# Integrated framework for rapid climate stress testing on a monthly timestep

Keirnan Fowler, Environmental Hydrology and Water Resources Group, University of Melbourne, fowler.k@unimelb.edu.au

Natasha Ballis, Environmental Hydrology and Water Resources Group, University of Melbourne

Avril Horne, Environmental Hydrology and Water Resources Group, University of Melbourne, avril.horne@unimelb.edu.au

Andrew John, Environmental Hydrology and Water Resources Group, University of Melbourne, Andrew.john@unimelb.edu.au

Rory Nathan, Environmental Hydrology and Water Resources Group, University of Melbourne, Rory.nathan@unimelb.edu.au

Murray Peel, Environmental Hydrology and Water Resources Group, University of Melbourne, mpeel@unimelb.edu.au

# Integrated framework for rapid climate stress testing on a monthly timestep

Keirnan Fowler[1]*, Natasha Ballis[1], Avril Horne[1], Andrew John[1], Rory Nathan[1], Murray Peel[1]

**Affiliations:**

[1]Environmental Hydrology and Water Resources Group, Department of Infrastructure Engineering, University of Melbourne, Parkville, Australia

*Correspondence to: fowler.k@unimelb.edu.au

## Highlights

- We present a freely available framework for climate stress testing of water systems
- The framework includes stochastic climate generation, rainfall-runoff modelling, and perturbation across five different stressors, all in an integrated set of code
- The monthly timestep allows rapid run-times, providing greater possibilities for a given computational budget

## Abstract

"Bottom-up" methods are increasingly used to assess the vulnerability of water systems to climate change. Central to these methods is the climate "stress test", where the system is subjected to various climatic changes to test for unacceptable outcomes. We present a framework for climate stress testing on a monthly timestep, suitable for systems whose dominant dynamic is seasonal or longer (eg. water resource systems with carry-over storage). The framework integrates multi-site stochastic climate generation with perturbation methods and in-built rainfall runoff modelling. The stochastic generation includes a low frequency component suitable for representing multi-annual fluctuations. Multiple perturbation options are provided, ranging from simple delta change through to altered seasonality and low frequency dynamics. The framework runs rapidly, supporting comprehensive multi-dimensional stress testing without recourse to supercomputing facilities. We demonstrate the framework on a large water resource system in southern Australia. The Matlab/Octave framework is freely available for download from https://doi.org/10.5281/zenodo.5617008.

**Keywords:** *climate stress test, stochastic generation, water resources, climate change, bottom-up*

## 1   Introduction

Climate change is a key risk to water systems, and recent years have seen many new methods to characterise this risk (see eg. review by Wilby and Murphy, 2019). The vulnerability of a given water system to climate change depends on many factors, including location-specific projections of climate change and runoff, the robustness of the system to these changes, and the specific aspect of system performance in view (eg. Brown and Wilby 2012; Nathan et al., 2019). Vulnerability assessments must represent each of these broad factors to inform risk management and help establish robust policy.

Methods of climate change risk assessment can be categorized as either "top-down" or "bottom-up". Top down is the traditional approach and uses Global Climate Model (GCM) projections as the starting point in a chain of modelling steps that typically includes simulations of the rainfall-runoff response and system operations (eg. reservoir release rules, allocations, etc.). The way the GCM projections are used in top-down methods varies: in many studies the bias-corrected climate projections are used directly as forcing data for subsequent modelling

steps (eg. Christensen et al., 2004, Chiew et al., 2009, García-Ruiz et al., 2011); in others the *distribution* of GCM projections inform the distribution and/or range of future climate changes tested. Given the potential for scenario selection to directly impact the robustness of outcomes (McPhail et al, 2020), there is increasing effort to ensure appropriate scenario choices (Guidici et al 2020), although the impact of scenario choice on management decisions themselves (as opposed to metrics) is still an active research area (McPhail et al., 2020).

In contrast, bottom-up methods consider a range of plausible future climate sequences selected to investigate the system of interest and identify conditions which cause stress to the system. Stochastic methods are often used to generate future climatic sequences in bottom-up studies (eg. Brown et al., 2012; Turner et al., 2014; Henley et al., 2019). A broad range of plausible futures may be tested, with the ranges of key variables (or "stressors", eg. change in precipitation) often extended beyond the envelope of GCM scenarios – a so called "scenario neutral" approach (Prudhomme et al., 2010). Rather than starting with a GCM scenario and asking *"what does this scenario mean for the system?"*, the bottom-up procedure starts with the system and asks *"what combination of future conditions cause stress to this system?"*. Thus, the procedure is often called a "climate stress test" (Brown and Wilby, 2012). Having identified the combinations of future conditions that result in stakeholder-defined unacceptable outcomes, the probability of these outcomes may be assessed by checking what proportion of GCMs contain similar conditions in their projections (eg. Brown et al., 2012). Thus, GCM projections can have an important role in bottom-up assessments, but the way they are used is different to the top-down approach.

