Incorporating Climate Change Projections into Risk Measures of Index-Based Insurance

Zhuoli Jin and Robert J. Erhardt

Department of Mathematics and Statistics, Wake Forest University

jinz16@wfu.edu, erhardrj@wfu.edu

February 27, 2018

Abstract

We present a complete working example of using regional climate model projections to estimate the changing risks of temperature index-based insurance products defined for a series of locations in California. This region is a major agricultural producer for the US and the world. The climate model projections are an ensemble of six regional climate model runs obtained from the North American Regional Climate Change Assessment Program. Hindcasts for the period of 1971-2000 are compared to historical observed temperature data for bias and variance corrections. Adjusted future model projections are used to estimate distributions of cooling degree days for 2041-2070, which are then used to estimate risk measures for index-based insurance products defined for cooling degree day indices. More broadly, this paper provides a transparent illustration of climate data processing, climate model bias correction, and the use of climate data products to explore actuarial risks.

Keywords: climate risk, NARCCAP, regional climate model, weather derivative

1 Introduction

Index-based insurance is a non-traditional form of insurance which specifies payments to the policyholder based on the settlement value of some external index. The contract specifies the index, the mathematical function which describes payments for all possible settlement values of the index, and the time period over which the insurance is in force. Whereas traditional insurance compensates policyholders based on the individual actual dollar loss realization from a covered peril, index-based insurance compensates policyholders with an amount based on an external trigger, regardless of the realized actual loss of any policyholder. As an example, consider the Livelihood Protection Plan from the Munich Climate Insurance Initiative (MCII, 2013). This index-based insurance in the Caribbean specified one type of payment to policyholders triggered if the total rainfall exceeded some high threshold. The motivation was to issue payments for severe weather events likely to cause physical damage and economic disruption, rather than issuing payments for confirmed losses. Rainfall was monitored by the Danish Hydrological Institute. Thus in this case, the index and random variable was total rainfall, and payments were made depending on the settlement value of this index.

The advantages of this approach include reduced administrative costs for the insurer – moral and morale hazards are virtually eliminated as the index is objectively verified and beyond anyone's control or manipulation, and there is no need to confirm the actual policyholder losses – and payments can be issued remotely and quickly. An obvious downside is the potential mismatch between actual losses and triggered payments (i.e. payments triggered when no true loss occurred, or true losses incurred without triggering index values leading to payments). Despite this limitation, index-based insurance has been widely touted in developing regions underserved by traditional insurance markets; see, for example, the review paper by Collier et al. (2009), and examples from Barnett and Mahul (2007) or Linnerooth-Bayer and Mechler (2006). Other index-based insurance products have been defined with

triggers based on precipitation extremes (Fischer et al., 2012), the normalized vegetation depth index (a satellite-based measurement of vegetation) which correlates with crop risk (Turvey and Mclaurin, 2012), temperature (Turvey, 2005), a temperature-humidity index designed to correlate with dairy production risk (Deng et al., 2007), and more complex indices meant to capture livestock risk (Chantarat et al., 2013). Because payments from the insurer to the insured are deterministic for a given index settlement value, obtaining a distribution of the settlement index value is equivalent to obtaining a distribution for the payments. Therefore, actuarial pricing is made possible by a distribution of the index (Jewson and Brix, 2005).

This paper will explore the actuarial analysis of a temperature index-based insurance product in California. That is, the index which determines payments will be based only on daily temperature values, and described later. For now, think only of some necessary pieces of data needed to price such a product. Historical data on daily temperature can be obtained from publicly available sources (e.g. the National Centers for Environmental Information (NCEI)), and these data could be used to estimate the distribution of possible settlement values of the temperature index. When using a few decades worth of historical temperature data, one must ask (Q1) if the distribution of temperatures has been *stationary* over this time period – that is, can we regard all historical data as being draws from the *same* distribution of daily temperatures, or is it possible that the distribution has been changing over time? Furthermore, if the index-based insurance produce is defined for some period far enough in the future, one must also ask (Q2) if the distribution in the future will match the distribution in the present, or if the climate is presently changing and further adjustments to the future date will be needed.

Historical data show that for many weather variables, distributions have not been stationary over the past few decades. Consider Figure 1 below, which shows global mean temperature anomalies temperature from 1880 - present. A temperature anomaly is a year's mean temperature minus the average taken over 1900 - 2000. As is immediately evident, average

temperatures have been trending up over the past few decades, and therefore the mean of the distribution of of temperature has been trending up. For an excellent visualization of shifting temperature distributions, see Hansen et al. (2010) and Hansen et al. (2012). The Actuaries' Climate Index (ACI, `http://actuariesclimateindex.or/home/`) similarly looks at historical data to identify shifts in climate. Revisiting Q1 from the previous paragraph, we conclude that the distribution of global mean temperatures has not been stationary, and therefore some adjustment may be needed if one uses temperature spanning back several decades. Some sources discussing climate non-stationarity or climate forecasts and insurance include Mills (2005), Chang et al. (2012), Carriquiry and Osgood (2012), and Erhardt (2017).

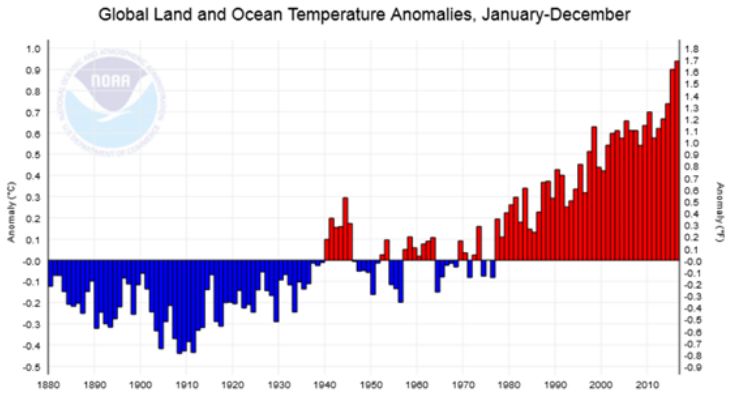

Figure 1: Global average temperature anomalies (land and ocean) from 1880-2016, obtained from the National Centers for Environmental Information (https://www.ncdc.noaa.gov/cag/). Anomalies are defined as a year's mean surface temperature minus the average temperature taken from 1900 - 2000. Years shown in red are warmer than the 20th average, whereas years shown in blue are cooler.

Investigation of Q2 is the primary focus of this paper. To ask question Q2 is to ask what scientists project forward for climate trends, and how those projections could be utilized by actuaries. If, for example, temperatures are projected to increase at a slight but non-

negligable rate each year, than an actuarial correction is needed to price any temperature index-based product in the future. Numerous scientific agencies build models which project that temperatures should continue to increase in the near future in response to anthropogenic carbon emissions and other feedbacks – see for example Pachauri et al. (2014), any Intergovernmental Panel on Climate Change (IPCC, `http://www.ipcc.ch/`) assessment report, or Erhardt and Von Burg (2018). Part of the evidence for this is captured by climate models, which are computer simulations of the Earth's climate, run for both the recent past (for verification) and for the near future, under different carbon emissions scenarios. Their output can be used to determine future climates in response to different carbon scenarios. This paper will consider a set of projections from the National Center for Atmospheric Research (NCAR), the Canadian Centre for Climate Modelling and Analysis, and the Hadley Centre in the UK.

