

Stochastic decadal climate simulations for the Berg and Breede Water Management Areas, Western Cape province, South Africa

Arthur M. Greene,¹ Molly Hellmuth,¹ and Trevor Lumsden²

Received 15 July 2011; revised 20 April 2012; accepted 21 April 2012; published 2 June 2012.

[1] A method is described for the generation of multivariate stochastic climate sequences for the Berg and Breede Water Management Areas in the Western Cape province of South Africa. The sequences, based on joint modeling of precipitation and minimum and maximum daily temperatures, are conditioned on annualized data, the aim being to simulate realistic variability on annual to decadal time scales. A vector autoregressive (VAR) model is utilized for this purpose and reproduces well those statistical attributes, including intervariable correlation and serial autocorrelation in individual variables, most relevant for the regional climate in this setting. The sequences incorporate nonlinear climate change trends, inferred using an ensemble of global climate models from the Coupled Model Intercomparison Project (CMIP5). Subannual variability is simulated using a block resampling scheme based on the k-nearest-neighbor approach, preserving both temporal patterns and spatial correlations. Downscaling to a network of quinary-level catchments enables distributed runoff, streamflow, and crop simulations and the assessment and integration of impacts. Final output takes the form of daily sequences, structured for driving the ACRU agrohydrological model of the University of KwaZulu-Natal, South Africa.

Citation: Greene, A. M., M. Hellmuth, and T. Lumsden (2012), Stochastic decadal climate simulations for the Berg and Breede Water Management Areas, Western Cape province, South Africa, *Water Resour. Res.*, 48, W06504, doi:10.1029/2011WR011152.

1. Introduction

[2] Interest in regional climate change and its potential impacts has increasingly come to focus on decadal time horizons, this perspective sometimes referred to as “near term.” Such a time scale is often felt to be of more immediate relevance than the centennial scales that have typically been considered in assessment reports of the *Intergovernmental Panel on Climate Change (IPCC)* [2007]. However, decadal climate forecasting is still very much a nascent science [Meehl *et al.*, 2009; Mehta *et al.*, 2011], with much current research attempting to simply characterize the degree to which the atmosphere-ocean system is potentially predictable on decadal time horizons [e.g., Boer and Lambert, 2008; Teng and Branstator, 2011]. Thus, the prospect of well-validated regional decadal forecasts, particularly for terrestrial regions, has not yet been achieved.

[3] A potentially useful alternative lies in the creation of stochastic sequences: synthetic data series having suitable decadal-scale statistical properties. Appropriately down-scaled and including a climate change component, these may be used to drive hydrology or other impacts models (possibly integrated or chained together) to explore the re-

silience of planned adaptation measures to a range of plausible climate variations, or scenarios, for the next few decades. Indeed, such assessments should continue to be useful if and when skillful decadal forecasts become a reality, given the uncertainty inherent in all forecasts. At least one study [Boer, 2009] suggests that, as climate warms, potential predictability on decadal time scales may decrease. Should this turn out to be case, the importance of scenario generation can be expected to grow accordingly.

[4] The utilization of stochastic simulations at the *weather* timescale is not new (see Wilks and Wilby [1999] for a review); there have also been attempts to incorporate climate change information in such sequences [e.g., Semenov and Barrow, 1997; Wilks, 1999; Kilsby *et al.*, 2007]. Application with the decadal time scale in mind, the focus of the work described here, has also been undertaken [Prairie *et al.*, 2008; Kwon *et al.*, 2007, 2009], but constitutes a less well explored domain. A focus on the decadal scale shifts the emphasis toward regional low-frequency variability and its potential for augmenting (or compensating) secular, forced climate change on decadal time horizons.

[5] The simulations discussed herein are generated on an annual time step for the study area as a whole, then down-scaled to individual locations and daily time resolution for application in ACRU [Schulze, 1995], which requires daily values for precipitation and maximum and minimum temperatures. (ACRU, formerly the Agricultural Catchments Research Unit model, has been generalized to include non-agricultural areas and is now known simply by its acronym.) Dependence among the variables requires that they be simulated jointly; accordingly, a multivariate modeling framework is adopted. Secular, anthropogenically forced trends are inferred using an ensemble of global climate

¹International Research Institute for Climate and Society, Earth Institute at Columbia University, Palisades, New York, USA.

²School of Agricultural, Earth and Environmental Sciences, University of KwaZulu-Natal, Pietermaritzburg, South Africa.

Corresponding author: A. M. Greene, International Research Institute for Climate and Society, Earth Institute at Columbia University, Lamont Campus, Palisades, NY 10964, USA. (amg@iri.columbia.edu)

models (GCMs) from the CMIP5 project [Taylor *et al.*, 2012].

[6] ACRU simulates runoff but also incorporates some crop modeling capabilities, and for impact studies its output will be used to drive an economic model. The present work does not extend to this model coupling but focuses on the methodology of simulation, the first link in this chain. Warburton *et al.* [2010] provide an assessment of present-day simulations using ACRU.

[7] In section 2 an overview is provided of the physical, hydrological and economic setting. Section 3 describes the data utilized and section 4 the methodology. Model validation is discussed in section 5 and application in section 6. Some trailing issues are considered in section 7, and a summary is provided in section 8.

2. The Study Area

2.1. Description

[8] Situated in the Western Cape region of South Africa, the study area (Figure 1) comprises the Berg Water Management Area (WMA) and parts of the neighboring Breede WMA, and covers approximately 19,000 km². The Berg and Breede Rivers flow northwestward and eastward, draining into the Atlantic and Indian Oceans, respectively. The two WMAs are treated jointly because they are linked by interbasin transfers, the region's water resources being managed as an integrated system. The study area is characterized by steep gradients in both altitude and rainfall: Mountainous areas along the divide between the WMAs reach elevations of nearly 2000 m and may receive annual mean precipitation in excess of 3000 mm, while near the mouth of the Berg annual means fall below 200 mm.

[9] With the aid of intensive irrigated agriculture, the Western Cape produces high-value crops such as wine and table grapes, deciduous fruits and citrus; the region makes a significant contribution to South Africa's national economy. At the same time urban water demand, notably in the city of Cape Town, has tripled since the late 1970s and con-

tinues to increase, by about 4% annually [Louw and van Schalkwyk, 2000], and there is now direct competition for water between urban and agricultural sectors.

2.2. Delineation of Subcatchments

[10] For operational decision making, South Africa has been divided into quaternary catchments, these being a fourth level subdivision of the 22 primary catchments that cover South Africa, Lesotho and Swaziland. Schulze and Horan [2011] have further subdivided these into quinary catchments, on the basis of natural breaks in altitude. The study area comprises 171 quinary (or 57 quaternary) catchments.

3. Data

[11] Two types of data are employed, (1) a set of local time series associated with the quinary-level catchments and (2) temperature and precipitation simulations from a 14-member ensemble of global climate models (GCMs) from CMIP5. Observed regional-scale variability, on which the annual-to-decadal component of the simulations is conditioned, is estimated from the spatially averaged quinary time series. The GCM outputs are utilized both globally, to characterize the "forced" anthropogenic climate change signal and regionally, to estimate future temperature and precipitation responses to global warming.

3.1. Local Observations

[12] The quinary catchment data comprise 171 trivariate daily time series, for precipitation and maximum and minimum temperatures (pr, Tmax, and Tmin, respectively) and span the years 1950–1999. The data derive from a network of stations associated with the quaternary catchments, the same rainfall record being used for three, or occasionally more, neighboring quinary catchments [Schulze *et al.*, 2005, 2011]. In all, there are 44 independent precipitation records among the quinary data. When ACRU is run, adjustment factors, based on empirical relationships, are applied to the daily values, resulting in the creation of unique rainfall time series for each quinary catchment.

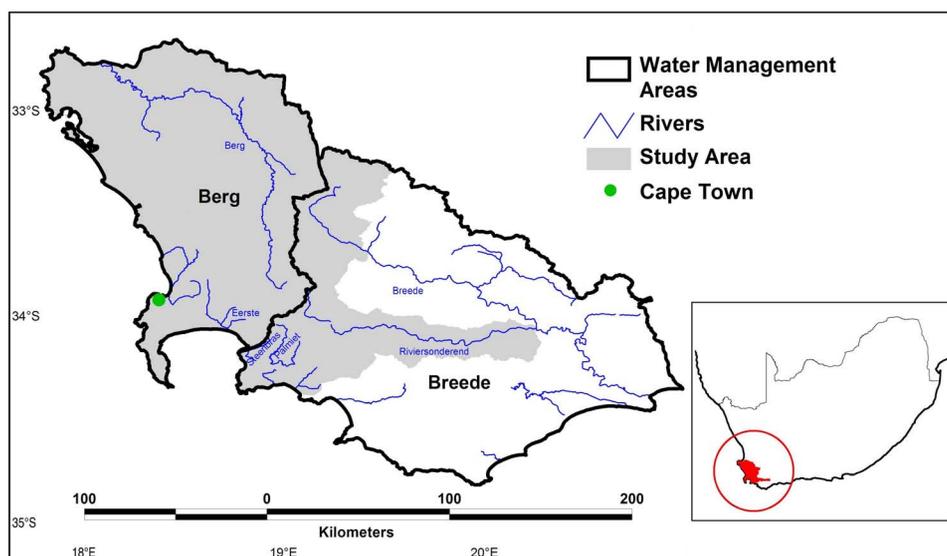


Figure 1. The Berg and Breede Water Management Areas and the study area (shading). The inset map shows the location within South Africa.