The bottom-up approach has significant advantages. The focus on specific system vulnerabilities is often more relevant to policy decisions (Brown and Wilby, 2012; Wilby and Murphy, 2019). The approach facilitates systematic assessment of policy options by their robustness to future changes of varying degree (eg. Weaver et al., 2012). Further, as newer GCM results become available they can be rapidly mapped onto existing analyses, in contrast to the top-down approach which may require the entire modelling chain to be re-run. Because of these advantages, bottom-up methods have been applied in many systems in the past decade, including in the United States (eg. Poff et al., 2015), United Kingdom (eg. Prudhomme et al., 2010), continental Europe (eg. Culley et al., 2016), Africa (eg. Ghile et al., 2013), Asia (eg. Yang et al., 2016) and Australia (eg. Turner et al., 2014; Culley et al., 2019).

Despite these advantages, implementing bottom-up methods presents technical challenges, particularly regarding the generation of future climatic sequences, which is the focus of the present framework. A minority of studies use simple methods such as scaling the historic timeseries (eg. Ghile et al., 2013; Francois et al., 2018), but most adopt stochastic methods because they provide a greater range of possible future sequences and thus a more exhaustive stress test (Brown and Wilby, 2012). Many stochastic methods are available (see eg., Wilks and Wilby, 1999; Srikanthan and McMahon, 2001), with some implementations specifically designed for use within bottom-up assessments (Steinschneider and Brown, 2013; Guo et al., 2017; Bennett et al., 2021). Key technical issues that need to be resolved when applying these methods include: (i) how to generate different data types with appropriate cross-correlations (for example, precipitation and temperature at a site may be correlated in time); (ii) how to represent spatially correlated behaviour (if the system is heterogeneous); (iii) how to represent climatic fluctuations that occur at lower (multi-year) frequencies; and (iv) if the boundary condition being perturbed is at the hydroclimatic level (precipitation, temperature, etc.), how to convert climatic sequences into streamflow, accounting for non-linear responses. Note that these are not limitations of stress testing itself, but rather are broader issues not specific to stress testing. The multi-faceted nature of these technical issues and tasks can be daunting for new investigators, possibly impeding uptake of bottom-up methods, particularly in budget-

conscious publicly funded research projects. One way to lessen this impediment is to publicly release frameworks that demonstrate how all the parts of the analysis fit together and which can be applied to new systems with relatively minimal effort. This is the motivation of the present article and framework.

Here we focus on the monthly timestep, which differs from the daily timestep used in most bottom-up studies (eg. Whateley et al., 2014; Mukundan et al., 2019; Freeman et al., 2020). The choice of timestep depends on the context. For example, a study of future flooding requires a timestep of daily or shorter because floods typically occur on timescales of days to weeks; however, in many water resource systems, particularly those involving carry-over storage, the dominant dynamic is seasonal or longer (Nathan et al, 2019). In such cases the greater detail provided by a daily timestep may not translate into greater decision-relevant information obtained from the stress test, relative to a longer timestep (Helgeson et al., 2020). Given the significant computational burden of stochastic climate stress testing using water resource models (John et al., 2021c), a timestep consistent with important system dynamics, and no finer, frees up computational resources and allows greater possibilities within the test itself (Wang et al., 2011). These possibilities could include a more extensive test (eg. a greater number of stressors or longer sequences) possibly allowing a more defensible assessment of uncertainty (Helgeson et al., 2020); and/or allowing the analysis to be undertaken without recourse to expensive supercomputing facilities. Recent studies also demonstrate the feasibility of hybrid timestep approaches, where computationally expensive parts of the process are run on longer timesteps and disaggregated to shorter timesteps only where required (John et al., 2021a, 2021c). The suitability of the monthly timestep for water resources studies and the relative lack of methods on this timestep suggests a gap which we seek to fill.

This paper outlines a freely available and integrated framework for generating stochastic data on a monthly timestep and perturbing it for climate stress testing. The framework includes both climatic data and streamflow data; it generates climatic data first and includes conversion to streamflow using a built-in monthly-timestep rainfall-runoff model. System modelling (eg. river operations) is context specific and out of scope but must be on a monthly timestep in order to be paired with this framework. Likewise, site-specific calibration of the rainfall-runoff model is out of scope, although we do outline the methods used in the case study in Supplementary Material Section S5. The framework is intended to be applicable from the scale of small catchments up to large river basins (approximately $10^1$ to $10^6$ km²).

The paper is structured as follows. Section 2 provides a technical description of the framework, while Section 3 describes a case study application on a large water resources system in southern Australia. Section 4 briefly discusses limitations of the framework and issues relating to the appropriate level of complexity and future research directions in reducing computational effort for climate stress testing.