We will use an ensemble of six regional climate models from the North American Climate Change Assessment Program (Mearns et al., 2007) to estimate changing risks of temperature index-based insurance products in central California. This index will be based on cooling degree days, an overall measure of the warmth of a time period, and global climate models will show projected changes to this index due to warmer future temperatures. We first describe how to obtain, visualize, and manipulate NARCCAP regional climate model output. We demonstrate the need for bias and variance correction by incorporating historical observations, and show how these corrections are applied to future climate model projections to estimate future distributions of the temperature index values. As this index serves as the only random variable for temperature index-based insurance, we demonstrate how to estimate a few common risk measures and quantify how these risks are projected to increase as a result of climate change. This paper therefore can serve as a guide to those who wish to explore the performance of index-based insurance products under future climates as described by climate model output.

City	% missing days	Mean	Std Dev	Skewness	Kurtosis
Bakersfield	0.00	65.5	14.0	0.05	1.99
Fresno	0.00	63.8	13.9	0.03	1.93
Los Angeles	0.00	63.3	6.4	0.09	2.82
Sacramento	0.18	61.1	11.9	-0.02	2.09

Table 1: Summary statistics for the 4 locations in this study.

2 Exploratory Data Analysis

2.1 Historical Data

We obtained historical daily temperature for the cities Sacramento, Bakersfield, Fresno, and Los Angeles in California for the period Jan 1^{st}, 1971 to Dec 31^{st}, 2000. These locations are shown as SA, BA, FR and LA in Figures 2 and 3. The selected cities are mainly located in Central Valley, one of the most productive agricultural regions in the United States. This time period was selected to match the climate model hindcasts from our ensemble of regional climate models, and thus allowed us compare regional climate model output to historical data for bias and variance corrections (described in next section). These data were obtained from the National Centers for Environmental Information (NCEI) Climate Data Online (CDO, `https://www.ncdc.noaa.gov/cdo-web/`). This portal provides free access to the National Climate Data Center (NCDC) archive of global historical weather and climate data in addition to station history information. Data include daily, monthly, seasonal and yearly measurements of weather, such as temperature, precipitation, wind, etc, and are available in various formats. Summary statistics for the four cities are provided in Table 1.

2.2 Regional Climate Model Data

Regional climate model (RCM) output comes from the North American Regional Climate Change Assessment Program (NARCCAP, `https://www.narccap.ucar.edu/`). An introduction and overview of this program can be found in Mearns et al. (2005) and Mearns et al. (2007). NARCCAP is an international program to produce climate projection outputs for North America in the year 2041 to 2070 from multiple regional climate models (RCMs) nested within multiple general circulation models (GCMs). Background about RCMs and GCMs for an actuarial purpose can be found in Erhardt and Von Burg (2018). What is relevant for this paper is that RCMs are climate models designed to allow for the study of more localized, regional consequences of projected climate change. Climate model output from NARCCAP is available for two time periods, a hindcast period from 1971 - 2000 and a future period from 2041-2070. We obtained both hindcast and future data on temperature for a rectangular region encompassing California, as shown in Figures 2 and 3. The hindcast is compared to historical observed data for verification, and the future period is used to explore a future climate.

Each climate model parameterizes the climate variables in slightly different ways. To incorporate model uncertainty, we chose an ensemble of six different GCM/RCM combinations to explore a range of models. This ensemble contains (1) the Canadian Regional Climate Model nested within the Community Climate System Model (CRCM-CCSM); the (2) Canadian Regional Climate Model nested within the Third Generation Coupled Global Climate Model (CRCM CGCM3); the (3) Weather Forecasting Research Group model nested within the Community Climate System Model (WRFG-CCSM); the (4) Weather Forecasting Research Group model nested within the Third Generation Coupled Global Climate Model (WRFG-CGCM3); the (5) the Hadley Regional Model 3 nested within the Hadley Centre Coupled Model, version 3 (HRM3-HadCM3); and the (6) PSU/NCAR mesoscale model nested within the Hadley Centre Coupled Model, version 3 (MM5I-HadCM3). All six GCMs were run under

the SRES A2 emissions scenario (see http://www.narccap.ucar.edu/about/emissions.html or Nakicenovic et al. (2000) for more detail). We will highlight results from the CRCM-CCSM model run to illustrate the data (Gent et al., 2011), but the others are qualitatively similar.

In the CRCM-CCSM model, every grid cell has a longitude and latitude, marked as two 140 by 115 matrices $lon[xc, yc]$ and $lat[xc, yc]$, where xc and yc are coordinates of the projection. Each pair of (xc, yc) fixes a grid cell in the climate model. We extracted the grid cells covering California for illustrative purposes (see figure 2). The output for one year of data had dimension $16 \times 33 \times 2920$, where the first two dimensions are longitude and latitude and the third is time, with a measurement every 3 hours for 365 days (all instances of February 29 are dropped).

2.3 Cooling Degree Days

The index we consider for the index-based insurance product in this paper is based on cooling degree days (CDDs), which measure the excess temperature above a certain threshold. For a single day d, a CDD is defined as

$$CDD_d = max(0, \bar{T}_d - 65),$$

where $\bar{T}_d$ is the average of daily maximum temperature and minimum temperature measured in degrees Fahrenheit. We sum CDDs over some time period denoted by $\mathcal{D}$, and call this sum the cumulative cooling degree days (cCDDs) for a place/location in time $\mathcal{D}$, defined as

$$\begin{aligned} cCDD &= \sum_{d \in D} CDD_d \\ &= \sum_{d \in D} max(0, \bar{T}_d - 65). \end{aligned} \tag{1}$$

cCDDs are a measure of overall temperatures in excess of 65 degrees Fahrenheit, without

regard to precisely how or when the temperature exceeded the threshold. They serve as a convenient proxy for the overall warmth of a season, and are widely used by utility companies to measure demand for heating and cooling energy, and by agricultural producers to measure the warmth of a particular season. Erhardt (2015) explored projections of trends in cCDDs in the mid-twenty-first-century using NARCCAP data. Here we will explore changes in the overall distribution of cCDDs. Following the example in Fleege et al. (2004) to correspond with agricultural risk, we define $\mathcal{D}$ as the period from May 1 to July 31, for both current and future years. We computed the annual cCDDs for each grid cell in Figure 2 by summing the daily CDDs over the period of May 1 through July 31. This figure shows the average of these cCDDs taken over the 30 year current period (left panel) and future period (center panel). Owing to warmer projected temperatures in the future, the future period shows higher average cCDDs, best seen by the third image (right panel) which shows the difference between the other two panels. For every grid cell in the study region, the cCDDs are projected to increase in the CRCM-CCSM model. This projected increase is largest in the southeast and central valley of CA, with only mild increases projected for the Pacific coast.