[13] The daily temperature data used to represent the quinary catchments were extracted from a gridded database of daily maximum and minimum temperatures (resolution one arc minute of latitude/longitude, corresponding to about 1.85/1.55 km) developed by *Schulze and Maharaj* [2004]. The mapping of quinary to grid points [*Schulze et al.*, 2011] produces unique records of maximum and minimum temperature for each catchment.

3.2. Information From GCMs

[14] The expected climate change signal, in this case the regional response for each of the three modeled variables to anthropogenic forcing, is estimated with the aid of an ensemble of 14 GCMs from CMIP5 (Table 1). Both global and regional temperature as well as regional precipitation fields are utilized. Twentieth century values are taken from the “historical” simulations, those for the 21st century from the RCP4.5 experiment [*Taylor et al.*, 2012], which represents a middle ground with respect to future emissions. Regional precipitation is computed as an average over 30°–35°S, 17°–23°E, an area encompassing the Western Cape and exhibiting a relatively homogeneous spatial pattern of precipitation change in most of the GCMs.

4. Method

[15] The simulation plan assumes three “process classes” that contribute on different time scales, and in differing degrees, to the variability expressed at each of the quinary catchments. These are the forced or anthropogenic component, a “natural” annual-to-decadal component and a sub-annual component, including the seasonal cycle as well as day-to-day variations. The first two of these are modeled at annual time resolution, the trend inferred using the CMIP5 ensemble and the annual-to-decadal component simulated

by the VAR model. Subannual variability is generated by preferentially resampling the observational data in 1 year blocks. The k-NN approach utilized for this step allows for shifts in daily rainfall statistics that may come about as a result of climatic changes. A possible increase in interannual precipitation variability with global temperature was investigated, but was not corroborated by the CMIP5 ensemble and is not modeled.

[16] Figure 2 shows a method schematic. In order that the overall flow be clearly illustrated, some of the symbols on this diagram have been assigned multiple levels of significance: The Input node represents an individual catchment record with respect to the trend path (along the top of Figure 2), but the entire ensemble of catchment records with respect to the resampling path (along the bottom); elements of both the analysis (e.g., leading to the VAR node) and synthesis (the Simulate path) are shown; trend is treated differently for precipitation and for the temperature variables, and so on. These procedural details are clarified in sections 4.1–4.4.

4.1. Trends, Past and Future

[17] Local trends are modeled as functions not of time but of global mean temperature, the motivating idea being that trend should represent a response to anthropogenic forcing, rather than simply a shift in the mean level with time. The global mean surface temperature, suitably computed, has been shown to be an effective proxy for the forced climate response [*Ting et al.*, 2009]. Note that various features of the regional climate, including atmospheric circulation, may change as the planet warms. The global mean temperature may be thought of as an index of this warming.

[18] Recent work suggests that the observed global temperature record may be contaminated to some degree by internal decadal variations [*Ting et al.*, 2009; *DelSole et al.*, 2011]. For this reason the signal is computed here by averaging over the CMIP5 ensemble. The individual GCM simulations are low-passed using a fifth-order Butterworth filter [*Smith*, 2003] having half power at a frequency of 0.1 yr^{-1} , then averaged. (Results are not sensitive to the precise method of filtering.) Ensemble averaging has the effect of attenuating unforced climate fluctuations, since these are uncorrelated from model to model, while enhancing that part of the signal that the individual GCMs have in common, namely, the response to anthropogenic forcing. The smoothing reduces both residual interannual variability and the effects of short-lived transients such as volcanic eruptions. For further discussion of this procedure see *Greene et al.* [2011a].

[19] When a local record is regressed on the smoothed multimodel mean signal the fitted values, representing the trend, take the form of a scaled, shifted version of the regressand. Three such fits appear as the trend lines in Figure 3, which shows the annualized regional (i.e., catchment-averaged) observational series. Note that this procedure implicitly removes any remaining additive bias in the multimodel mean temperature record. Detrending is accomplished by subtracting the fitted values. The regression coefficients are used, in conjunction with the future multimodel mean temperature signal, for forward projection, now in terms of the response to future anthropogenic forcing.

Table 1. GCMs in the CMIP5 Archive Utilized^a

Center	Model
BCC (China)	BCC-CSM1.1
CCCma (Canada)	CanESM2
CNRM-CERFACS (France)	CNRM-CM5
CSIRO (Australia)	CSIRO-Mk3-6
INM (Russia)	INM-CM4
IPSL (France)	IPSL-CM5A-LR
MIROC (Japan)	MIROC-ESM
MIROC (Japan)	MIROC-ESM-CHEM
MIROC (Japan)	MIROC5
MOHC (United Kingdom)	HadGEM2-ES
MPI-M (Germany)	MPI-ESM-LR
NASA GISS (United States)	GISS-E2-R
NCAR (United States)	CCSM4
NCC (Norway)	NorESM1-M

^aInstitutions are as follows: BCC, Beijing Climate Center, China Meteorological Administration; CCCma, Canadian Centre for Climate Modeling and Analysis; CNRM-CERFACS, Centre National de Recherches Météorologiques/Centre Européen de Recherche et Formation Avancée en Calcul Scientifique; CSIRO, Commonwealth Scientific and Industrial Research Organisation and the Queensland Climate Change Centre of Excellence; INM, Institute for Numerical Mathematics; IPSL, Institut Pierre Simon Laplace; MIROC, Atmosphere and Ocean Research Institute (University of Tokyo), National Institute for Environmental Studies, and Japan Agency for Marine-Earth Science and Technology; MOHC, Met Office Hadley Centre; MPI-M, Max Planck Institute for Meteorology; NASA GISS, NASA Goddard Institute for Space Studies; NCAR, National Center for Atmospheric Research; NCC, Norwegian Climate Centre.

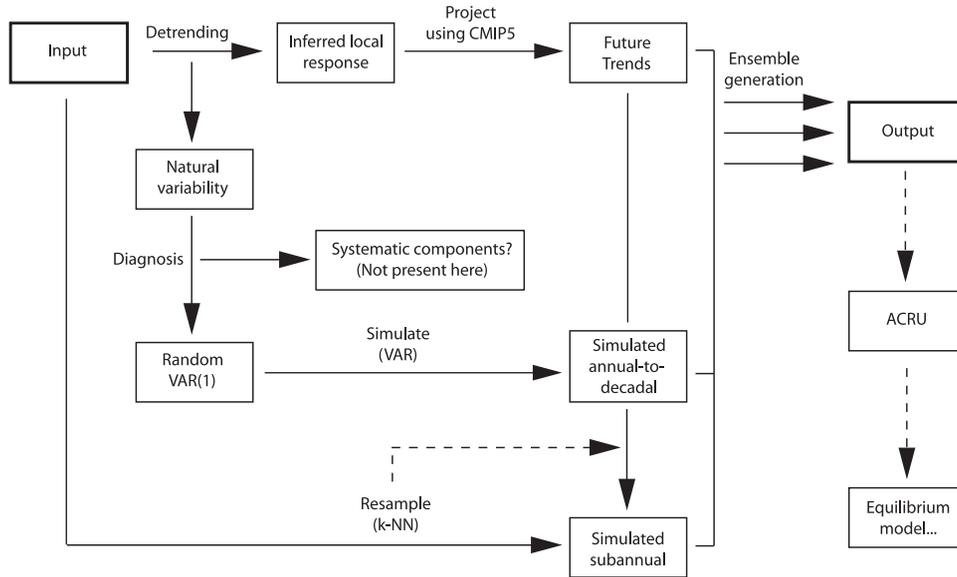


Figure 2. Method schematic: Nodes (boxes) represent signals, and edges (lines) the processes that link them. The dashed line from Resample indicates implicit trend-subannual coupling via the k-nearest neighbors (k-NN) scheme. Dashed lines from Output and ACRU represent processes outside the scope of the present study.

[20] In going from the 20th to the 21st century regional temperature and precipitation variables, as simulated in the CMIP5 ensemble, behave quite differently, as shown in Figure 4. Regional temperature projects consistently on the global mean, as evidenced by the uniform slope across centuries (Figure 4a). Precipitation (Figure 4b) exhibits considerably greater variability, but 20th century values (from the “historical” simulations) do not trend significantly, while the 21st century regression (based on the RCP4.5 experiments) is significant at $p = 10^{-4}$. Observed regional precipitation also lacks any significant trend for the period of record, 1950–1999. (To facilitate comparison with the

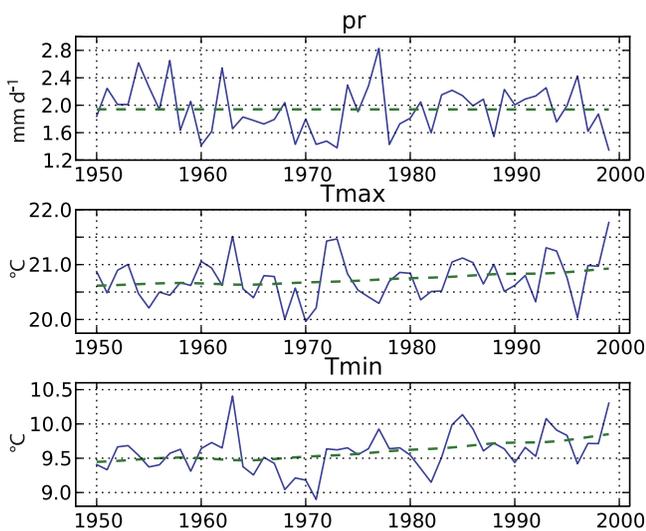


Figure 3. The three regional time series on which simulations are based. Dashed lines are the fitted trends, from regression on the low-passed CMIP5 multimodel mean global temperature record.

observational record, the 20th century values are shown only for this period; the “historical” data actually extend through 2005, the RCP4.5 simulations beginning in 2006.) Because of this difference in behavior, the trend component is treated differently for the temperature and precipitation variables.

