° have a tendency to be less likely to occur compared to shorter waves with $n = 3$ and threshold temperatures $a > 30^\circ$. There are marked differences between probabilities calculated for different months with heat waves of all types being significantly less probable to

occur in August. The model implied probabilities further indicate that the risks of extreme heat waves of both higher temperatures and longer durations than observed historically are non-negligible.

5. Concluding remarks

We have presented a framework for constructing and pricing parametric heat wave insurance contracts. The contracts are based on a general heat wave definition formulated in terms of a single temperature index. The definition, and hence the contract types, can be varied by considering different durations, threshold temperatures, temperature index and measurement periods. We exemplify the framework by considering basic contracts that give payment based on the number of heat waves occurring in the measurement period. The associated contract prices are determined from the probabilities of the insured heat events. The framework is based on stochastic modeling of the temperature index. By modeling the relevant temperature index directly the model specification is independent of the heat wave contract definition. This makes it possible to price a variety of different contracts using the same model thus creating a flexible and practically appealing framework.

The framework we propose relies on a reduced form model for a single underlying index. This makes pricing flexible and straightforward. As long as a time series of the temperature index employed in the heat wave definition can be constructed this approach can be used. We supply the necessary discrete time results to perform pricing also without the standard assumption of a zero market price of temperature risk. We do not explore this question further. To be able to do that would require price data on the type of contracts that we consider which is not currently available. This situation is similar to many other studies in this field. An alternative is to study pricing in an equilibrium model which however is a major undertaking and requires an entirely separate study. In the empirical section we use data for Berlin and employ a relatively standard heat wave definition based on consecutive days of daily maximum temperature exceeding a certain threshold. The historical heat wave patterns using this definition reveal increases in frequency, duration and intensity of Berlin heat waves. We perform model estimation and calculate heat wave probabilities for a range of scenarios with different thresholds, durations and measurement periods. Since market data is not available we limit ourselves to the case of zero market price of temperature risk in the empirical part.

The present study is related to an emerging field of climate risk management with focus on heat waves. More research is needed in this area and the work presented here can be complemented and extended in many directions. More empirical evidence using data for different locations is of natural interest. We have limited ourselves to a case study of Berlin. Exploring different heat wave definitions may also prove relevant. Health related effects of heat waves are e.g. believed to be better captured by definitions based on average daily temperatures that also take nighttime temperatures into account. Ultimately the choice of heat wave definition depends on how it relates to costs. Different heat wave definitions are relevant for different types of costs. This is an area that require more study. It also requires data that can be used to link costs to heat wave events. Even though businesses and governments already monitor their exposure to heat waves they would benefit from methodological advances to improve their assessments. These challenging questions are left for future studies.

CRedit authorship contribution statement

Karl Larsson: Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Project administration, Resources, Software, Validation, Visualization, Writing – original draft, Writing – review & editing.

Declaration of competing interest

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

Data availability

Data will be made available on request.

Appendix

In this appendix we provide details on the NIG distribution including the Esscher transform, estimation results and simulated heat wave probabilities.

The Normal Inverse Gaussian (NIG) have been used in applications to temperature modeling, see e.g. [Ahčan \(2012\)](#) and [Benth and Šaltytė-Benth \(2005\)](#). The NIG distribution has the probability density

$$f(x) = \frac{\alpha\delta}{\pi} e^{\delta\gamma + \beta(x-\mu)} \frac{K_1(\alpha q(x))}{q(x)}, \quad x \in (-\infty, \infty) \quad (\text{A.1})$$

where $K_1(x)$ is the modified Bessel function of the second kind, $q(x) = \sqrt{\delta^2 + (x - \mu)^2}$ and $\gamma = \sqrt{\alpha^2 - \beta^2}$. We use the notation $X \sim \text{NIG}(\alpha, \beta, \mu, \delta)$ where the parameters must fulfill the restrictions $\alpha > 0$, $\delta > 0$ and $\alpha > |\beta|$. In the model (7) we may take the