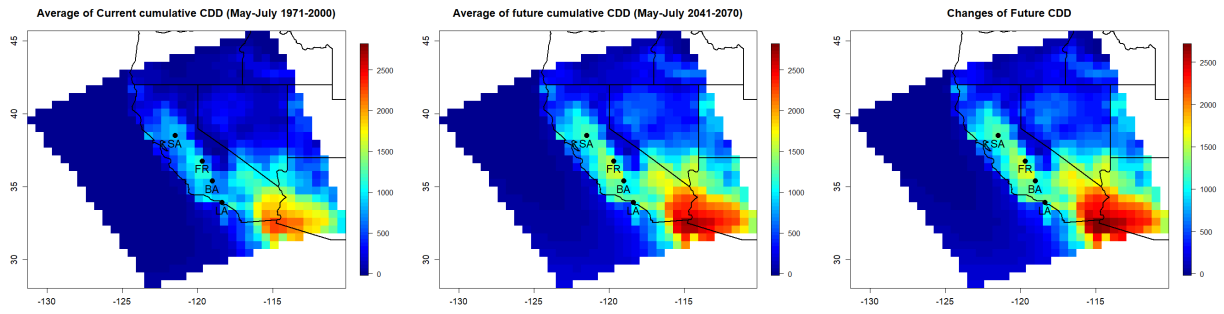

Figure 2: The left and center panels are average cCDDs for the current (1971-2000) and future (2041-2070) time periods for California, respectively. The third panel shows the difference in mean cCDDs between these two time periods (future cCDDs minus current cCDDs).

Figure 3 shows a pointwise 95% confidence surface for the projected increase in cCDDs, with the lower bound (left panel) and upper bound (right panel) computed simply as the difference $\pm$ 2 standard errors of the difference. Even after incorporating uncertainty in the two estimates, the lower bound is everywhere greater than zero.

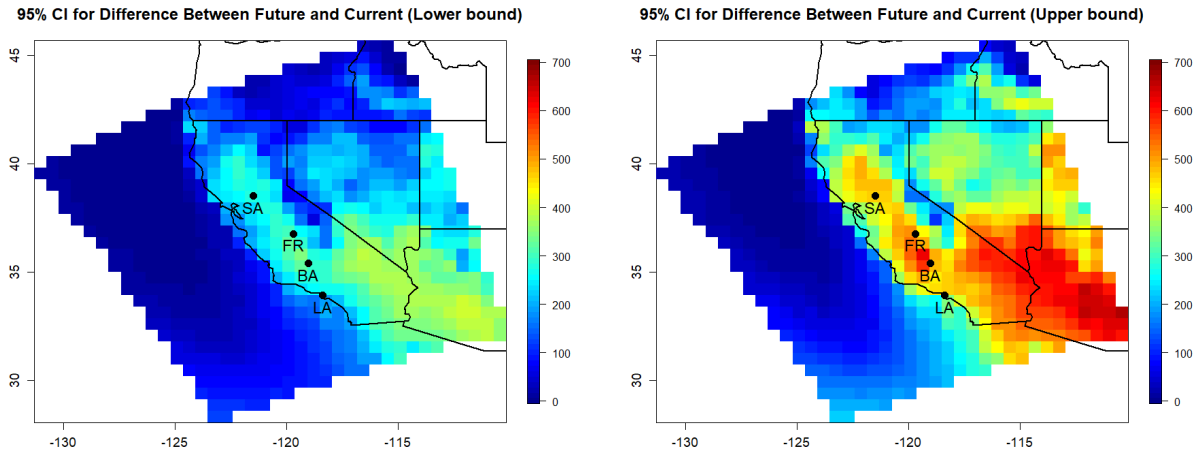

Figure 3: These two panels show pointwise 95% confidence surface for the projected change in cCDDs from the CCSM-CRCM model over California. The lower bound (left panel) is computed as the mean minus 2 standard errors, and the upper bound (right panel) is computed as the mean plus 2 standard errors.

3 Downscaling

3.1 Introduction to Statistical Downscaling

Regional climate model output may be biased relative to observed climate (Christensen et al., 2008). To help remove some of this bias, we compare RCM output for the past (1971-2000) to observed historical data temperature data (also 1971-2000) and compare the means

and standard deviations of the distributions of daily data from these two sources. From the differences in means and standard deviations, we can define and then apply bias and variance corrections to output from the RCM in the future period, to make a statistically downscaled adjustment which corrects for bias in the mean and variance. Here we describe this procedure.

For each of the four cities, we use a distribution of historical daily temperatures and a distribution of current regional climate model output for the grid cell containing the city. Since there has been mild warming over the period 1971 - 2000, we re-state all data to the year 2000. Similarly, for the future time period 2041 - 2070, we restate all temperatures to the year 2070 (ignoring this re-stating step would lead to slightly inflated estimates of the standard deviation of daily temperatures, as the distribution of daily temperatures has been shifting to the right over time). These two restating procedures are described in Algorithm 1.

Algorithm 1 Re-stating

Variables:
T: temperature before re-stating
Y_t: year t (in decimal form, e.g. 1975.7)
Y: base year (2000 or 2070)
$\tilde{T}$: re-stated temperature
$\hat{\beta}$: linear coefficient
$\hat{\epsilon}$: fitting error

1: Fit linear models:
$T = \beta Y_t + \epsilon$, estimate $\hat{\beta}$
2: $\tilde{T} = T + \hat{\beta}(Y - Y_t)$

Next, we have an algorithm to compute biases in the mean and variance between historical observed temperatures from 1971-2000 and RCM output over the same period. First, we re-state both historical temperatures as well as RCM output to year 2000 levels using Algorithm 1. Then we downscale RCM output while correcting for bias in the mean and variance. This computation is described in Algorithm 2. Finally, the downscaled and bias-corrected

RCM output is re-stated to the proper year by running what is essentially the inverse of Algorithm 1.

Algorithm 2 Downscaling and Bias Correction

Variables:
$T_{historical,t}$ Re-stated historical temperature in year t
$T_{current,t}$ Re-stated current temperature in year t
$T_{future,t}$ Re-stated future temperature in year t
$\bar{T}_{historical}$ Average of re-stated historical temperature
$\bar{T}_{current}$ Average of re-stated current temperature
$\bar{T}_{future}$ Average of re-stated future temperature
$\mathbb{T}_{current,t}$ Downscaled current temperature before re-trending
$\mathbb{T}_{future,t}$ Downscaled future temperature before re-trending
$\sigma_{historical}$ Standard deviation of re-stated historical temperature
$\sigma_{current}$ Standard deviation of re-stated current temperature
ψ Bias correction
ϕ Variance correction

1: Compute $\psi = \bar{T}_{current} - \bar{T}_{historical}$
$\phi = \sigma_{historical}/\sigma_{current}$
2: $\mathbb{T}_{current,t} = (T_{current,t} - \bar{T}_{current}) \times \phi + \bar{T}_{current} - \psi$,
$\mathbb{T}_{future,t} = (T_{future,t} - \bar{T}_{future}) \times \phi + \bar{T}_{future} - \psi$.

3.2 Results of Downscaling

Here we demonstrate both the need for bias and variance correcting downscaling, along with the impacts of these algorithms. To begin, Figure 4 shows raw data for historical temperature (solid line), current RCM output (dashed line) and future RCM output (dot-dashed line) for the four cities under consideration. It is immediately clear that the distributions of future temperatures have shifted to the right from current temperatures; however, it is also clear that the distributions of RCM data for 1971-2000 does not match the distribution of observed temperatures. The discrepancy between then solid and dashed lines indicate a need for downscaling bias-correction. This same downscaling correction then should be applied

to the future distributions of temperature as well.