## 2 Framework description

### 2.1 Overview

Figure 1 provides an overview of the framework, moving from historic timeseries, through to stochastic data generation, rainfall runoff modelling and perturbation methods to generate the hydroclimatic data sets required for stress testing. The starting point is the set of historic timeseries data for the area of interest. The user can optionally divide this area into portions, termed *sub-areas*, in which case historic climate data (spatially lumped) must be provided for each sub-area. Subdivision can useful to: (i) create demarcations that are necessary for a given application
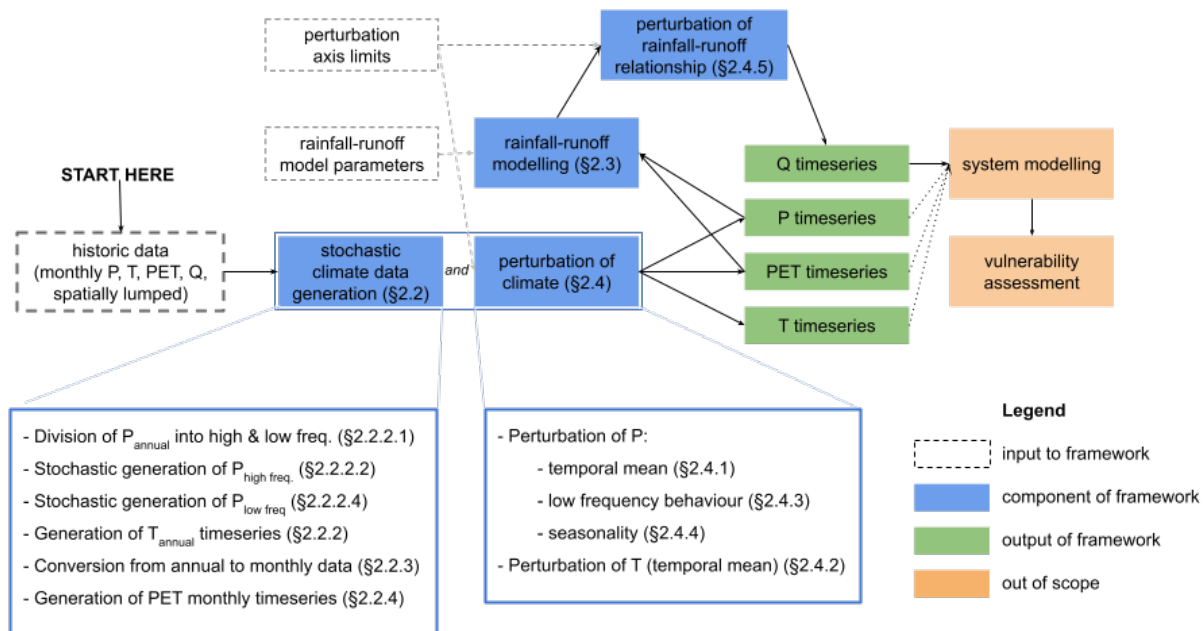
*Figure 1: Flowchart of framework, marked with the relevant sections (§) of this article. Abbreviations: P=precipitation; T=temperature; PET=potential evapotranspiration; Q=streamflow. Note, the steps 'stochastic climate data generation' and 'perturbation of climate' are intertwined in practice (see §2.4.6).*

(eg. upstream versus downstream of a reservoir); (ii) divide a heterogeneous study area into more homogeneous sub-areas; and (iii) capture distinct but correlated behaviour, eg. in different tributaries. All these principles of subdivision are demonstrated in the case study (see Section 3.4). The framework is designed for spatially averaged data across sub-areas, and as such: (i) the user must pre-process raw gauged data to obtain spatial averages prior to input into the framework; (ii) spatial fields are not directly digestible by the framework; spatial averages must be taken to produce monthly timeseries for each sub-area.

The bulk of the framework lies in the generation of perturbed stochastic climate data. Steps are undertaken to first stochastically generate annual timeseries for each sub-area, and later disaggregate to monthly timestep. Although the steps for perturbation of climate are shown in Figure 1 as distinct from the steps for stochastic generation, in practice they are intertwined (Section 2.4.6). The generation of streamflow data is the final step in the framework and has two sub-steps: rainfall-runoff modelling using a monthly rainfall-runoff model, and perturbation of the rainfall-runoff response.

The generated streamflow data can be input into a model of the system of interest, and a vulnerability assessment undertaken. Readers are directed to Brown et al. (2019) or Wilby and Murphy (2019) for guidance in these later steps and also for putting the climate generation into the context of decision making.

As with any stochastic technique, the adopted period of historic data is important because it forms the basis for subsequent stochastic generation (ie. generated data is expected to have similar statistics to this data, see next subsection). Furthermore, perturbations will be relative to this data – eg. a perturbation of -10% will be relative to the mean over the chosen period (strictly speaking, it will be relative to the baseline stochastic series which should have a close but not exact match with the selected period). Thus, users should make every effort to ensure the input data is representative of historical conditions. Usually, longer historic timeseries are preferable as they cover a greater range, and provide a better sample, of past conditions.

## 2.2 Baseline stochastic climate generation

The aim of stochastic data generation is to provide arbitrarily long synthetic timeseries with similar statistical properties to historical data. The relevant statistical properties include averages, variability on different timescales (monthly, annual, multiannual), autocorrelation, and cross correlations between different locations in space. Here, we are particularly concerned with low-frequency (ie. multi-annual or longer) variability and persistence, such as that driven by climate teleconnections (see, eg. Kiem, 2004), because this may cause stress in systems with carry over storage (McMahon et al., 2007), even in the absence of climate change. However, unless special consideration is given to low frequency behaviour, stochastic schemes tend to miss this dynamic entirely (Thyer and Kuczera, 2000; Srikanthan and McMahon, 2001). Therefore, as discussed in the next subsection, we adapt an earlier scheme (McMahon et al., 2008) which explicitly separates out low frequency behaviour in precipitation and generates it using a separate stochastic scheme to that used for the high frequency component.

Please note:

- the process of stochastic climate generation begins with P. Later, T and PET are generated as a function of P, based their historic correlations. Thus, P is described first (Section 2.2.1), then T (2.2.2) and PET (2.2.4).
- Although the final timestep of the stochastic data is monthly, many of the operations are done on a longer timestep which is subject to later disaggregation (Section 2.2.3). The annual timestep is the default choice for the longer timestep, and this article is written assuming this choice, but alternative choices are possible, as discussed in Section 2.2.2.1.