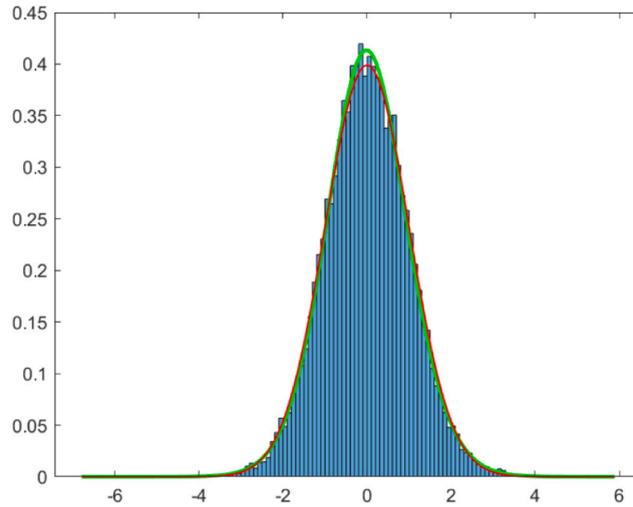


Fig. A.1. Normalized histogram of ε_t with probability density for the $N(0,1)$ distribution (red curve) and probability density function for the fitted standardized NIG distribution (green curve). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

stochastic sequence ε_t to be NIG-distributed but we also require that $E[\varepsilon_t] = 0$ and $\text{Var}(\varepsilon_t) = 1$. The first two moments of the NIG-distribution are $E[\varepsilon_t] = \mu + \beta\delta/\gamma$ and $\text{Var}(\varepsilon_t) = \delta\alpha^2/\gamma^3$. A normalized version is thus obtained from the parameter restrictions

$$\delta^* = \gamma^3/\alpha^2 \quad \text{and} \quad \mu^* = -\delta^*\beta/\gamma. \tag{A.2}$$

We hence assume that $\varepsilon_t \sim \text{NIG}(\alpha, \beta, \mu^*, \delta^*)$ which has mean zero and unit variance and only depends on the parameters α and β . The moment generating function for ε_t has the following simple form

$$\varphi_{\varepsilon_t}^{\varepsilon}(s) = e^{\mu^*s + \delta^*(\gamma - \sqrt{\alpha^2 - (\beta+s)^2})} \tag{A.3}$$

Using (A.3) and (19) we can therefore conclude that

$$\varphi_{\varepsilon_t}^Z(s; \mathbb{Q}) = e^{\eta_t s + \mu^* \sigma_t s + \delta^* \left(\gamma_t^{\mathbb{Q}} - \sqrt{\alpha^2 - (\beta_t^{\mathbb{Q}} + s \sigma_t)^2} \right)}$$

and hence that $Z_t | \mathcal{F}_{t-1} \sim \text{NIG}(\alpha/\sigma_t, \beta_t^{\mathbb{Q}}/\sigma_t, \eta_t + \sigma_t \mu^*, \delta^*/\sigma_t)$ where $\beta_t^{\mathbb{Q}} = \beta + \theta_t \sigma_t$ and $\gamma_t^{\mathbb{Q}} = \sqrt{\alpha^2 - (\beta_t^{\mathbb{Q}})^2}$. The NIG-distribution is preserved under linear transformation in the following manner: if $X_t \sim \text{NIG}(\alpha, \beta, \mu, \delta)$ then $Y = aX + b \sim \text{NIG}(\alpha/|a|, \beta/a, a\mu + b, |a|\delta)$. It then follows from standard calculations that the time series dynamics under \mathbb{Q} can be written

$$Z_t = \eta_t + \sigma_t \varepsilon_t^{\mathbb{Q}}$$

where $\varepsilon_t^{\mathbb{Q}} \sim \text{NIG}(\alpha, \beta_t^{\mathbb{Q}}, \mu^*, \delta^*)$. The distribution of $\varepsilon_t^{\mathbb{Q}}$ no longer have zero mean and unit variance. The first two conditional moments of Z_t under \mathbb{Q} are

$$E^{\mathbb{Q}} [Z_t | \mathcal{F}_{t-1}] = \eta_t + \delta^* \left(\frac{\beta_t^{\mathbb{Q}}}{\gamma_t^{\mathbb{Q}}} - \frac{\beta}{\gamma} \right) \quad \text{and} \quad \text{Var}^{\mathbb{Q}} (Z_t | \mathcal{F}_{t-1}) = \sigma_t^2 \left(\frac{\gamma}{\gamma_t^{\mathbb{Q}}} \right)^3$$

Hence the conditional mean is again augmented with a compensation for risk term but also the variance is changed. Even though the measure change only alters the β -parameter the time dependency that is introduced is much more complex than in the Gaussian case. Note that only if both θ_t and σ_t are constant will $\varepsilon_t^{\mathbb{Q}}$ constitute an iid sequence. The dependency of $\beta_t^{\mathbb{Q}}$ on $\theta_t \sigma_t$ will e.g. also make the conditional skewness and kurtosis of Z_t time dependent under \mathbb{Q} .