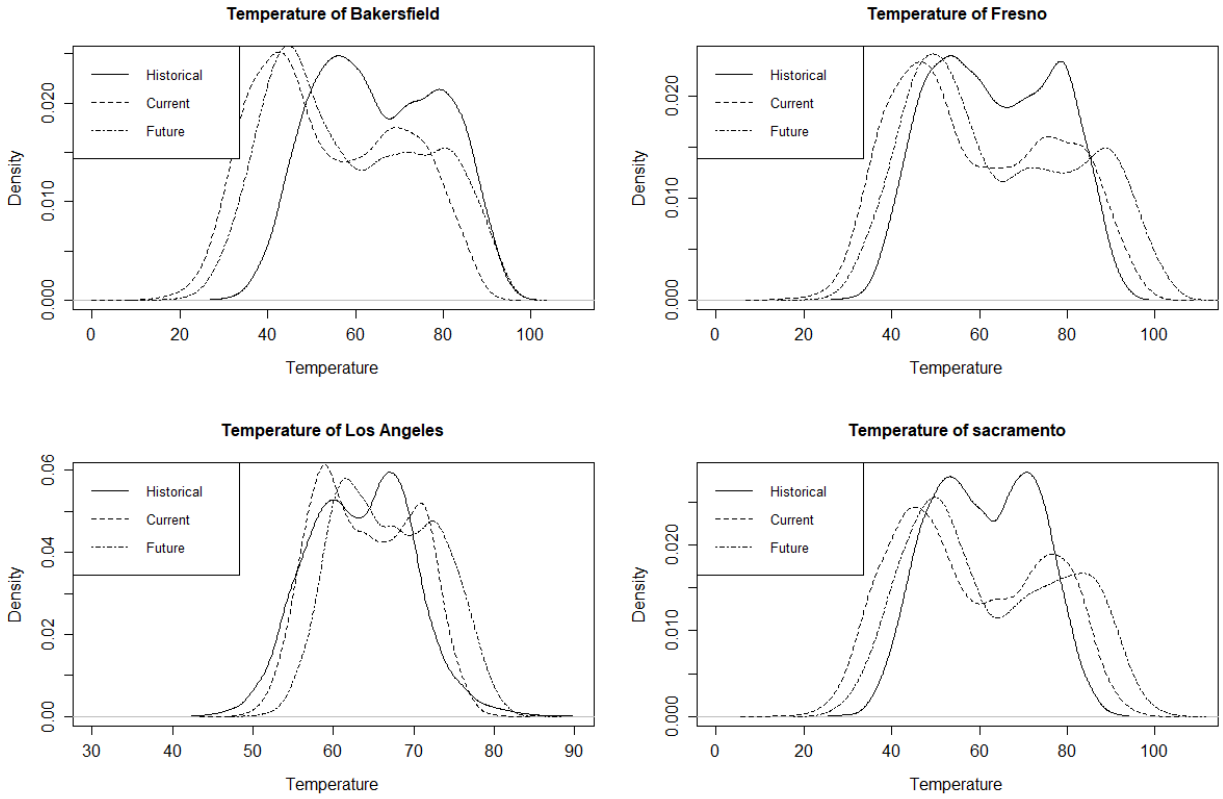

Figure 4: Density plots for historical (solid line), current (dashed line), and future (dot-dashed line) temperature data before downscaling.

Figure 5 shows the impact of the downscaling procedure. The left column shows historical data (dot-dash) along with RCM data both before (dashed) and after (solid) the downscaling procedure. Observe that the downscaling shifts the biased dashed distribution to match observed in both mean and variance. The right column of figures show the same downscaling procedure applied to future RCM output which shifts the dashed curve to the solid, downscaled curve; for comparison purposes only, historical data is once again shown (dot-dash) which highlights that even after downscaling is applied, future regional climate model output

show a shifted distribution of temperatures from a warmer climate. Finally, Figure 6 shows just the shift from downscaled current to downscaled future, to make clear the impact of climate change from the period of 1971-2000 to 2041-2070 under this climate model projection. As a reminder, all of these figures described are only for the CRCM-CCSM climate model, and the same exact procedure is applied to the other five RCMs described earlier. From equation 1, we calculate and compare annual cCDDs for each city in both current and future time periods, shown in Figure 7. Exploring the impact of changing CDDs is the subject of the Application in the next section.