### 2.2.1 Generation of annual precipitation (P) timeseries

Figure 2 provides an overview of the method for stochastic generation of annual precipitation. The following subsections explain each step:

- Section 2.2.1.1 outlines the process of splitting the precipitation into high and low frequency components;
- Section 2.2.1.2 describes the stochastic generation of the high frequency component;
- Section 2.2.1.3 describes the model used for stochastic generation of the low frequency component; and
- Section 2.2.1.4 focusses specifically on how the low frequency model parameters are chosen to preserve relevant historical statistics, and how this is tailored to each individual sub-area.

#### 2.2.1.1 Division of annual precipitation into high and low frequency components

The tendency for low frequency dynamics to be missed by common stochastic generation techniques has prompted the development of specialized schemes to overcome this limitation (eg. Thyer and Kuczera, 2000; McMahon et al., 2008; Steinschneider and Brown, 2014). Here, we adapt McMahon et al. (2008)'s scheme to explicitly define and separate low frequency behaviour using a variant of the Empirical Mode Decomposition technique, and stochastically generate the high and low frequency components separately. Note, "low frequency" in this context means oscillations which occur with period of approximately ten years or longer, but it is noted that the distinction between high and low frequency is largely under the user's control, as discussed below.

Empirical Mode Decomposition (EMD, Huang et al., 1998; Torres et al., 2011) is a non-parametric technique capable of analyzing non-stationary timeseries such as those stemming from non-linear systems. The EMD process iteratively removes variation from a time series,

highest frequencies first, as shown in Figure 3. In each iteration, the removed variation is called an 'intrinsic mode function', or IMF, and each IMF has a mean of zero. The sum of all IMFs plus the residual gives the original time series. The advantages of using EMD in hydroclimatic contexts include that it *"can handle both nonlinear and nonstationary time series. This is a major advantage over other decomposition techniques, such as Fourier analysis or wavelets"* (McMahon et al., 2008).
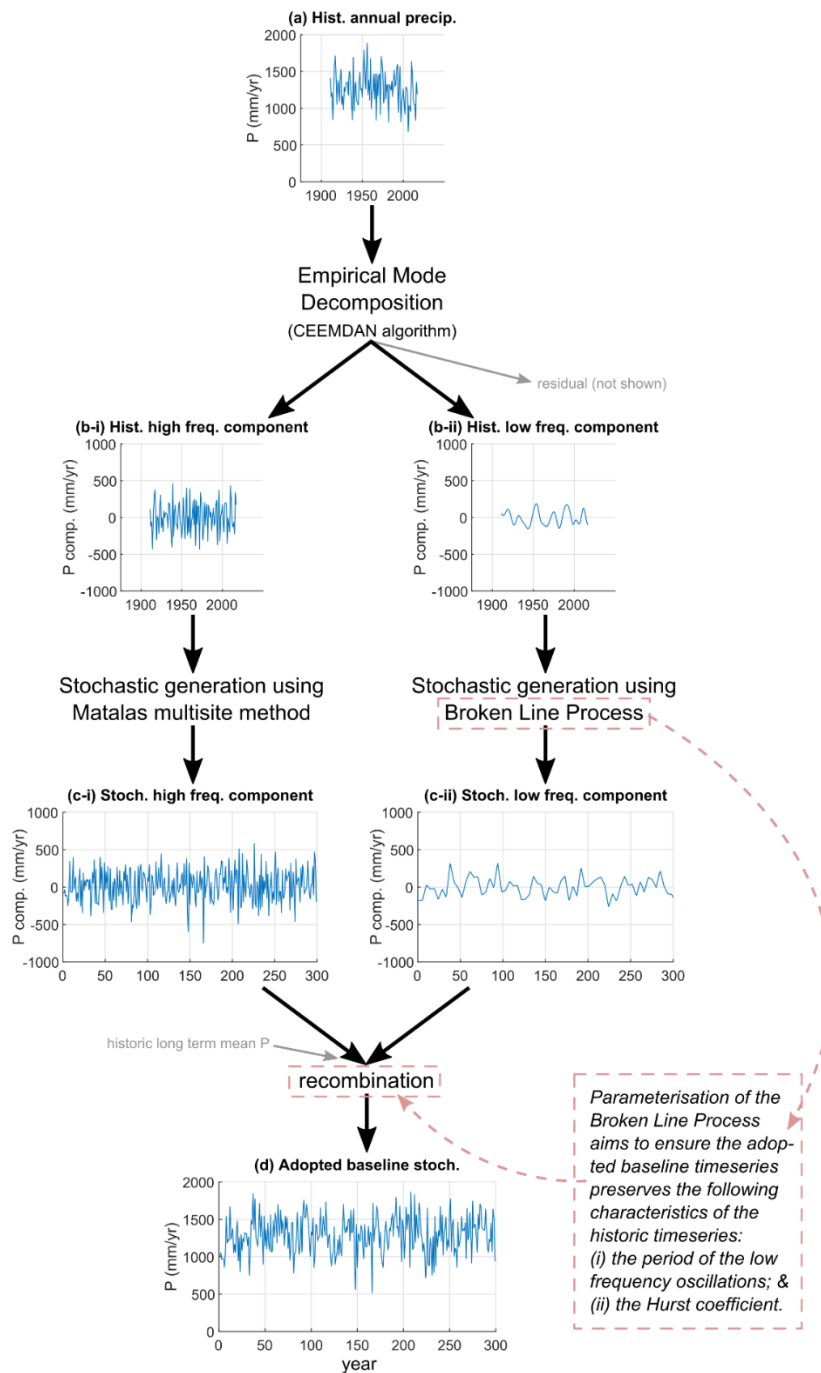