[21] Future trends are generated at the catchment level. For maximum and minimum temperatures, this is accomplished by regressing the annualized catchment records on the smoothed multimodel mean signal, as described above, then applying the resulting coefficients to the future multimodel mean temperature record. This enforces a consistent relationship between 20th and 21st century behavior, with respect to the global mean.

[22] The temperature response differs among catchments and between Tmax and Tmin (catchment means of $0.51 \pm 0.22(1\sigma)$ and 0.62 ± 0.30 , respectively). Forward projection in this manner thus implies a rather complex set of changes in surface temperature gradients over time, while the more rapid increase of Tmin, compared with Tmax, leads to a mean reduction of the diurnal temperature range (DTR). Divergent temperature tendencies could eventually evoke compensating behaviors, such as small-scale circulation adjustments that would act to reduce local gradients. However the reduction in DTR could represent a shift toward a new equilibrium state [see, e.g., Braganza et al., 2004].

[23] Because of this complexity, and because the simulations under discussion extend just a few decades into the future, we do not attempt to include compensating mechanisms for temperature trends in the simulation model. This could be done, for example, by relaxing catchment trends toward a common mean, insuring that local gradients do not become unrealistically large. However there is some spatial dependence in temperature trends, for which an additional level of modeling would be required.

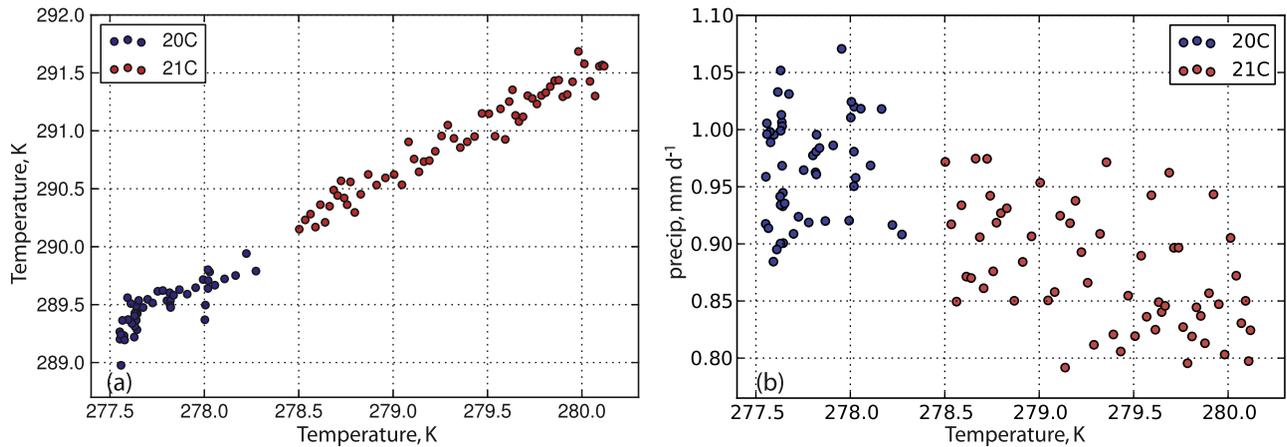


Figure 4. Scatterplots of (a) regional mean temperature and (b) precipitation against global mean temperature, CMIP5 ensemble means. Annual mean values are shown for 1950–1999, corresponding to the observational period of record, and from 2006, when the RCP4.5 simulations begin, through 2065; the region is 30°–35°S, 17°–23°E.

[24] Annualized catchment-level precipitation records are similarly regressed on the smoothed multimodel mean, but the resulting coefficients are utilized in a different manner. To begin with, the regional 21st century precipitation response differs among GCMs. The distribution of this response is shown in Figure 5, along with a Gaussian fit. Trends are computed in log space, the coefficients then representing the fractional change in regional precipitation (shown in Figure 5 as percentage change) per degree of global temperature increase. The distribution has mean -6.7% per degree warming, with a standard deviation of 6.9% . (These values refer to the 2006–2065 period in the RCP4.5 experiment; the 2000–2005 interval, belonging to the historical simulations, has an intermediate character and its future precipitation trend is interpolated between 20th and 21st century values.) Three of the 14 models become wetter with warming temperatures, suggesting a nonnegligible probability ($\sim 17\%$ in terms of the fitted Gaussian) of

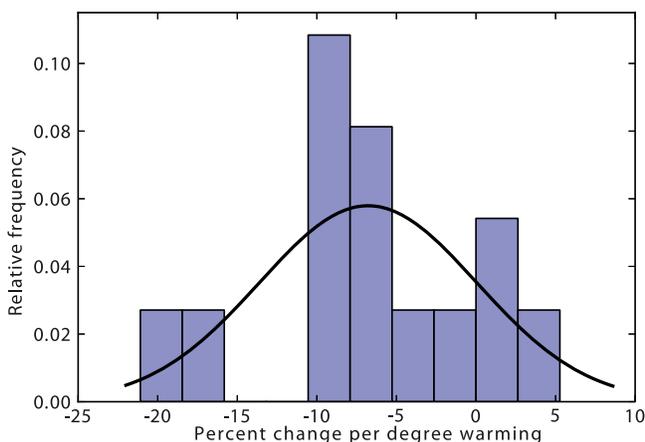


Figure 5. Distribution of the regional precipitation response to global mean temperature change in the CMIP5 ensemble utilized here. Regression is carried out in log space: the abscissa shows the response as the percent change per degree of global temperature increase. The curve is a Gaussian fit to the data.

such an outcome. Note that absolute GCM precipitation values are not utilized in these computations, bypassing a potential source of bias.

[25] In projecting the precipitation trend, a desired quantile is first specified and the corresponding value calculated using the fitted Gaussian. Catchment-level trends are then computed as

$$T_{c21} = T_r + \alpha T_{c20}, \quad (1)$$

where T_{c21} is the future catchment trend, T_r the quantile-based regional trend, T_{c20} the 20th century catchment trend and α is a factor, here set provisionally at 0.5. Thus, the catchment-averaged future trend will correspond to the imposed quantile-based value, while individual catchment trends will scatter around this value according to their 20th century behavior. (Recall that the average 20th century catchment trend is not significantly different from zero.) The degree of scatter, α , is at the operator's disposal, but is attenuated here in order that study area precipitation not become overly “disorganized” as the simulations are projected into the future. Catchment-level 20th century trends show no dependence on either altitude or location within the study area, suggesting a significant random component in their dispersion.

[26] From the physical perspective there is some reason to believe that the Western Cape will dry in coming decades, owing to poleward migration of the dry subtropical belts and midlatitude storm tracks. Indeed, some of this migration has already been observed [Seidel *et al.*, 2008; Yin, 2005]. The phenomenon is also suggested in Figure 11.2 of the IPCC Fourth Assessment Report [IPCC, 2007, p. 869], which shows that the drying projected for the Western Cape is not an isolated regional phenomenon but part of a coherent global pattern.

4.2. Regional Annual-to-Decadal Variability

[27] The annual-to-decadal component of the simulation model is based on January–December means of the three variables, averaged over the study area (i.e., the series shown in Figure 3, but after detrending). Owing to the

winter (JJA) maximum in Western Cape rainfall, time averaging in this way does not bisect the rainy season. Area averaging reduces small-scale “noise” that is uncorrelated across catchments, while enhancing whatever large-scale, quasi-regional signal the catchments share. Climate variability on annual-to-decadal scales may be expected to arise chiefly from large-scale oceanic or coupled ocean-atmosphere processes [Schlesinger and Ramankutty, 1994; Trenberth and Hurrell, 1994; Mantua *et al.*, 1997]. Terrestrial climate variations, conditioned by such processes through atmospheric teleconnections, would tend to exhibit relatively large-scale spatial signatures [Hurrell, 1996; Enfield *et al.*, 2001; Glantz *et al.*, 1991].

4.2.1. Data Attributes

[28] Tyson *et al.* [2002] discussed a pervasive 18 year oscillation in southern African climate. The claimed signal, present in both instrumental and paleorecords, was strongest in the 20°–30°S latitude band, but detectable to the southern extremity of the continent. Our regional record was tested for the presence of such a signal, using both singular spectrum analysis (SSA), a technique well suited for detecting quasiperiodic signals in short, noisy time series [Ghil *et al.*, 2002], as well as wavelet analysis [Torrence and Compo, 1998]. Neither method confirms the presence of such an oscillation, as illustrated by the precipitation wavelet spectrum shown in Figure 6a. (Oscillations at the claimed frequency would fall outside the “cone of influence” in this plot but should nevertheless be visible if present.) Likewise, SSA and wavelet analyses of the Tmax and Tmin records (wavelet spectra shown in Figures 6b and 6c) do not suggest the presence of significant oscillatory components. There are many possible reasons for such a discrepancy, including spatial or temporal inhomogeneity of the 18 year signal and the analysis of disparate data sets. Without disputing the claims made by Tyson *et al.*, it is concluded that such an oscillation is not present in the data utilized here.

[29] Lag 1 autocorrelation coefficients for the regional variables are significant at 0.10 and 0.05 for Tmax and Tmin, respectively (one-sided test) but not for precipitation. The Durbin-Watson test for serial autocorrelation is a bit more confident, yielding p values of 0.06 and 0.01 for Tmax and Tmin, respectively. Thus the “decadal” (i.e., persistent) component of the observational record appears to reside in the temperature variables, the precipitation sig-

nal being essentially indistinguishable from white noise. The wavelet spectra of Figure 6, despite the presence of some episodic activity in the 10 year band, do not suggest (via the presence of significant peaks in the global spectra) the presence of systematic signal components, i.e., components that differ from AR(1) in character.