We evaluate the fit to the standardized residuals ε_t using the standardized NIG distribution. The parameters α and β are estimated using maximum likelihood. The estimates are $\hat{\alpha} = 3.1628$ and $\hat{\beta} = 0.1136$ with standard errors 0.2176 and 0.0652 respectively. The estimates are significant and implies a skewness and kurtosis of 0.0341 and 3.3022 which compares well to the empirical counterparts of 0.0335 and 3.4203. Fig. A.1 presents the histogram together with the densities of the standard Gaussian and standardized NIG distributions. Both distributions appear to fit the data well. To examine if using the NIG distribution instead of the Gaussian has any impact on heat wave contracts we calculate simulated heat wave probabilities. These are presented in Table A.1 and can be compared to the same set of probabilities calculated using the Gaussian $N(0,1)$ distribution presented in Table 5. The differences are very small even for the more extreme heat wave events considered and indicate that using the more complicated NIG distribution does not provide any marked advantage over the Gaussian.

Table A.1

Heat wave probabilities $100 \times \xi^{\text{H}}(n, a, \tau^{(4)}, Z)$ (%) for the summer measurement period ($\tau^{(4)}$) and different durations (n) and threshold temperatures (a). Probabilities are calculated by Monte Carlo simulation using 500,000 simulations based on the fitted standardized NIG distribution. Corresponding standard errors are given in parenthesis. The valuation date is set to 2022-01-01.

n	$a = 30^\circ$	$a = 32^\circ$	$a = 34^\circ$	$a = 36^\circ$	$a = 38^\circ$
3	85.00 (0.0505)	62.70 (0.0684)	40.65 (0.0695)	25.08 (0.0613)	15.35 (0.0510)
5	54.77 (0.0704)	31.93 (0.0659)	17.72 (0.0540)	9.87 (0.0422)	5.69 (0.0328)
7	31.46 (0.0657)	15.95 (0.0518)	8.11 (0.0386)	4.29 (0.0287)	2.39 (0.0216)
9	17.64 (0.0539)	8.21 (0.0388)	3.97 (0.0276)	2.03 (0.0199)	1.10 (0.0148)
11	9.99 (0.0422)	4.35 (0.0288)	2.01 (0.0199)	0.99 (0.0140)	0.53 (0.0103)