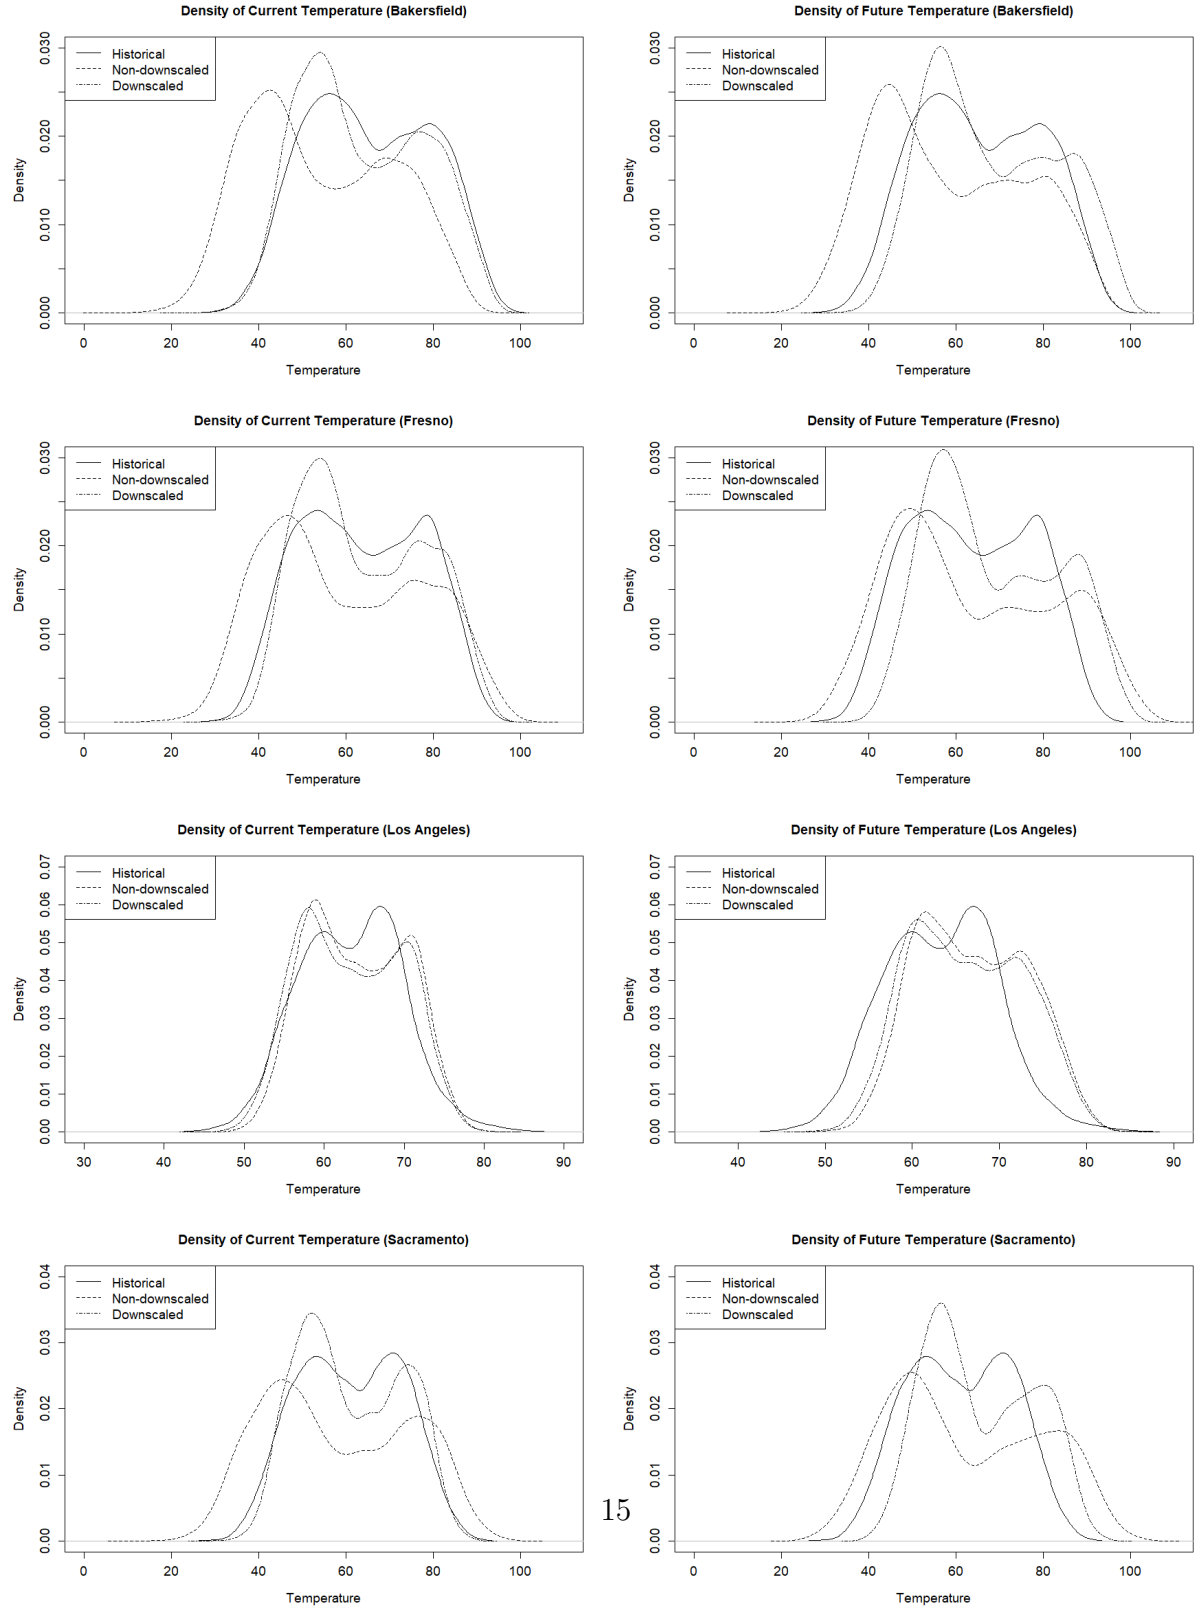

Figure 5: Historical temperatures and current/future temperatures for BA, FR, LA and SA. From the plot, we can see downscaled temperature has a more similar shape to historical temperature.

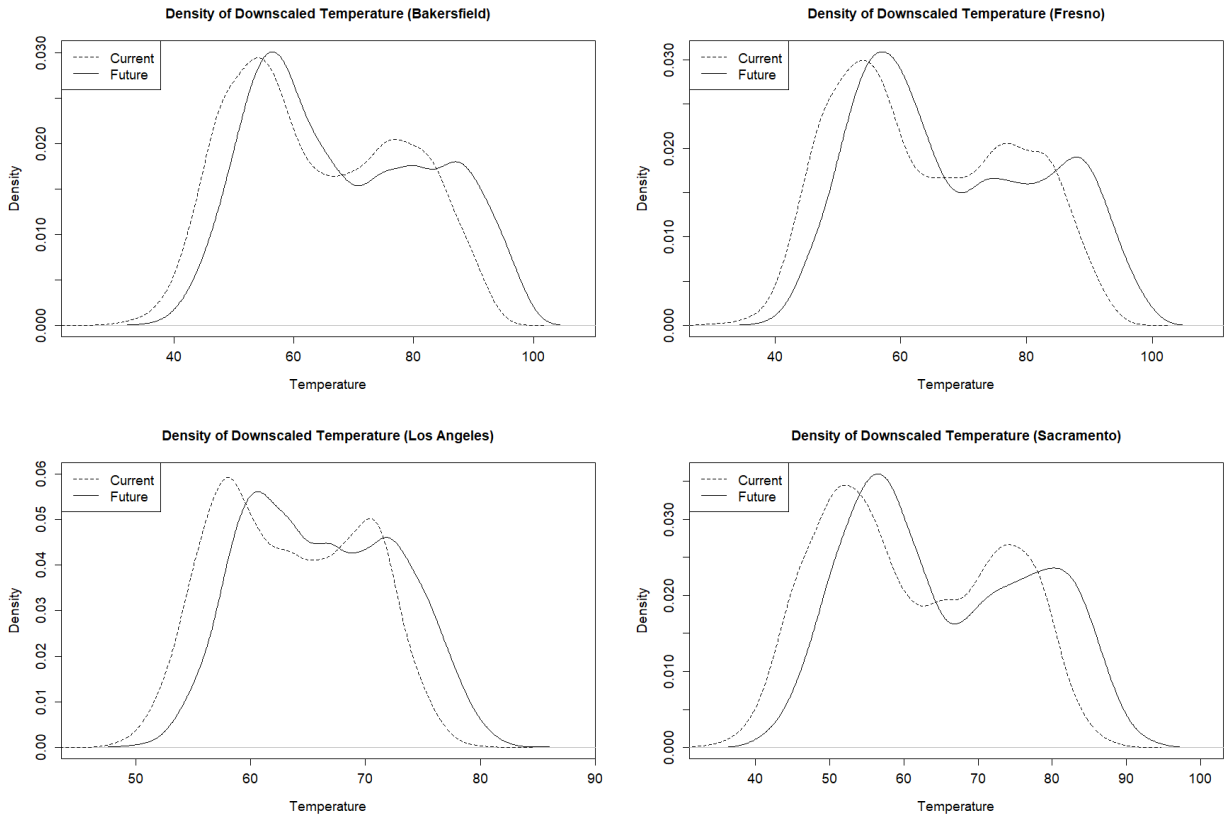

Figure 6: Downscaled current and future temperatures for BA, FR, LA and SA.