[30] Owing to the well-defined water year, the modeling of seasonal (JJA) values was considered. However the ACRU model, because it includes the memory effects of soil moisture, requires full years of simulated climate. Additionally, water stresses tend to be highest during the summer months, when rainfall is low and evaporation high, while area physiography limits the potential for buffering via the construction of new dams. Thus, behavior outside the rainy season is also of significant interest. Preliminary inspection of the JJA statistics indicates that they do not differ greatly from those derived from annual values.

[31] A composite scatterplot (Figure 7) shows a negative correlation ($r = -0.45$, significant at 0.001) for pr and Tmax, and a stronger positive correlation ($r = 0.73$) for Tmax and Tmin. The first of these may reflect rain-associated cloudiness and/or reduction of the Bowen ratio by surface moistening. Both mechanisms involve a significant insolation component, which may explain the lack of correlation between pr and Tmin. In any event, failure to represent these relationships correctly would likely bias the simulations, and, by extension, the resulting outputs from ACRU.

[32] It has been hypothesized that global warming will bring about an increase in year-to-year precipitation variability, owing to the exponential dependence on temperature of water saturation vapor pressure. Figure 8 shows the CMIP5 ensemble distributions of interannual precipitation variance for the 30°–35°S, 17°–23° E domain for 1950–1999 and 2046–2065, by which time global mean temperature has increased by $\sim 1.5^\circ\text{C}$. Almost no change in the distribution of variance is observed, so there would seem to be little justification for modeling such a dependence. The CMIP3 simulations are in accord on this point.

4.2.2. Statistical Model

[33] The absence of significant peaks in the three regional spectra (for pr, Tmax and Tmin), together with the serial correlation exhibited by the temperature variables, suggests the deployment of a vector autoregressive (VAR) model, a multivariate generalization of the classical AR

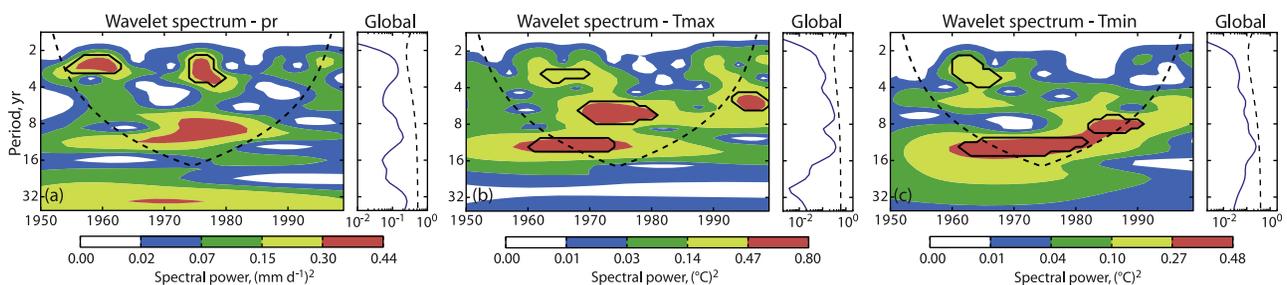


Figure 6. Wavelet power spectra for the detrended regional series: (a) precipitation, (b) Tmax, and (c) Tmin. Level boundaries correspond to the 25th, 50th, 75th, and 95th percentiles of spectral power, and the solid black contours correspond to the 0.05 red noise significance level. The thick dashed line delineates the “cone of influence,” outside of which edge effects become important. For each variable the panel at the right shows the global wavelet spectrum (solid line) and the 0.05 red noise significance level (dashed line).

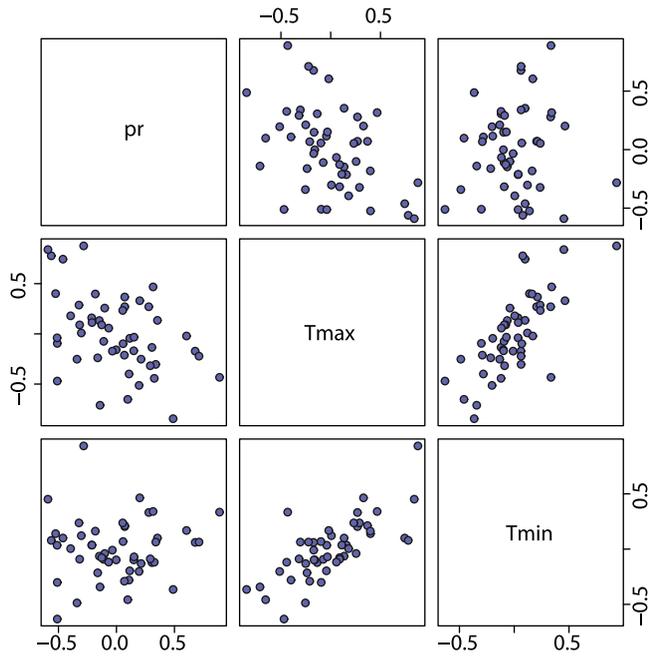


Figure 7. Scatterplots for the (detrended) regional series. Units are mm d^{-1} for pr and degrees Celsius for Tmax and Tmin.

model, for the annual-to-decadal component of the simulations. VAR models have been widely utilized in both econometrics (see, e.g., Holden [1995] and other papers in that issue) and climate studies [e.g., Penland and Sardeshmukh, 1995; Newman, 2007], where a VAR model of order unity is known as a linear inverse model.

[34] The first-order VAR model can be written

$$\mathbf{y}_t = \mathbf{A}\mathbf{y}_{t-1} + \mathbf{u}_t, \quad (2)$$

where \mathbf{y}_t is the 3×1 (pr, Tmax, Tmin) climate vector at time t , \mathbf{A} is a 3×3 matrix of coefficients and \mathbf{u}_t is a 3×1 stationary white noise process with expectation $\mathbf{0}$ and covariance matrix $\Sigma_u = E(\mathbf{u}_t \mathbf{u}_t^T)$. Note that \mathbf{u} need not be white in parameter space, i.e., there may be some contemporaneous correlation between the variables, corresponding to nonzero off-diagonal elements in Σ_u .

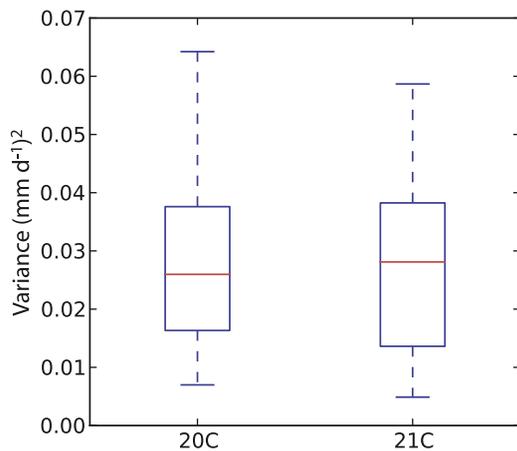


Figure 8. Interannual precipitation variance for the 20th and 21st century simulations. Distributions across the CMIP5 ensemble are shown.

[35] A model of the form (2) was fitted to the annualized regional series by least squares, using the dynamical systems estimation (dse) package [Gilbert, 1995] for the R programming language [Ihaka and Gentleman, 1996]. One-step-ahead forecasts are shown in Figure 9, where it can be seen that the fraction of variability accounted for by the predictive component of the model (the first term on the right-hand side of equation (2)) is modest. While some of this predictability arises from autocorrelation, some may also result from lagged cross correlations. To quantify contributions from the latter the Granger causality test [Granger, 1969] was applied. This pairwise test assesses the predictability above and beyond that arising from serial autocorrelation that is contributed by lagged cross-variable dependence. Results suggest a limited degree of additional predictability for precipitation and Tmin based on the inclusion of Tmax and precipitation, respectively, at the preceding time step (p values of 0.17 and 0.14, pooled across variables).

[36] Although one could perhaps make a case for a model lacking serial dependence, we retain the VAR structure, in part on the accumulated evidence but also because the hydrological significance of this dependence is not known a priori. It should therefore be instructive to compare ACRU outputs with those based on simulations from a model lacking the predictive term. Ultimately the VAR model is also more general, and can better serve as a prototype for application in a diverse range of settings. Neither Akaike’s information criterion (AIC [Akaike, 1973]) nor the Bayesian information criterion (BIC [Raftery, 1986]) suggest that there is anything to be gained by moving beyond the complexity of a first-order model.

4.3. Subannual Variations

[37] Subannual variability is generated by resampling the observations in 1 year blocks, using a modified k-NN scheme [see, e.g., Rajagopalan and Lall, 1999] in which the three-component feature vector consists of a single year’s simulated annual means of pr, Tmax and Tmin. The aim is to select, from among the 50 data years, one whose mean annual values approximate this vector. Since year-to-year dependence is already accounted for by the VAR model there is no role in this scheme for a “successor”: A particular year having been chosen from among the candidates, its subannual patterns of variability are appropriated for the year being simulated.

[38] Experimentation suggested that the use of $k = 5$ nearest neighbors provides a reasonable compromise between the generation of sufficient variety in the resultant sequences and the inclusion of too remote candidates (see section 5.2). The Mahalanobis distance metric [Mahalanobis, 1936] is utilized, with weights of $(\frac{2}{3}, \frac{1}{6}, \frac{1}{6})$ assigned to pr, Tmax, and Tmin, respectively, effectively weighting precipitation double the combined weights of the two temperature variables. These weights are based on past results from ACURU implicating precipitation as the most important predictor of runoff, the key application variable, but should be considered provisional subject to further experimentation. (Alternate weights may prove desirable for studies focusing on the summer dry season.)