*Figure 2: Steps involved in generation of stochastic annual precipitation timeseries in cases with only one sub-area. In cases of multiple sub-areas: (i) CEEMDAN is applied to each sub-area separately; (ii) the Matalas multisite method ensures spatial cross correlations are maintained between sub-areas; (iii) the historic low frequency is a weighted average of all the sub-areas; and (iv) a single Broken Line sequence is adopted across all sub-areas, but the amplitude of this sequence is set to ensure the historic Hurst value is preserved in each sub-area.*
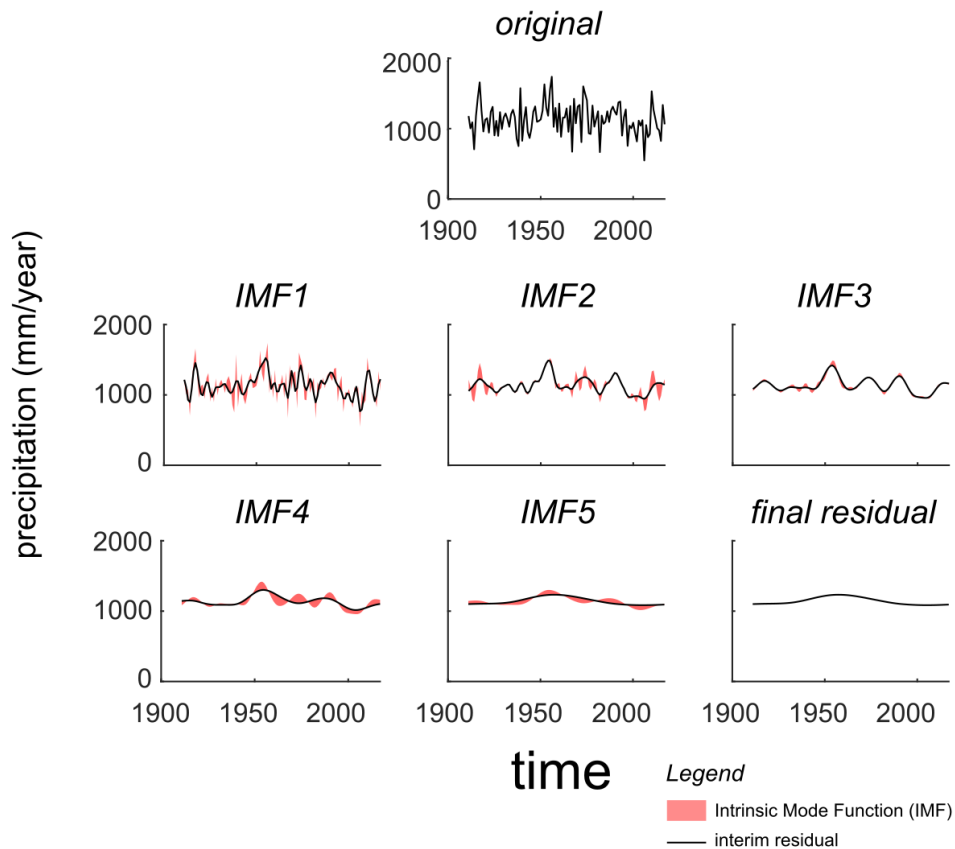
*Figure 3: Application of Complete Ensemble Empirical Mode Decomposition with Adaptive Noise (CEEMDAN) to annual precipitation for an example location, showing Intrinsic Mode Functions (IMFs) and residuals.*

Instead of the original EMD algorithm, we adopt *Complete Ensemble Empirical Mode Decomposition with Adaptive Noise,* or CEEMDAN (Torres et al., 2011; Colominas et al., 2014, http://bioingenieria.edu.ar/grupos/ldnlys/metorres/metorres_files/ceemdan_v2014.m, last accessed 22-Oct-2021) and incorporate its code into the framework. CEEMDAN has multiple additional features over the original EMD and subsequent Ensemble EMD (EEMD) that improve the robustness of IMF extraction, reduce information sharing (referred to as "mode mixing") between IMFs and better ensure the set of decomposition products sum exactly to the original series (unlike EEMD) (Colominas et al, 2012; 2014).

Here, categorisation into 'high' and 'low' frequency involves lumping the first *n* IMFs and categorizing these as high frequency, where *n* is a user-specified parameter. The remainder are then lumped to form the low frequency component. For example, McMahon et al. (2008) choose *n* so that 'high' captures all *intra*-decadal oscillations, whereas 'low' captures *inter*-decadal oscillations. Deciding *n* is subjective; the user needs to consider what might cause stress and expose vulnerabilities of the system, which depends on the strength of different modes of variation in observed climate and also on the ability of the system to buffer against these variations.