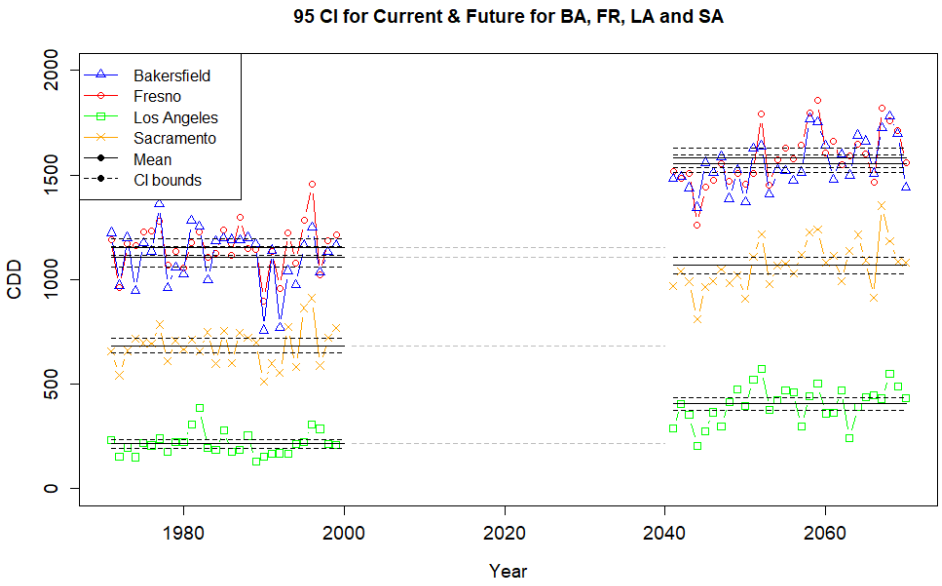

Figure 7: The triangles, circles, squares and crosses represent BA, FR, LA and SA, respectively. Current downscaled temperature are shown on the left, and future downscaled temperature are shown on the right. After accounting for annual variability, the distribution of downscaled future cCDDs is noticeably shifted from the current downscaled distribution.

4 Application

A CDD index can be used to price weather derivatives. Zeng (2000) introduces basic concept of call, put and swap contracts. Alaton et al. (2002) prices a heating degree day (HDD) call option, and Deng et al. (2007) prices a temperature-humidity index (THI) call option. In this part, a CDD call option is designed, and analyzed. Define a call option that makes a payment P conditional on the realization of the cCDDs over some period of interest $\mathcal{D}$ according to the following:

$$P(cCDD \mid cCDD_{threshold}, k) = \begin{cases} 0 & cCDD < cCDD_{threshold}, \\ k(cCDD - cCDD_{threshold}) & cCDD_{threshold} \leq cCDD, \end{cases} \tag{2}$$

where P is the payment, which is a function of cCDDs given the threshold $cCDD_{threshold}$. When $cCDDs$ exceed $cCDD_{threshold}$, this option triggers an indemnity. Thus, this option can be uniquely described by the parameter $cCDD_{threshold}$.

The premium on this CDD call option is a function of $cCDD_{threshold}$ and k, and the probability distribution of $cCDDs$. Deng et al. (2007) describes kernel density estimation for this probability distribution. In this paper, kernel smoothing is used to derive a probability density function $f(cCDD)$ from a collection of individual realizations of cCDDs. Given T realizations $cCDD_1, ..., cCDD_T$, the kernel density function of $cCDD$ is

$$f(cCDD) = \frac{1}{T\Delta} \sum_{t=1}^{T} K\left(\frac{cCDD - cCDD_t}{\Delta}\right), \tag{3}$$

where $K()$ is a normal density kernel function, and Δ is the bandwidth.

4.1 Risk Measures

The Value-at-Risk (VaR) is a measure of the upper tail of payments, defined simply as the p^{th} percentile

$$VaR_p = F_P^{-1}(p) \tag{4}$$

where F_P is the distribution function of random variable P. Typically, p is commonly chosen as a high value such as $p = 0.95$ and $p = 0.99$, which makes the VaR the amount exceeded in only 5% and 1% of cases, respectively. The conditional tail expectation (CTE) quantifies the expected value of the loss given that the loss has exceeded a high percentile. Mathematically, it is

$$CTE_p = E(P|P > F_P^{-1}(p)), \tag{5}$$

with all notation the same as above.

4.2 Simulation-based distributions of cCDDs and Payments

This subsection describes a statistical model fit to historical temperature data, and this model then serves as a generator for realizations of possible temperature outcomes under the future climate. These realizations are used to compute cCDDs and then payments for two hypothetical temperature derivatives. We demonstrate this model fit for Fresno, CA, but the same process is used for all cities.

We first fit a linear model for temperature using data from January 1st, 1971 to April 30th, 2000 with periodic terms and ARMA(3,3) residuals:

$$T_d = c_0 + c_1 d + \sum_{r=1}^{3} \alpha_r \sin\left(\frac{2\pi \cdot r \cdot d}{365}\right) + \sum_{r=1}^{3} \beta_r \cos\left(\frac{2\pi \cdot r \cdot d}{365}\right) + w_d,$$

where T_d is temperature for day d, $c_0, c_1, \alpha_1, \alpha_2, \alpha_3, \beta_1, \beta_2, \beta_3$ are coefficients, and error term $w_d \sim ARMA(3,3)$ with Gaussian errors. This model was chosen as it minimized BIC

when compared to similar models with other orders for the number of Fourier terms and ARMA(p,q) structure. Using the fitted model, we were able to predict 92 days ahead, i.e., daily temperatures in May 1st to July 31st, and then we computed payments using equation 2. Setting $k = 1$ for simplicity, the simplified call option is then

$$P(CDD \mid CDD_{threshold}) = \begin{cases} 0 & cCDD < cCDD_{threshold}, \\ cCDD - cCDD_{threshold} & cCDD_{threshold} \leq cCDD, \end{cases} \tag{6}$$

To determine a reasonable $cCDD_{threshold}$, we first bootstraped predicting temperatures 10,000 times, and computed 10,000 cCDDs. Then, we computed the 75^{th} and 95^{th} percentile of cCDDs, and let them be two thresholds for two payments which we define as Payment #1 and Payment#2. In this way, we defined two call options for Fresno:

$$P1(CDD \mid 1117) = max(0, cCDD - 1117), \tag{7}$$

and

$$P2(CDD \mid 1132) = max(0, cCDD - 1132). \tag{8}$$

Then, for each of the computed 10,000 cCDDs, we calculated 10,000 payments, and compute the VaR and CTE risk measures with $p = 0.95$ and $p = 0.99$.

Table 4 shows the difference in downscaled future and downscaled current means. The downscaling parameters to get the downscaled temperature are shown in table 2. To predict payments under the future climate, we simply added these amounts to each daily temperature value for spring (first 31 days) and summer (next 62 days). Then, we can apply equation 7 and 8 to compute future payments, and equation 4 and 5 to compute risk measures.

Our aim is to combine payments from these six regional climate models into one overall distribution, whose variance reflects both model uncertainty (RCM) as well as temperature

uncertainty. To achieve this, we randomly select one of the six RCMs, and then randomly select one payment from within that model. Repeating 10,000 times, we construct a distribution of 10,000 payments drawn from across all six RCMs. Kernel density estimation (Equation 3) is used to draw density plots of payments for each city. Figure 8 shows results for Fresno, CA, using thresholds of the 75^{th} (left) and 95^{th} percentiles from cCDDs, and low bandwidth Δ (top) and high bandwidth Δ (bottom) for smoothing.

and adjust standard deviations in estimations to obtain more smooth curves, as shown in figure 8.