[39] A monotonically decreasing resampling kernel, with values $\propto 1/j, j = 1, 2, \dots, 5$, is utilized to select from among

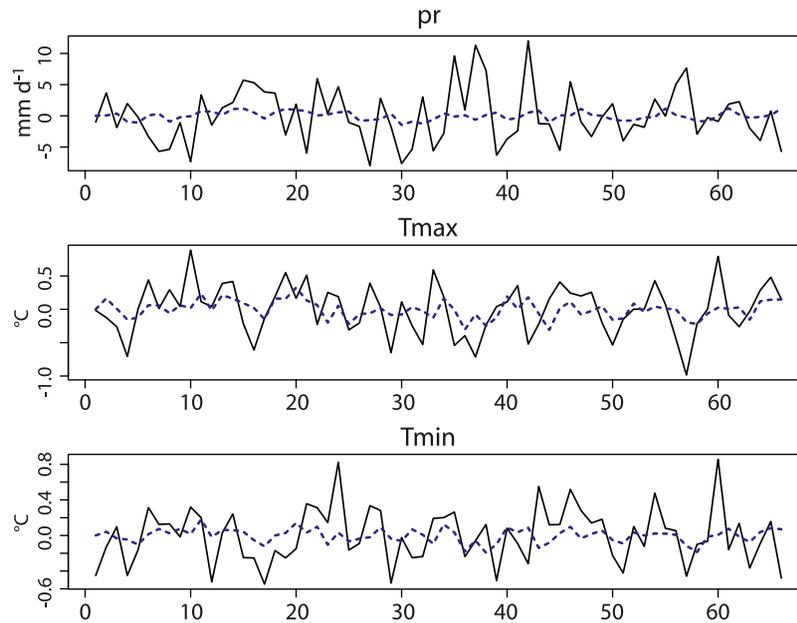


Figure 9. One-step-ahead predictions (dashed lines) for the VAR(1) model. Solid lines show the observed regional series.

the nearest neighbors. The resampling scheme explicitly links the climate change and subannual time scales, in principle enabling the realization of climatically induced changes in daily rainfall statistics, as secular trends in the mean state induce shifts in the population from which resampled statistics are drawn. The three variables are resampled jointly and across the entire study area, preserving spatial coherence (including potential shifts in spatial patterns driven by large-scale mean changes), high-frequency covariation and seasonal cycle shape.

4.4. Downscaling

[40] Figure 10 shows coefficient distributions for regressions of the detrended catchment-level variables on the corresponding regional series. Because the regional signals represent catchment means, the average coefficient for each

of the variables is unity, guaranteeing that the catchment-averaged response will reproduce the imposed simulation sequence. The plots give some idea of the degree to which annualized catchment variations follow those of the regional signal, or, put another way, the degree to which the regional signal is expressed at each of the catchments.

[41] In the downscaling step, simulated annual-level variations are propagated to the catchments using these coefficients. Uncorrelated noise is added to bring variances into agreement, emulating the observed variability. The resulting signals are substituted for intrinsic annual-level catchment variations by adjusting the resampled catchment values, in 1 year blocks. This is done additively for the temperature variables, multiplicatively for precipitation. ACRU is driven by the resulting daily sequences, superimposed on the CMIP5-derived trends.

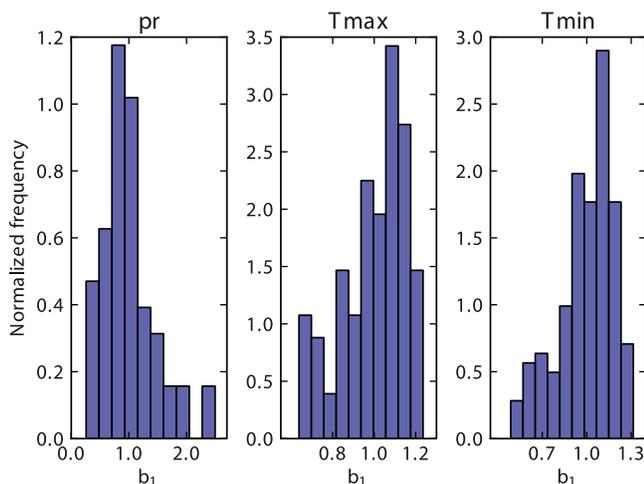


Figure 10. Distributions of the coefficient b_1 for regressions of the annualized catchment-level variables on the corresponding regional time series.

5. Simulation and Model Checking

[42] The fitted VAR model is initially used to generate a single very long sequence at the annual time step ($N = 500,000$). This is 10,000 times the length of the observational record and provides many more realizations than would be needed (or practical) for driving ACRU. The profusion of data is useful, however, in that it provides both precise estimates of simulation statistics and the opportunity to select, from a large ensemble of possible realizations, a small set having well-constrained properties.

5.1. Annual-Level Simulation Statistics

[43] Lag 1 autocorrelation coefficients for the detrended regional time series and the long simulation are shown in Table 2, where it can be seen that the coefficients for the simulated series mimic their targets fairly closely. Differences between the observed and simulated coefficients are considerably smaller than coefficient sampling variability, which approximates 0.14 for all variables. Correlation matrices for the observations and simulation are given in Table 3,

Table 2. Lag 1 Autocorrelation Coefficients for pr, Tmax, and Tmin for the Regional Observations and Long Simulation

Source	pr	Tmax	Tmin
Observations	0.003	0.166	0.293
Simulation	-0.010	0.175	0.299

which shows that quite a close correspondence has been achieved. As noted, these correlations may arise through either the modeled lag 1 dependencies or the innovations term u_t (i.e., either of the terms on the right-hand side of equation (2)). The VAR model thus captures well the two characteristics of the regional series that are of importance for the simulation of annual-to-decadal variations in this setting.

5.2. Subannual Simulation Characteristics

5.2.1. Spell-Related Properties

[44] Figure 11 compares observed and simulated distributions for dry and wet spell counts and wet spell amounts, in the form of quantile-quantile plots. Comparisons are made on the 1950–1999 period, statistics being computed over the 44 catchments having unique precipitation records. Simulations based on three nonoverlapping segments, randomly sampled from the long simulation sequence, are plotted. Figure 11 shows that the simulated distributions closely approximate the observed. Plots for spell lengths (not shown) show similarly high degrees of correspondence; examination of additional sequences suggests that these are representative.

[45] The quantiles shown in Figure 11 are computed on spell data pooled over catchments, obscuring the degree to which distributions may or may not correspond at the catchment level. This question is addressed, taking a slightly different perspective, in Figure 12, which shows coefficients for linear regressions of log-transformed catchment wet spell amounts on annual mean regional precipitation. The response of subannual statistics to mean changes is of interest because simulation values are propagated (via k-NN and subsequent adjustment) in terms of annual means. For the observations only one 50 year “sample” is available for computing these coefficients, and serves here as a reference. Fifty simulations were generated on the basis of (nonoverlapping) segments drawn at random from the long simulation sequence, yielding the distributions plotted at each of the catchments.

[46] Figure 12 shows that coefficient distributions computed from the simulations bracket the observed response at each of the catchments. Somewhat more than half (30 of 44) of the observed coefficients lie within the interquartile ranges at the individual catchments. This number ranges from 28 to 38 over the five spell parameters investigated,

Table 3. Contemporaneous Correlations for the Regional Observations and Long Simulation

	Observations			Simulation		
	pr	Tmax	Tmin	pr	Tmax	Tmin
pr	1.000			1.000		
Tmax	-0.449	1.000		-0.445	1.000	
Tmin	-0.066	0.732	1.000	0.068	0.733	1.000

wet and dry spell counts and lengths and wet spell mean daily amounts, suggesting that the spread of the simulated distributions might be somewhat too large. However, this does not take into account uncertainty in the observed coefficients themselves, so the simulated coefficient distributions may in fact not be unreasonably broad.

5.2.2. Precipitation Extremes

[47] Although the k-NN scheme selects for years with climate characteristics that are “close” to those being simulated, some adjustment of annual means of the three variables is still required in order that they replicate the imposed, simulated values. For precipitation this adjustment takes the form of a multiplicative scaling. Given the parameters adopted for the k-NN scheme, the resampled annual sequences prior to this scaling account for an average of 90% of the variance in the values to be simulated (based on the 50-sequence sample), so the required adjustments are relatively small. Nevertheless it is of interest to investigate whether the scaling might distort extreme rainfall distributions. To this end, the dependence of 3 day block maxima (i.e., highest 3 day rainfall total in each year) on annual mean regional precipitation was examined. As in the analysis of spell-related properties only the 44 unique precipitation records were considered.

[48] Regressions were conducted on both the actual and log-transformed data, yielding coefficients that reflect absolute and fractional change in 3 day precipitation totals, given absolute and fractional changes in annual mean precipitation, respectively. In absolute terms, 3 day totals were found to increase 40 ± 21 mm (1σ) per unit annual mean increase (expressed as mm d^{-1}) for the observations and 41 ± 23 mm for the simulations, the distributions taken over catchments. In fractional terms the corresponding values are 0.90 ± 0.27 for the observations, 0.97 ± 0.38 for the simulations. The observed and simulated coefficients for individual catchments are well correlated ($r = 0.86$, mean over the 50 simulations), indicating that differential catchment sensitivity in block precipitation maxima is well replicated in the simulations.