It is important to note that EMD and CEEMDAN have limitations which may be particularly relevant in systems where El Niño-Southern Oscillation (ENSO) variation is the dominant multi-year dynamic. ENSO oscillations typically take between four and seven years, and in the shorter case the oscillations may be too rapid to be properly characterised. This is because

CEEMDAN (along with other EMD-like algorithms) works by identifying local minima and maxima in timeseries, and ideally a sequence of at least five points (mid-low-mid-high-mid) is required for CEEMDAN to resolve the cycle. In cases where ENSO dynamics are expected to expose important system vulnerabilities, a shorter aggregated timestep may be required, such as the six-month timestep adopted by McMahon et al. (2008). Most of the framework method is agnostic to this choice, except that the disaggregation step (Section 2.2.3) would require extra care – for example, if the 6-month periods were Nov.-Apr. and May-Oct., historic patterns from Nov.-Apr. should not be 'donated' to May-Oct., nor vice versa.

After applying CEEMDAN to each sub-area separately, the high frequency components and low frequency components are treated differently. The high frequency components are stochastically generated such that each sub-area is distinct but spatial cross-correlations are maintained, as discussed in the next subsection. In contrast, the low frequency components are aggregated together and largely treated as one, with a specialized method of stochastic generation, as discussed in Section 2.2.1.3 and Section 2.2.1.4.

### 2.2.1.2   Stochastic generation of high frequency component

The high frequency component of annual precipitation is stochastically generated using the multi-site method of Matalas (1967). This method, described as 'weakly stationary' by the original authors, preserves lag-1 autocorrelation at each site individually, in addition to the lag-0 cross correlations between sites (note, in this context a site means a sub-area). For more detail, the reader is directed to the original paper (Matalas, 1967) and subsequent reviews such as Srikanthan and McMahon (2001). In cases where the user elects not to divide the study area into sub-areas but rather treat it as a single entity, the Matalas method is the same as a lag-1 autocorrelation (AR(1)) model.

As with many statistical techniques, the performance of the Matalas method improves if the input data are approximately normally distributed. To achieve this, the framework provides the option to normalize the high-frequency component data prior to the Matalas step, using a Box-Cox transformation (Box and Cox, 1964). The Box-Cox parameter is chosen on a per-sub-area basis, to minimize the absolute value of the skew. Post generation, the same parameter is used to reverse the transformation.

### 2.2.1.3   Stochastic generation of low frequency component: adopted model

The earlier demonstration of a similar workflow by McMahon et al. (2008) adopted an AR(1) model for generating the low frequency component of precipitation. However, as shown in Figures 4a and 4b, initial testing revealed that the AR(1) model was unable to replicate the dynamics of the low frequency component of our test site, even in the case where the generated timeseries has similar standard deviation and autocorrelation to the historic low frequency timeseries. The AR(1) model tends to exhibit sustained departures from the mean (in this case zero), resulting in a small number of inappropriately long spells, seen as a "fat tail" in the distribution in Figure 4d-i. To make matters worse, when the AR1 timeseries happens to be close to the mean, its tendency for short-term fluctuations leads to a large number of very short spells above and below the mean, as seen by the high frequency of one-to-five year spells in Figure 4d-i.