References

- Ahčan, A., 2012. Statistical analysis of model risk concerning temperature residuals and its impact on pricing weather derivatives. *Insurance Math. Econom.* 50 (1), 131–138. <http://dx.doi.org/10.1016/j.insmatheco.2011.10.005>.
- Alaton, P., Djehiche, B., Stillberger, D., 2002. On modelling and pricing weather derivatives. *Appl. Math. Finance* 9 (1), 1–20.
- Arnfield, A.J., 2003. Two decades of urban climate research: a review of turbulence, exchanges of energy and water, and the urban heat island. *Int. J. Climatol.* 23 (1), 1–26.
- Benth, F.E., Härdle, W.K., López Cabrera, B., 2011. Pricing Asian temperature risk. In: Cizek, P., Härdle, W., Weron, R. (Eds.), *Statistical Tools for Finance and Insurance*. Springer-Verlag, pp. 163–199.
- Benth, F.E., Šaltytė-Benth, J., 2005. Stochastic modelling of temperature variations with a view towards weather derivatives. *Appl. Math. Finance* 12 (1), 53–85.
- Benth, F.E., Šaltytė-Benth, J., 2013. *Modeling and Pricing in Financial Markets for Weather Derivatives*. World Scientific.
- Benth, F.E., Šaltytė-Benth, J., Koekebakker, S., 2008. *Stochastic Modelling of Electricity and Related Markets*. World Scientific.
- Björk, T., 2019. *Arbitrage Theory in Continuous Time*, fourth ed. Oxford University Press.
- Campbell, S.D., Diebold, F.X., 2005. Weather forecasting for weather derivatives. *J. Amer. Statist. Assoc.* 100 (469), 6–16.
- Chen, Y., Li, Y., 2017. An inter-comparison of three heat wave types in China during 1961–2010: Observed basic features and linear trends. *Sci. Rep.* 7 (45619), <http://dx.doi.org/10.1038/srep45619>.
- Chorro, C., Guégan, D., Ielpo, F., 2012. Option pricing for GARCH-type models with generalized hyperbolic innovations. *Quant. Finance* 12 (7), 1079–1094.
- Christoffersen, P., Jacobs, K., Ornathanal, C., 2012. Dynamic jump intensities and risk premia: evidence from S & P500 returns and options. *J. Financ. Econ.* 106 (3), 447–472.
- Fenner, D., Holtmann, A., Krug, A., Scherer, D., 2019. Heat waves in Berlin and Potsdam, Germany long-term trends and comparison of heat wave definitions from 1893 to 2017. *Int. J. Climatol.* 39, 2422–2437.
- Fenner, D., Meier, F., Scherer, D., Polze, A.A., 2014. Spatial and temporal air temperature variability in Berlin, Germany, during the years 2001–2010. *Urban Clim.* 10, 308–331. <http://dx.doi.org/10.1016/j.uclim.2014.02.004>.
- García-León, D., Casanueva, A., Standardi, G., Burgstall, A., Flouris, A.D., Nybo, L., 2021. Current and projected regional economic impacts of heatwaves in Europe. *Nature Commun.* 12, 5807. <http://dx.doi.org/10.1038/s41467-021-26050-z>.
- Gerber, H.U., Shiu, S.W., 1994. Option pricing by esscher transforms. *Trans. Soc. Actuar.* 46, 99–191.
- Goering, L., 2020. As Cities Bake on a Warming Planet, Insurers Cook Up Heatwave Cover. Reuters, Retrieved from <https://www.reuters.com/article/us-heatwave-insurance-climatechange-anal-idUSKBN26Q02Y>, last accessed 2021-08-24.
- Guégan, D., Ielpo, F., Lalaharison, H., 2013. Option pricing with discrete time jump processes. *J. Econom. Dynam. Control* 37, 2417–2445.
- Heston, S.L., Nandi, S., 2000. A closed-form GARCH option valuation. *Rev. Financ. Stud.* 13, 585–625.
- Lamm, T., Blumberg, L., Elkind, E.N., 2020. *Insuring Extreme Heat Risks*. Policy report, Berkley Law Center for Law, Energy and the Environment, December 2020. Available at: <https://www.law.berkeley.edu/wp-content/uploads/2020/11/Insuring-Extreme-Heat-Risks-Dec-2020.pdf>. Last accessed 2021-09-03.
- Matzarakis, A., Laschewski, G., Muthers, S., 2020. The heat health warning system in Germany application and warnings for 2005 to 2019. *Atmosphere* 11 (170), <http://dx.doi.org/10.3390/atmos11020170>.
- Ornathanal, C., 2014. Lévy jump risk: evidence from options and returns. *J. Financ. Econ.* 112, 69–90. <http://dx.doi.org/10.1016/j.jfineco.2013.11.009>.
- Prodhomme, C., Materia, S., Ardilouze, C., White, R.H., Batté, L., Guemas, V., Fragkoulidis, G., Garcia-Serrano, J., 2021. Seasonal prediction of European summer heatwaves. *Clim. Dynam.* <http://dx.doi.org/10.1007/s00382-021-05828-3>.
- Robinson, P.J., 2001. On the definition of a heat wave. *J. Appl. Meteorol.* 40, 762–775. [http://dx.doi.org/10.1175/1520-0450\(2001\)0400762:OTDOAH>2.0.CO;2](http://dx.doi.org/10.1175/1520-0450(2001)0400762:OTDOAH>2.0.CO;2).
- Šaltytė-Benth, J., Benth, F.E., 2012. A critical view on temperature modelling for application in weather derivatives markets. *Energy Econ.* 34, 592–602.
- Schiller, F., Seidler, G., Winner, M., 2012. Temperature models for pricing weather derivatives. *Quant. Finance* 12 (3), 489–500.
- Smith, T., Zaitchik, B., Gohlke, J., 2013. Heat waves in the United States: definitions, patterns and trends. *Clim. Change* 118, 811–825. <http://dx.doi.org/10.1007/s10584-012-0659-2>.
- Smoyer-Tomic, K.E., Kuhn, R., Hudson, A., 2003. Heat wave hazards: an overview of heat wave impacts in Canada. *Nat. Hazards* 28, 465–486. <http://dx.doi.org/10.1023/a:1022946528157>.
- Stewart, I.D., 2010. A systematic review and scientific critique of methodology in modern urban heat island literature. *Int. J. Climatol.* 31 (2), 200–217.
2021. *The Atlas of Mortality and Economic Losses from Weather, Climate and Water Extremes (1970-2019)*. World Meteorological Organization.