[49] Simulated and observed responses are effectively indistinguishable (in the aggregate), and in fractional terms are indistinguishable from unity, suggesting that, at least in this setting, precipitation extremes do scale linearly with the mean. This proportional behavior supports the use of multiplicative precipitation adjustments in the downscaling step.

[50] As long as projected annual means for the three variables lie within or near their 20th century ranges, the adjustment of subannual variations will constitute a second-order correction, following the resampling step. Given the high interannual variability of precipitation, the most important of the three variables, compared with projected mean changes, this condition is likely to be met in many cases, in particular for wet extremes in the expected drying climate of the Western Cape.

6. Toward Application

6.1. Specification and Screening

[51] The length of the simulated sequence allows for screening on multiple criteria, thus the exploration of

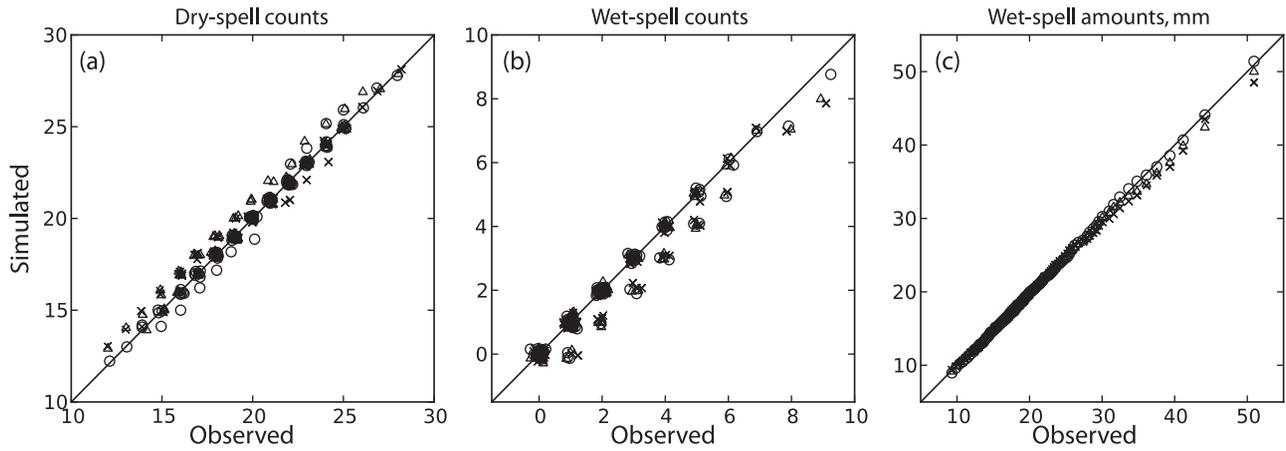


Figure 11. Quantile-quantile plots for (a) dry and (b) wet spell counts and (c) wet spell mean amounts (mm d^{-1}) for three randomly selected simulations compared with observations over the 44 unique catchment records. Wet (dry) spells must be at least 3 (5) days in length, and rainfall during wet spells must be at least 10 mm on each day. Markers are plotted at each 0.01 quantile from 0.01 to 0.99, and in Figures 11a and 11b they have been shifted by small random amounts (“jittered”) for readability.

potential impacts in an efficient manner. Figure 13 shows annual values for two simulations, both for the same quinary catchment, which lies at an elevation of 870 m and has mean annual precipitation for 1950–1999 of 2.0 mm d^{-1} . For both simulations the 50th percentile, corresponding to a 6.7% precipitation decrease per degree global temperature increase, was specified for the regional 21st century response. The 20th century trend at this catchment, inferred from the observations, is small but negative, a reduction of 1.2% per degree warming. When combined according to (1) the net response is a reduction of 7.3% per degree global temperature increase, or -0.19 mm d^{-1} by the middle of the 2041–2050 decade, given a warming by that time of 1.3°C . The choice of decade in these examples is arbitrary, but for illustrative purposes it was desired that the precipitation reduction be significant.

[52] Simulated values in Figure 13 begin in year 2000. To select the particular sequences shown, the long simulation was screened on four criteria: First, it was desired to obtain some sense of the relative contributions of trend and decadal variability over the near term. As it turns out,

decadal anomalies lying at the 5th and 95th percentiles correspond to deviations of $\pm 0.18 \text{ mm d}^{-1}$, nearly matching the secular drying at this catchment during the target decade. Catchment precipitation follows the regional signal closely, with a regression coefficient of 0.94 (see Figure 10), so there is very little dilution of the simulated signal in the downscaling step. The initial screening specification is thus that 10 year mean precipitation be constrained to lie near the 5th (or 95th) percentile.

[53] Figure 7 tells us that pr and Tmax are anticorrelated: Years when these variables are both anomalously high (or low) are relatively unlikely. Since we are interested here in the “typical” impacts of precipitation fluctuations, it was further required that the temperature variables lie within reasonable ranges of their conditional expected values, given the specified precipitation departure. Finally, because there is some hydrologic memory of past conditions, anomalous precipitation during the preceding decade was required to be reasonably small. Screening for the specified values ± 0.1 , ± 0.6 , ± 0.6 , and ± 1.0 standard deviations, for decadal precipitation, Tmax, Tmin, and the

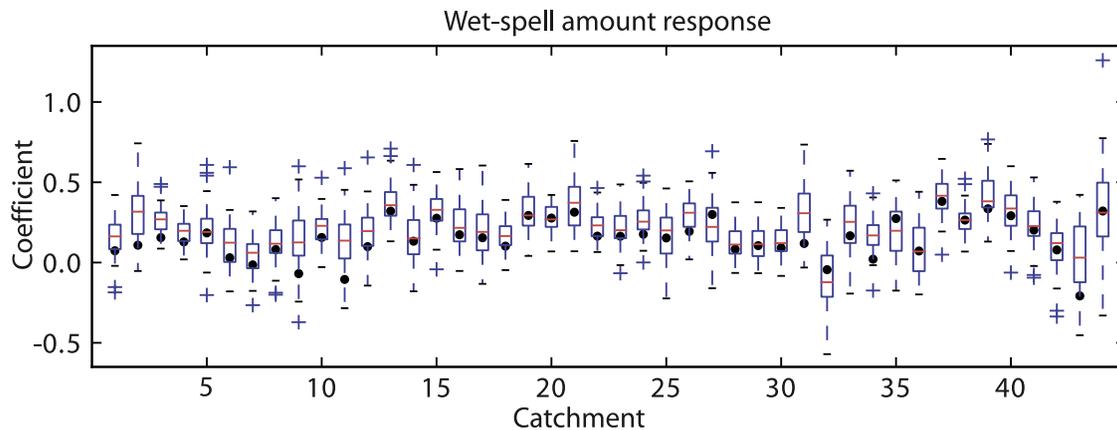


Figure 12. Catchment-level coefficients for log wet spell mean amounts regressed on annual mean precipitation. Distributions are taken over 50 simulations. Black markers show observed values.

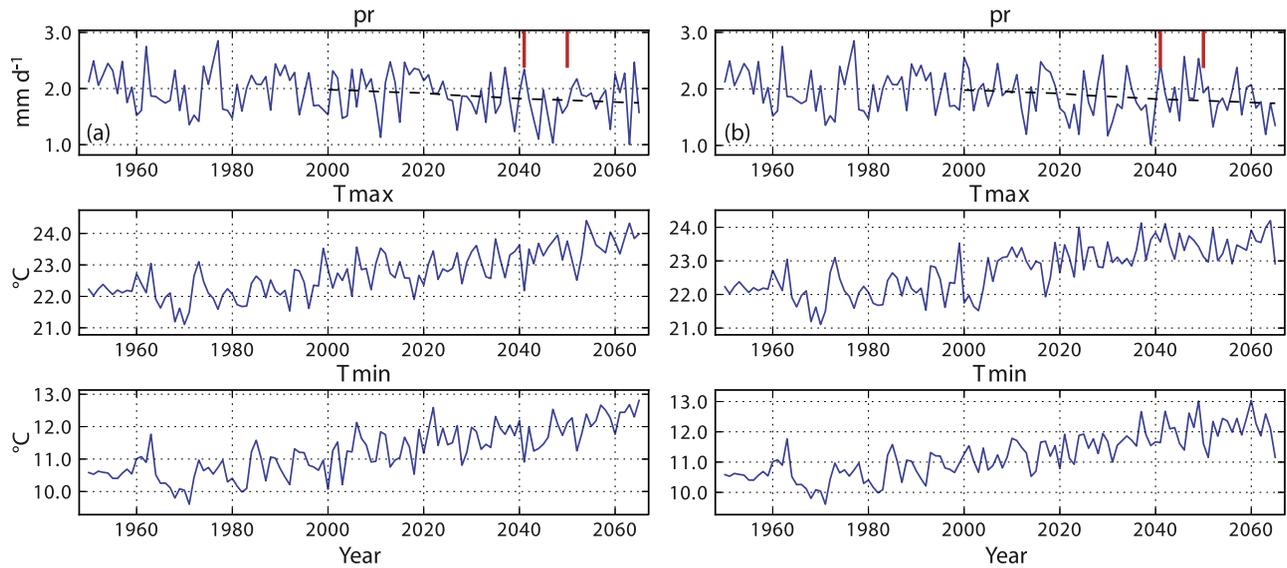


Figure 13. Two simulations for the same quinary catchment and with the same precipitation trend but opposing decadal fluctuations. Decadal anomalies were specified for the 2041–2050 period, indicated by red bars. Observed, rather than simulated, values are shown for 1950–1999.

preceding decade's precipitation, respectively, yielded several hundred simulation candidates, from which the examples shown were selected. Finally, the long sequence was sliced so that the specified decadal fluctuations occur during the 2041–2050 decade.