For these reasons, we instead adopt a Broken Line Process (Mejia et al. 1972) for stochastic generation of the low frequency component. As the name implies, a Broken Line Process is simply a set of connected linear segments, with an equal (usually multi-year) length of time assigned to each segment, and the value at each transition point being selected randomly from a normal distribution with zero mean. The two parameters of the Broken Line Process are the
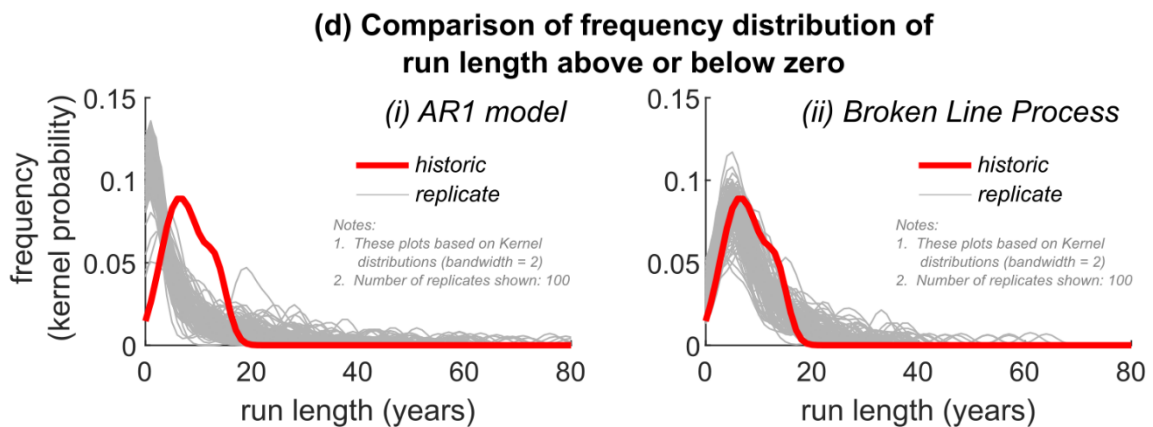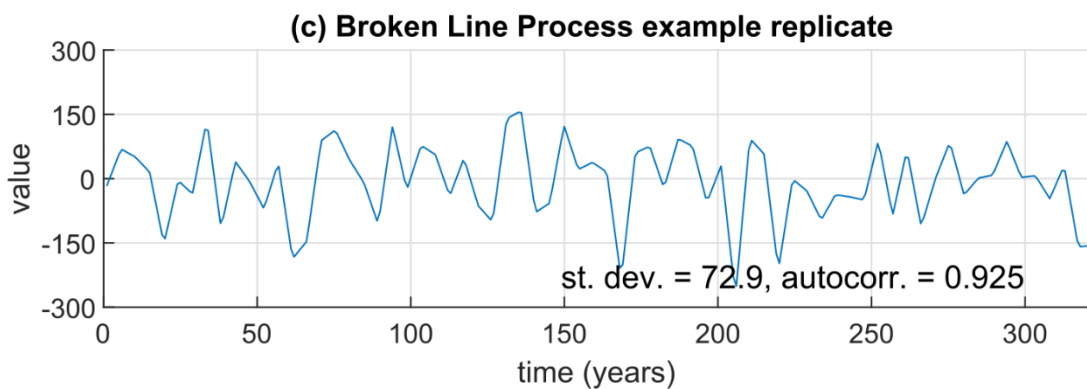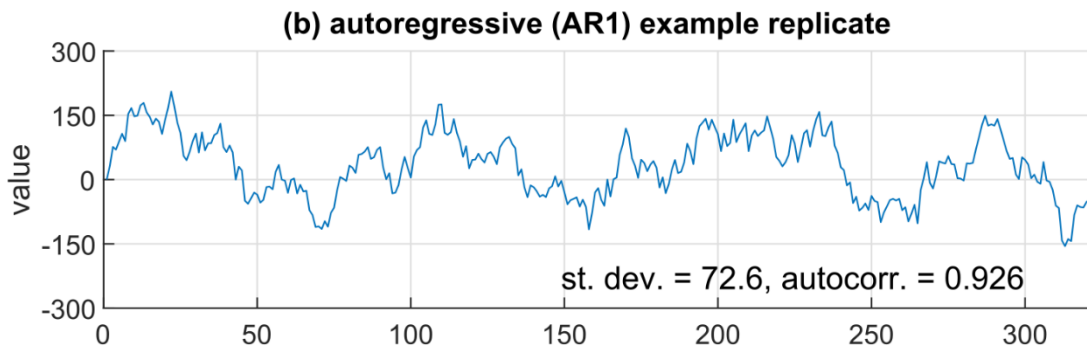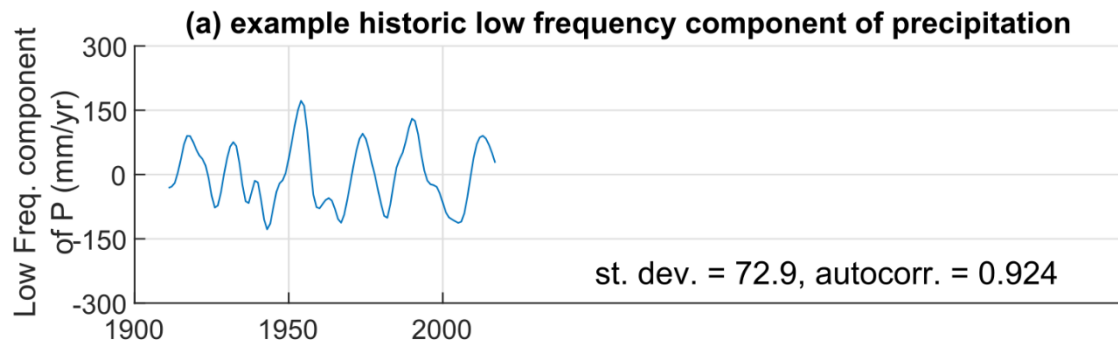
*Figure 4: Comparison of example fits for AR1 and Broken Line Process models against historic annual timeseries: (a) low frequency component of an example historic series; (b) example output of an autoregressive model, parameterised so that the standard deviation and autocorrelation match historic values; (c) example output of a Broken Line Process, parameterised in a similar fashion; (d) differences in run length distribution for the models in (b) and (c), where a run can be either above zero or below zero. Plot (d) quantifies these differences using 100 stochastic replicates for each model.*

segment duration and the amplitude (standard deviation) of the process. Both parameters can take any positive real number. The values at a given transition point are independent of all the others (ie. zero autocorrelation), but because the segment duration is usually multiple years, the annual timeseries exhibits non-zero autocorrelation (Figure 4c). Importantly, in this initial test, a Broken Line Process parameterised to match the standard deviation and autocorrelation matched the run length distribution much more closely than the AR(1) model did (ie. the most frequent run length is similar to historic, and the tail is less "fat" than the AR(1) (Fig. 4d-ii).

In concept, the Broken Line Process is related to the Broken Line Model (Garcia et al., 1972); specifically, a Broken Line Model is composed of multiple Broken Line Processes added together, each operating at different frequencies. However, given the promising replication of statistics including run length, a single process is considered sufficient for the present purpose. That said, if users wish to swap the Broken Line Process for a different model (whether the Broken Line Model or another long memory model) this would be straightforward to code.