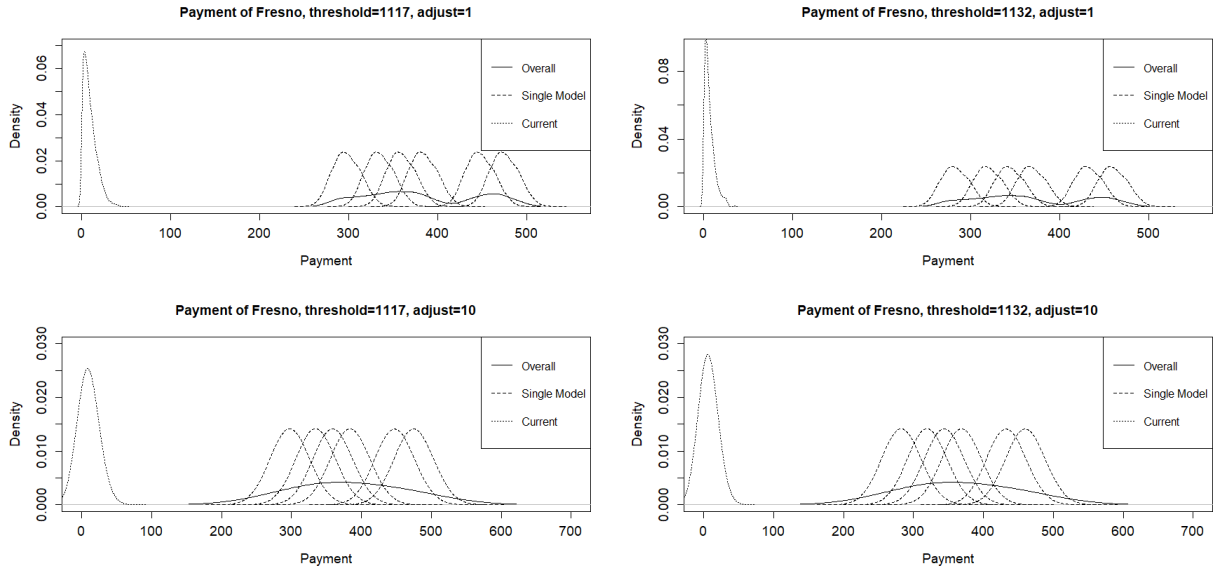

Figure 8: Density plots for future payments in Fresno. Left two panels shows payments when the threshold is 1117, and right two panels show payments when the threshold is 1132. Solid lines represent the overall density of payments, dashed lines represent six climate models, and the dotted lines represent payments under the current climate. The top row uses a smaller standard deviation for the kernel density estimation (adjust=1), while the bottom row uses a larger standard deviation (adjust=10) and therefore results in a smoother, wider multi-model distribution shown as the solid line.

	CRCM-CCSM	CRCM-CGCM3	WRFG-CCSM	WRFG-CGCM3	HRM3-HaDCM3	MM5I-HaDCM3
Bakersfield						
biasc.spr	-14.6	-14.0	-6.0	-6.6	0.1	-9.4
varc.spr	0.9	0.9	0.9	1.0	0.9	1.2
biasc.sum	-6.3	-11.4	-6.4	-9.8	8.1	-5.4
varc.sum	1.0	1.0	1.0	0.9	0.8	1.0
Fresno						
biasc.spr	-6.9	-7.1	-6.5	-7.4	-0.5	-9.4
varc.spr	0.9	0.9	0.9	1.0	0.9	1.2
biasc.sum	3.1	-11.4	-6.4	-9.8	8.1	-5.4
varc.sum	0.9	1.0	1.0	0.9	0.8	1.0
Los Angeles						
biasc.spr	-0.2	-14.0	-6.0	-6.6	0.1	-9.4
varc.spr	1.2	0.9	0.9	1.0	0.9	1.2
biasc.sum	1.5	-11.4	-6.4	-9.8	8.1	-5.4
varc.sum	1.1	1.0	1.0	0.9	0.8	1.0
Sacramento						
biasc.spr	-6.9	-14.0	-6.0	-6.6	0.1	-9.4
varc.spr	0.9	0.9	0.9	1.0	0.9	1.2
biasc.sum	3.1	-11.4	-6.4	-9.8	8.1	-5.4
varc.sum	0.9	1.0	1.0	0.9	0.8	1.0

Table 2: Downscaling parameters in spring and summer for each city, when matching current temperature to historical temperature. Each column represents a different climate model. Within each city, the first two rows are bias corrections and variance corrections in spring, and the last two rows are these corrections in summer.

	75^{th} *percentile*	95^{th} *percentile*
Fresno	1117	1132
Los Angeles	203	209
Sacramento	497	575
Bakersfield	1013	1026

Table 3: The 75^{th} percentile and 95^{th} percentile of cCDDs calculated from bootstrapping current temperatures, respectively, used as thresholds in payments. The same fixed threshold is used for payments under both the current and future climate. The top row matches the thresholds used in Figure 2.

	CRCM-CCSM	CRCM-CGCM3	WRFG-CCSM	WRFG-CGCM3	HRM3-HaDCM3	MM5I-HaDCM3
Bakersfield						
biasc.spr	4.1	3.8	3.0	3.6	3.5	3.3
biasc.sum	6.6	4.8	3.8	3.4	5.4	4.4
Fresno						
biasc.spr	3.5	3.2	2.9	2.9	3.9	3.2
biasc.sum	6.2	4.9	4.2	3.6	5.6	4.4
Los Angeles						
biasc.spr	2.7	2.7	2.7	2.8	3.3	3.6
biasc.sum	3.0	2.4	3.0	2.3	3.3	3.7
Sacramento						
biasc.spr	3.8	2.8	2.8	2.6	3.6	3.3
biasc.sum	5.7	4.6	4.2	3.2	5.4	4.9

Table 4: This table shows the mean of downscaled future temperature minus mean of downscaled current temperature. This amount is then added to each single bootstrapping temperature to calculate future bootstrapping temperature.

Payment 1		Current	Future	CRCM-CCSM	CRCM-CGCM3	WRFG-CCSM	WRFG-CGCM3	HRM3-HaDCM3	MM5I-HaDCM3
Bakersfield	95% VaR	14	492	510	424	339	332	450	379
	95% CTE	21	502	517	431	346	338	457	386
	99% VaR	25	508	521	436	351	343	461	391
	99% CTE	31	516	527	442	357	349	467	397
Fresno	95% VaR	15	485	502	411	362	325	475	386
	95% CTE	22	495	509	418	369	332	482	393
	99% VaR	25	501	513	422	372	335	485	397
	99% CTE	31	508	519	427	378	341	491	403
Los Angeles	95% VaR	6	296	235	201	234	198	275	311
	95% CTE	9	305	241	207	241	204	281	317
	99% VaR	11	310	245	211	244	208	285	321
	99% CTE	13	316	251	216	250	213	290	326
Sacramento	95% VaR	79	458	516	417	396	331	491	452
	95% CTE	114	505	554	455	434	369	529	490
	99% VaR	136	532	579	479	458	393	553	514
	99% CTE	173	572	662	522	502	436	597	558

Table 5: Risk measures under Payment #1 (Equation 6). The Current column shows risk measures computed under the current climate. The Future columns shows risk measures computed for the future climate, as defined across the multi-model combination of the six regional climate models shown. The six columns on the right show risk measures under the future climate described by that model along.

5 Discussion

We recap the methodology here before mentioning a few points of discussion. In this paper we fit a statistical model to daily temperature, with the aim of forecasting daily temperatures to simulate distributions of cCDDs in the following 3 months. We performed these forecasts under the current climate, and then again under a hypothetical future climate described by one of six regional climate models. Regional climate models were first downscaled to correct for biases in the mean and variance. After downscaling, these future climates are overall warmer than then current climate under the carbon emissions scenario considered, and therefore distributions of cCDDs were shifted notably to the right. We combined distributions from each of these six RCMs into an overall future distribution through kernel density estimation, and computed risk measures such as the VaR and CTE for $p = 0.95$ and $p = 0.99$ levels.