[54] For application purposes it may be useful to calculate the quantile corresponding to a particular decadal precipitation value, resulting from the joint specification of trend and fluctuation percentiles. Knowing the distributions of both trend and fluctuation, the desired value can be computed by first expressing the trend as a precipitation anomaly, as we have done above. The sum of trend in this form with the decadal fluctuation can then be evaluated in terms of the sum of their distributions.

6.2. Balance of Trend and Variability

[55] In the scenario just described, the 5th-percentile decadal fluctuation produces a near doubling of the precipitation reduction owing to trend (Figure 13a), while the 95th-percentile anomaly (Figure 13b) results in an almost complete cancellation. Because the trend is computed in fractional terms while fluctuations take the form of anomalous precipitation rates, this balance will differ across catchments. Mean precipitation at the example catchment is close to the regional value (see Figure 3), so the balance there can be expected to approximate that of the study area as a whole. Catchment-level factors that may modify this balance, aside from mean precipitation, include variations in 20th century trend and the degree to which the catchment subscribes to the regional signal.

[56] From the regional viewpoint and assuming the GCM mean precipitation response, Figure 14 indicates that by about 2040 mean annual precipitation would be expected to decline to the level of the present-day 5th percentile for decadal means, i.e., the level demarcating the driest five percent of decades in the present climate. The reduction comes about in the context of significant interannual variability,

such that this level is reached fairly frequently, even at present, in individual years, about one in three. The GCM mean response amounts to a reduction in annual mean precipitation of about 10% by midcentury, given the RCP4.5 global temperature increase ($\sim 1.5^\circ\text{C}$) projected by this time. A shift of this magnitude would be expected to have significant impacts. However a more comprehensive modeling framework is required in order to fully understand the implications of trend, variability and their interplay in the complex physical and socioeconomic setting of the Western Cape. ACRU represents the next step in this process.

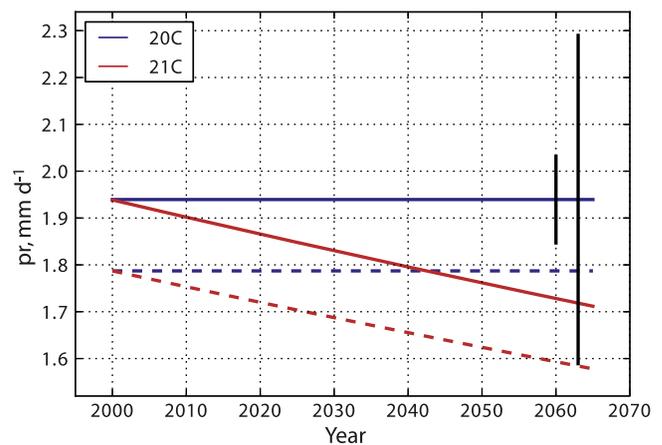


Figure 14. GCM mean precipitation response in the context of regional variability. Blue solid and dashed lines indicate the 20th century regional mean and 5th percentile for decadal means, respectively. Red lines show the mean CMIP5 response applied to these values, given the projected RCP4.5 temperature increase. Shorter and longer black bars at the right show $\pm 1\sigma$ for 10 year and annual means, respectively.

6.3. Application in ACURU

[57] A set of screened sequences, including both high and low trend and decadal fluctuation quantiles, will be used to drive ACURU, enabling the exploration of a range of plausible impacts whose likelihoods can now be quantified. In addition to water availability, ACURU represents crop-related responses such as evapotranspiration, irrigation demand and yields. Its simulations will be combined with economic modeling to produce impact assessments for agricultural and urban sectors as well as comparative evaluations of adaptation options, including demand management, infrastructure development and water trading. The results of these experiments will be reported in a separate publication.

7. Some Trailing Considerations

[58] Supplementing the observational record with tree ring or other paleodata could in principle permit more robust estimation of climatic behavior on the annual-to-decadal scale. The regionally relevant tree ring chronology that was examined by *Tyson et al.* [2002] apparently exhibits a rather complex dependence on climatic factors [*Dunwiddie and LaMarche*, 1980] rendering it difficult to interpret, and we are unaware of other, possibly less ambiguous proxy records in or near the study area. We have therefore not attempted to incorporate such data in the present analysis.

[59] The trend in our simulation example represents the average CMIP5 response (with a small addition from the intrinsic catchment trend). The distribution of these responses was discussed in section 4.1. *Gleckler et al.* [2008] show that over a wide range of both surface and atmospheric fields, the multi-model ensemble mean consistently outperforms individual models, so perhaps more confidence can be placed in the mean response, and less in the outlying GCMs, than the distribution of Figure 5 suggests. On the other hand, the 14-GCM ensemble utilized here may not sample the full range of possible regional variations, so the true distribution could in fact be broader than that suggested by Figure 5. Detailed process studies may aid in the evaluation of GCM behavior, eventually permitting a differential assignment of credibility among models and, consequently, more precise uncertainty estimates.

[60] A range of sophisticated downscaling schemes exist [e.g., *Mehrotra and Sharma*, 2010; *Greene et al.*, 2011b], and one may ask whether such a scheme might be substituted for the k-NN/rescaling deployed herein. On the one hand, our use of a relatively straightforward method preserves an emphasis on the manner of combining information across time scales that we believe constitutes the principal point of interest of this study. More fundamentally, however, the procedure used here arises naturally in conjunction with the other components of the method: Since regional variability down to the annual scale is prescribed by the VAR model, use of a scheme that generates such variability by other means (perhaps using GCM output) would result in overspecification. Thus the integration of any alternate methodology would require careful consideration, and the range of applicable methodologies may be limited.

8. Summary

[61] A method, based on the decomposition of variability into trend, annual-to-decadal, and subannual components,

is utilized to generate stochastic climate simulations for the Berg and Breede Water Management Areas, Western Cape province, South Africa. Simulations are produced on a daily time step and are structured for driving the ACURU agrohydrology model. Long-range trends are inferred using an ensemble of 14 of the CMIP5 GCMs, respecting the considerable spread they exhibit in regional precipitation response. Variability on annual-to-decadal time scales is simulated using a first-order vector autoregressive (VAR) model fit to annualized observed values of precipitation and minimum and maximum daily temperatures, while subannual variations are generated via a modified k-NN resampling algorithm. The simulated sequences preserve both spatial coherence and the temporal characteristics important for the annual-to-decadal time scale, and link subannual statistics, including spell-related behavior and precipitation extremes (evaluated here using 3 day block maxima), to climatic trends in a manner consistent with observed variability. The VAR model is used to generate a single long simulation sequence at the annual time step, from which short segments may be extracted by screening against multiple criteria. Follow-on modeling using ACURU will explore a range of scenarios, for which the long simulation provides an ample “library.”

[62] Because the simulated components of variability are superimposed on secular trends that may carry the system to states beyond the range of those visited in the past, climatic stresses may be produced, on a range of time scales, that exceed those previously experienced. It is hoped and expected that these simulations, utilized as inputs to ACURU, will ultimately prove useful in delineating the climate risks with which the Western Cape may have to contend in coming decades.

[63] **Acknowledgments.** We are grateful for the many helpful comments and suggestions provided by colleagues at both the International Research Institute for Climate and Society and the Lamont-Doherty Earth Observatory of Columbia University, by Roland Schulze at the University of KwaZulu-Natal, Mac Calloway at the UNEP Risø Centre, Roskilde, and by three anonymous reviewers, whose detailed comments helped measurably to improve the manuscript. This work was carried out with the aid of a grant from the International Development Research Centre, Ottawa, Canada, ID 104150-001. A.M.G.’s work was supported in part by a grant from NOAA, NA08OAR4320912. The views expressed herein are those of the authors and do not necessarily reflect those of NOAA or any of its subagencies.

References

- Akaike, H. (1973), Information theory and an extension of the maximum likelihood principle, in *Second International Symposium on Information Theory*, edited by B. N. Petrov and F. Csaki, pp. 267–281, Akad. Kiado, Budapest.
- Boer, G. J. (2009), Changes in interannual variability and decadal potential predictability under global warming, *J. Clim.*, 22, 3098–3109, doi:10.1175/2008JCLI2835.1.
- Boer, G. J., and S. J. Lambert (2008), Multi-model decadal potential predictability of precipitation and temperature, *Geophys. Res. Lett.*, 35, L05706, doi:10.1029/2008GL033234.
- Braganza, K., D. J. Karoly, and J. M. Arblaster (2004), Diurnal temperature range as an index of global climate change during the twentieth century, *Geophys. Res. Lett.*, 31, L13217, doi:10.1029/2004GL019998.
- DelSole, T., M. K. Tippett, and J. Shukla (2011), A significant component of unforced multidecadal variability in the recent acceleration of global warming, *J. Clim.*, 24(3), 909–926, doi:10.1175/2010JCLI3659.1.
- Dunwiddie, P. W., and V. C. LaMarche Jr. (1980), A climatically responsive tree-ring record from *Widdringtonia cedarbergensis*, Cape Province, South Africa, *Nature*, 286, 796–797.
- Enfield, D. B., A. M. Mestas-Nunez, and P. J. Trimble (2001), The Atlantic Multidecadal Oscillation and its relationship to rainfall and river flows in