The methodology described in this paper is meant primarily to demonstrate the use of an ensemble of regional climate model outputs for an actuarial purpose. Therefore, the primary "result" is simply a guided tour of how a climate model is stored, visualized, manipulated, utilized in conjunction with observed data for bias and variance correction, and combined with other models to form an ensemble which captures model uncertainty. However, even within this context the particular results in this example are illuminating for the user who wishes to utilize climate model projections for other purposes. Table 5 illustrates a number of important things. First and most obvious is the staggering difference between the risk measures in the "Current" and "Future" columns. Recall that the threshold for payments was set using the current climate, such that only 25% of bootstrapped cCDD runs exceeded it. Moving to the future climate, if models project changes upwards of 1000+ cCDDs over the year (and they do, see Figure 2), then of course these models can project increases in the range of 300-500 over the 92 day period considered in this paper. The fixed threshold is easily exceeded by realizations of cCDDs under the future climate, and payments grow accordingly.

Thus, the striking difference in Table 5 is a reminder that unless climate-defined thresholds or benchmarks move in tandem with climate change, exceedances can grow quickly.

The distinctions across the six regional models in Table 5 give a sense of model uncertainty, and underscore the importance of using a robust ensemble of regional climate models rather than 1-2 models of convenience. Every climate model is constructed slightly differently, and every model parameterizes local phenomena (such as clouds) differently. The methodology presented in this paper was replicated across six different models and those results combined into a multi-model ensemble distribution of payments, thus capturing not only the uncertainty within each climate model but also across them. In general, we recommend that actuaries always replicate analyses for an ensemble of models unless there is a particular and compelling reason to prefer a single model.

Acknowledgements

The authors wish to thank the Research Expanding Boundaries (REX) Funding Pool of the Society of Actuaries for their generous financial support, as well as the Project Oversight Group (POG) who provided valuable feedback and perspectives for this paper.

References

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

Barnett, B. J. and Mahul, O. (2007). Weather index insurance for agriculture and rural areas in lower-income countries. *American Journal of Agricultural Economics*, 89(5):1241–1247.

Carriquiry, M. A. and Osgood, D. E. (2012). Index insurance, probabilistic climate forecasts, and production. *Journal of Risk and Insurance*, 79(1):287–300.

Chang, C. W., Chang, J. S., and Lim, K. G. (2012). Global warming, extreme weather events, and forecasting tropical cyclones: A market-based forward-looking approach. *ASTIN Bulletin: The Journal of the IAA*, 42(1):77–101.

Chantarat, S., Mude, A. G., Barrett, C. B., and Carter, M. R. (2013). Designing index-based livestock insurance for managing asset risk in northern kenya. *Journal of Risk and Insurance*, 80(1):205–237.

Christensen, J. H., Boberg, F., Christensen, O. B., and Lucas-Picher, P. (2008). On the need for bias correction of regional climate change projections of temperature and precipitation. *Geophysical Research Letters*, 35(20).

Collier, B., Skees, J., and Barnett, B. (2009). Weather index insurance and climate change: opportunities and challenges in lower income countries. *The Geneva Papers on Risk and Insurance-Issues and Practice*, 34(3):401–424.

Deng, X., Barnett, B. J., Vedenov, D. V., and West, J. W. (2007). Hedging dairy production losses using weather-based index insurance. *Agricultural Economics*, 36(2):271–280.

Erhardt, R. (2015). Mid-twenty-first-century projected trends in north american heating and cooling degree days. *Environmetrics*, 26(2):133–144.

Erhardt, R. (2017). *Climate, Weather and Environmental Sources for Actuaries*. Society of Actuaries.

Erhardt, R. and Von Burg, R. (2018). *How do they know and What could we do? The science of 21st century climate projections and opportunities for actuaries*. Society of Actuaries.

Fischer, T., Su, B., Luo, Y., and Scholten, T. (2012). Probability distribution of precipitation extremes for weather index–based insurance in the zhujiang river basin, south china. *Journal of Hydrometeorology*, 13(3):1023–1037.

Fleege, T. A., Richards, T. J., Manfredo, M. R., Sanders, D. R., et al. (2004). The performance of weather derivatives in managing risks of specialty crops. In *NCR-134 Conference on Applied Commodity Price Analysis, Forecasting, and Market Risk Management, St. Louis, Missouri*, pages 19–20.

Gent, P. R., Danabasoglu, G., Donner, L. J., Holland, M. M., Hunke, E. C., Jayne, S. R., Lawrence, D. M., Neale, R. B., Rasch, P. J., Vertenstein, M., et al. (2011). The community climate system model version 4. *Journal of Climate*, 24(19):4973–4991.

Hansen, J., Ruedy, R., Sato, M., and Lo, K. (2010). Global surface temperature change. *Reviews of Geophysics*, 48(4).

Hansen, J., Sato, M., and Ruedy, R. (2012). Perception of climate change. *Proceedings of the National Academy of Sciences*, 109(37):E2415–E2423.

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

Linnerooth-Bayer, J. and Mechler, R. (2006). Insurance for assisting adaptation to climate change in developing countries: a proposed strategy. *Climate policy*, 6(6):621–636.

MCII (2013). *Climate risk adaptation and insurance. Reducing vulnerability and sustaining the livelihoods of low-income communities. Report No. 13. Bonn: United Nations University Institute for Environment and Human Security.* (UNU-EHS).

Mearns, L., McGinnis, S., Arritt, R., Biner, S., Duffy, P., Gutowski, W., Held, I., Jones, R., Leung, R., Nunes, A., et al. (2007). The north american regional climate change assessment program dataset. *National Center for Atmospheric Research Earth System Grid data portal, Boulder, CO*.

Mearns, L. O., Arritt, R., Boer, G., Caya, D., Duffy, P., Giorgi, F., Gutowski, W., Held, I., Jones, R., Laprise, R., et al. (2005). Narccap, north american regional climate change assessment program. In *16th Conference on Climate Variability and Change*.

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

Nakicenovic, N., Alcamo, J., Grubler, A., Riahi, K., Roehrl, R., Rogner, H.-H., and Victor, N. (2000). *Special Report on Emissions Scenarios (SRES), A Special Report of Working Group III of the Intergovernmental Panel on Climate Change*. Cambridge University Press.

Pachauri, R. K., Allen, M. R., Barros, V. R., Broome, J., Cramer, W., Christ, R., Church, J. A., Clarke, L., Dahe, Q., Dasgupta, P., et al. (2014). *Climate change 2014: synthesis report. Contribution of Working Groups I, II and III to the fifth assessment report of the Intergovernmental Panel on Climate Change*. IPCC.

Turvey, C. G. (2005). The pricing of degree-day weather options. *Agricultural Finance Review*, 65(1):59–85.

Turvey, C. G. and Mclaurin, M. K. (2012). Applicability of the normalized difference vegetation index (ndvi) in index-based crop insurance design. *Weather, Climate, and Society*, 4(4):271–284.

Zeng, L. (2000). Weather derivatives and weather insurance: concept, application, and analysis. *Bulletin of the American Meteorological Society*, 81(9):2075–2082.