- the continental U.S., *Geophys. Res. Lett.*, 28(10), 2077–2080, doi:10.1029/2000GL012745.
- Ghil, M., et al. (2002), Advanced spectral methods for climatic time series, *Rev. Geophys.*, 40(1), 1003, doi:10.1029/2000RG000092.
- Gilbert, P. D. (1995), Combining VAR estimation and state space model reduction for simple good predictions, *J. Forecast.*, 14, 229–250.
- Glantz, M. H., R. W. Katz, and N. Nicholls (Eds.) (1991), *Teleconnections Linking Worldwide Climate Anomalies*, 548 pp., Cambridge Univ. Press, Cambridge, U. K.
- Gleckler, P. J., K. E. Taylor, and C. Doutriaux (2008), Performance metrics for climate models, *J. Geophys. Res.*, 113, D06104, doi:10.1029/2007JD008972.
- Granger, C. W. J. (1969), Investigating causal relations by econometric models and cross-spectral methods, *Econometrica*, 37, 424–438.
- Greene, A. M., L. Goddard, and R. Cousin (2011a), Web tool deconstructs variability in twentieth-century climate, *Eos Trans. AGU*, 92(45), 397–398, doi:10.1029/2011EO450001.
- Greene, A. M., A. W. Robertson, P. Smyth, and S. Triglia (2011b), Down-scaling projections of Indian monsoon rainfall using a non-homogeneous hidden Markov model, *Q. J. R. Meteorol. Soc.*, 137, 347–359.
- Holden, K. (1995), Vector autoregression modelling and forecasting, *J. Forecast.*, 14, 159–166.
- Hurrell, J. W. (1996), Influence of variations in extratropical wintertime teleconnections on Northern Hemisphere temperature, *Geophys. Res. Lett.*, 23(6), 665–668, doi:10.1029/96GL00459.
- Ihaka, R., and R. Gentleman (1996), R: A language for data analysis and graphics, *J. Comput. Graphical Stat.*, 5(3), 299–314.
- Intergovernmental Panel on Climate Change (IPCC) (2007), *Climate Change 2007: The Physical Science Basis*, edited by S. Solomon et al., 996 pp., Cambridge Univ. Press, Cambridge, U. K.
- Kilsby, C., P. Jones, A. Burton, A. Ford, H. Fowler, C. Harpham, P. James, A. Smith, and R. Wilby (2007), A daily weather generator for use in climate change studies, *Environ. Modell. Software*, 22(12), 1705–1719, doi:10.1016/j.envsoft.2007.02.005.
- Kwon, H.-H., U. Lall, and A. F. Khalil (2007), Stochastic simulation model for nonstationary time series using an autoregressive wavelet decomposition: Applications to rainfall and temperature, *Water Resour. Res.*, 43, W05407, doi:10.1029/2006WR005258.
- Kwon, H.-H., U. Lall, and J. Obeysekera (2009), Simulation of daily rainfall scenarios with interannual and multidecadal climate cycles for south Florida, *Stochastic Environ. Res. Risk Assess.*, 23, 879–896, doi:10.1007/s00477-008-0270-2.
- Louw, D. B., and H. D. van Schalkwyk (2000), Water markets: An alternative for central water allocation? *Agrekon*, 39(4), 484–494.
- Mahalanobis, P. C. (1936), On the generalised distance in statistics, *Proc. Natl. Acad. Sci. India*, 2(1), 49–55.
- Mantua, N. J., S. R. Hare, Y. Zhang, J. M. Wallace, and R. C. Francis (1997), A Pacific interdecadal climate oscillation with impacts on salmon production, *Bull. Am. Meteorol. Soc.*, 78(6), 1069–1079.
- Meehl, G. A., et al. (2009), Decadal prediction: Can it be skillful? *Bull. Am. Meteorol. Soc.*, 90(10), 1467–1485.
- Mehrotra, R., and A. Sharma (2010), Development and application of a multisite rainfall stochastic downscaling framework for climate change impact assessment, *Water Resour. Res.*, 46, W07526, doi:10.1029/2009WR008423.
- Mehta, V., G. Meehl, L. Goddard, J. Knight, A. Kumar, M. Latif, T. Lee, A. Rosati, and D. Stammer (2011), Decadal climate predictability and prediction: Where are we?, *Bull. Am. Meteorol. Soc.*, 92(5), 637–640.
- Newman, M. (2007), Interannual to decadal predictability of tropical and North Pacific sea surface temperatures, *J. Clim.*, 20, 2333–2356, doi:10.1175/JCLI4165.1.
- Penland, C., and P. D. Sardeshmukh (1995), The optimal growth of tropical sea surface temperature anomalies, *J. Clim.*, 8, 1999–2024.
- Prairie, J., K. Nowak, B. Rajagopalan, U. Lall, and T. Fulp (2008), A stochastic nonparametric approach for streamflow generation combining observational and paleoreconstructed data, *Water Resour. Res.*, 44, W06423, doi:10.1029/2007WR006684.
- Raftery, A. E. (1986), Choosing models for cross-classification, *Am. Sociol. Rev.*, 51(1), 145–146.
- Rajagopalan, B., and U. Lall (1999), A k-nearest-neighbor simulator for daily precipitation and other weather variables, *Water Resour. Res.*, 35(10), 3089–3101.
- Schlesinger, M. E., and N. Ramankutty (1994), An oscillation in the global climate system of period 65–70 years, *Nature*, 367, 723–726, doi:10.1038/367723a0.
- Schulze, R. E. (1995), Hydrology and agrohydrology: A text to accompany the ACRU 3.00 agrohydrological modelling system, *WRC Rep. TT 69/95*, Water Res. Comm., Pretoria, South Africa.
- Schulze, R. E., and M. J. C. Horan (2011), Methods 1: Delineation of South Africa, Lesotho and Swaziland into quinary catchments, in *Methodological Approaches to Assessing Eco-Hydrological Responses to Climate Change in South Africa*, *WRC Rep. 1562/1/10*, edited by R. E. Schulze et al., chap. 6, pp. 55–62, Water Res. Comm., Pretoria, South Africa.
- Schulze, R. E., and M. Maharaj (2004), Development of a database of gridded daily temperatures for southern Africa, *WRC Rep. 1156/2/04*, 83 pp., Water Res. Comm., Pretoria, South Africa.
- Schulze, R. E., M. Warburton, T. G. Lumsden, and M. J. C. Horan (2005), The Southern African Quaternary Catchments Database: Refinements to, and links with, the ACRU system as a framework for modelling impacts of climate change on water resources, in *Climate Change and Water Resources in Southern Africa: Studies on Scenarios, Impacts, Vulnerabilities and Adaptation*, *WRC Rep. 1430/1/05*, edited by R. E. Schulze, chap. 8, pp. 111–139, Water Res. Comm., Pretoria, South Africa.
- Schulze, R. E., M. J. C. Horan, R. P. Kunz, T. G. Lumsden, and D. M. Knoesen (2011), Methods 2: Development of the Southern African Quinary Catchments Database, in *Methodological Approaches to Assessing Eco-hydrological Responses to Climate Change in South Africa*, *WRC Rep. 1562/1/10*, edited by R. E. Schulze et al., chap. 7, pp. 63–73, Water Res. Comm., Pretoria, South Africa.
- Seidel, D. J., Q. Fu, W. J. Randel, and T. J. Reichler (2008), Widening of the tropical belt in a changing climate, *Nat. Geosci.*, 1, 21–24, doi:10.1038/geo.2007.38.
- Semenov, M. A., and E. M. Barrow (1997), Use of a stochastic weather generator in the development of climate change scenarios, *Clim. Change*, 35(4), 397–414, doi:10.1023/A:1005342632279.
- Smith, S. W. (2003), *Digital Signal Processing: A Practical Guide for Engineers and Scientists*, 650 pp., Newnes, Burlington, Mass.
- Taylor, K. E., R. J. Stouffer, and G. A. Meehl (2012), An overview of CMIP5 and the experiment design, *Bull. Am. Meteorol. Soc.*, 93, 485–498, doi:10.1175/BAMS-D-11-00094.1.
- Teng, H., and G. Branstator (2011), Initial-value predictability of prominent modes of north Pacific subsurface temperature in a CGCM, *Clim. Dyn.*, 36(9–10), 1813–1834, doi:10.1007/s00382-010-0749-7.
- Ting, M., Y. Kushnir, R. Seager, and C. Li (2009), Forced and internal twentieth-century SST trends in the North Atlantic, *J. Clim.*, 22, 1469–1481, doi:10.1175/2008JCLI2561.1.
- Torrence, C., and G. P. Compo (1998), A practical guide to wavelet analysis, *Bull. Am. Meteorol. Soc.*, 79(1), 61–78.
- Trenberth, K., and J. Hurrell (1994), Decadal atmosphere-ocean variations in the Pacific, *Clim. Dyn.*, 9, 303–319.
- Tyson, P. D., G. R. J. Cooper, and T. S. McCarthy (2002), Millennial to multi-decadal variability in the climate of southern Africa, *J. Climatol.*, 22, 1105–1117, doi:10.1002/joc.787.
- Warburton, M., R. Schulze, and G. Jewitt (2010), Confirmation of ACRU model results for applications in land use and climate change studies, *Hydrol. Earth Syst. Sci.*, 14, 2399–2414, doi:10.5194/hess-14-2399-2010.
- Wilks, D. S. (1999), Multisite downscaling of daily precipitation with a stochastic weather generator, *Clim. Res.*, 11, 125–136.
- Wilks, D. S., and R. L. Wilby (1999), The weather generation game: A review of stochastic weather models, *Prog. Phys. Geogr.*, 23, 329–357, doi:10.1177/030913339902300302.
- Yin, J. H. (2005), A consistent poleward shift of the storm tracks in simulations of 21st century climate, *Geophys. Res. Lett.*, 32, L18701, doi:10.1029/2005GL023684.