### 2.2.1.4 Stochastic generation of low frequency component: parameterisation and treatment across different sub-areas

As noted, most aspects of the treatment of the low frequency component of precipitation are conducted in common across all sub-areas. For model fitting, the historic low frequency component timeseries for all sub-areas are combined into a single timeseries using a weighted areal average (or optionally, a bespoke user-defined weighting). The same segment duration parameter is adopted across all sub-areas, calculated from the number of zero crossings in the combined historic timeseries, using the formula in Supplementary Material S1. Also, the same sequence of random deviates is used across all sub-areas so that the generated low frequency timeseries for a given sub-area varies from the others only in amplitude, but not in timing; ie., all sub-areas move together into and out of multi-year wet and dry periods. The justification for this is the assumption that the processes driving low frequency oscillations, such as climate teleconnections, are acting on all sub-areas in common across the study area.

The amplitude parameter is chosen for each sub-area separately so that the observed historic value of the Hurst coefficient (Hurst, 1951) is replicated in the baseline precipitation data. The Hurst coefficient measures the persistence in the timeseries; high persistence means consecutive values are more dependent upon one another (Koutsoyiannis, 2005) and thus multi-year droughts will tend to be longer, as will multi-year runs of wet conditions. Note, during perturbation the Hurst coefficient value can be directly perturbed as one of the five stressors in the stress test, in which case the perturbed value will be preserved rather than the historic (see Section 2.4.3).

The following two clarifications may help to avoid confusion:

1. For a given sub-area, the two values being matched are: (i) the Hurst value of the historic precipitation timeseries (ie. the raw observed data, _not_ the low frequency component); and (ii) the Hurst value of the stochastic timeseries after the high and low frequency components have been added together. Since the high frequency component is essentially just (red) noise, the persistence of the final timeseries will tend to increase with increases in the amplitude assigned to the low frequency component. In short, higher amplitude means higher Hurst coefficient.

2. Although the synthetic replicates may be hundreds or thousands of years long, it is not wise to calculate Hurst coefficient values over this duration because the calculations become sensitive to low frequency oscillations whose period exceeds the historic record. To sidestep this issue, the stochastic timeseries is divided into segments of equal length to the historic record. For example, if the stochastic replicate is 2000 years long and the historic

record 100 years, the replicate is divided into twenty fragments, and a separate Hurst value calculated for each fragment. The key value is then the *average* Hurst value among this set of twenty values, and the solver seeks to match this value to the target Hurst value.

After the application of EMD, Matalas and Broken Line methods, the result is:

- Each sub-area has its own stochastic timeseries of annual precipitation;
- The sub-areas move together through multi-year periods of wet and dry conditions;
- For a given sub-area, the historic Hurst coefficient (or a perturbed value, if relevant) is replicated, assessed over sub-periods of the same length as the historic data;
- The historic autocorrelations and cross correlations between sub-areas are typically close, but this is not exact. The Matalas method can exactly replicate these properties, but the Matalas method governs the high frequency component only. The low frequency component method focusses on ensuring Hurst values are preserved in the combined (high+low) synthetic timeseries, and this may result in the cross correlations and autocorrelations varying from historic values. Generally, these departures are small.

### 2.2.2   Generation of annual temperature (T) timeseries

Figure 5 provides an overview of the remainder of the stochastic generation process. The first step is generation of annual T timeseries for each sub-area, which is discussed in the next paragraph. We note the following two points: (1) the temperature information can be based on daily maximums, daily minimums, or daily means, in line with user preference; and (2) the process used for T (ie. developing a relationship with P based on historical observations and then using this to create synthetic T from synthetic P) is relatively generic and could be used to generate any factor that is dependent upon P.

Annual T is generated based on the stochastic timeseries of annual P and the historic relationship between T and P (assumed to be linear). Often, there is considerable scatter in the historic relationship, so that for a given P value there is a distribution of T values, as shown in Figure 5. The residuals typically exhibit spatial cross-correlation and autocorrelation. To represent this behaviour, the first step is to create a single timeseries of autocorrelated random numbers of zero mean, with the autocorrelation set to the average across all sub-areas (ie. the autocorrelation of the residuals is calculated in each sub-area separately, and then these numbers are averaged). Then for each sub-area: (i) an initial timeseries of T is generated using the historic line of best fit; (ii) the random number sequence is scaled to have the same variance as the residuals of the historic relationship, for that sub-area; and (iii) the two timeseries are added together to give the final timeseries of T. It is recommended that the user check the normality of the historic residuals before applying this process. This process assumes the cross-correlation in the residuals is perfect; it was considered that this was an appropriate simplification given that the cross correlations in P are preserved in the Matalas process (see above) and P is the basis for T. Also, this simple scheme for T is considered appropriate for water resources systems as they are generally less sensitive to T than to P.

### 2.2.3   Conversion from annual to monthly data

Conversion from annual to monthly timestep is achieved by the method of fragments (eg. Srikanthan and McMahon, 2001) for both P and T. We note that, if desired, other methods could also be incorporated into the framework by future users, such as the non-parametric k-nearest neighbour (kNN) bootstrap sampling approach of Lall and Sharma (1996). For each synthetic year, the method of fragments randomly selects a year from history to 'donate' its pattern. For P, the donated pattern is the proportion of precipitation in each month, whereas for T, the donated pattern is the variation of monthly T (in °C) about the average for that year.
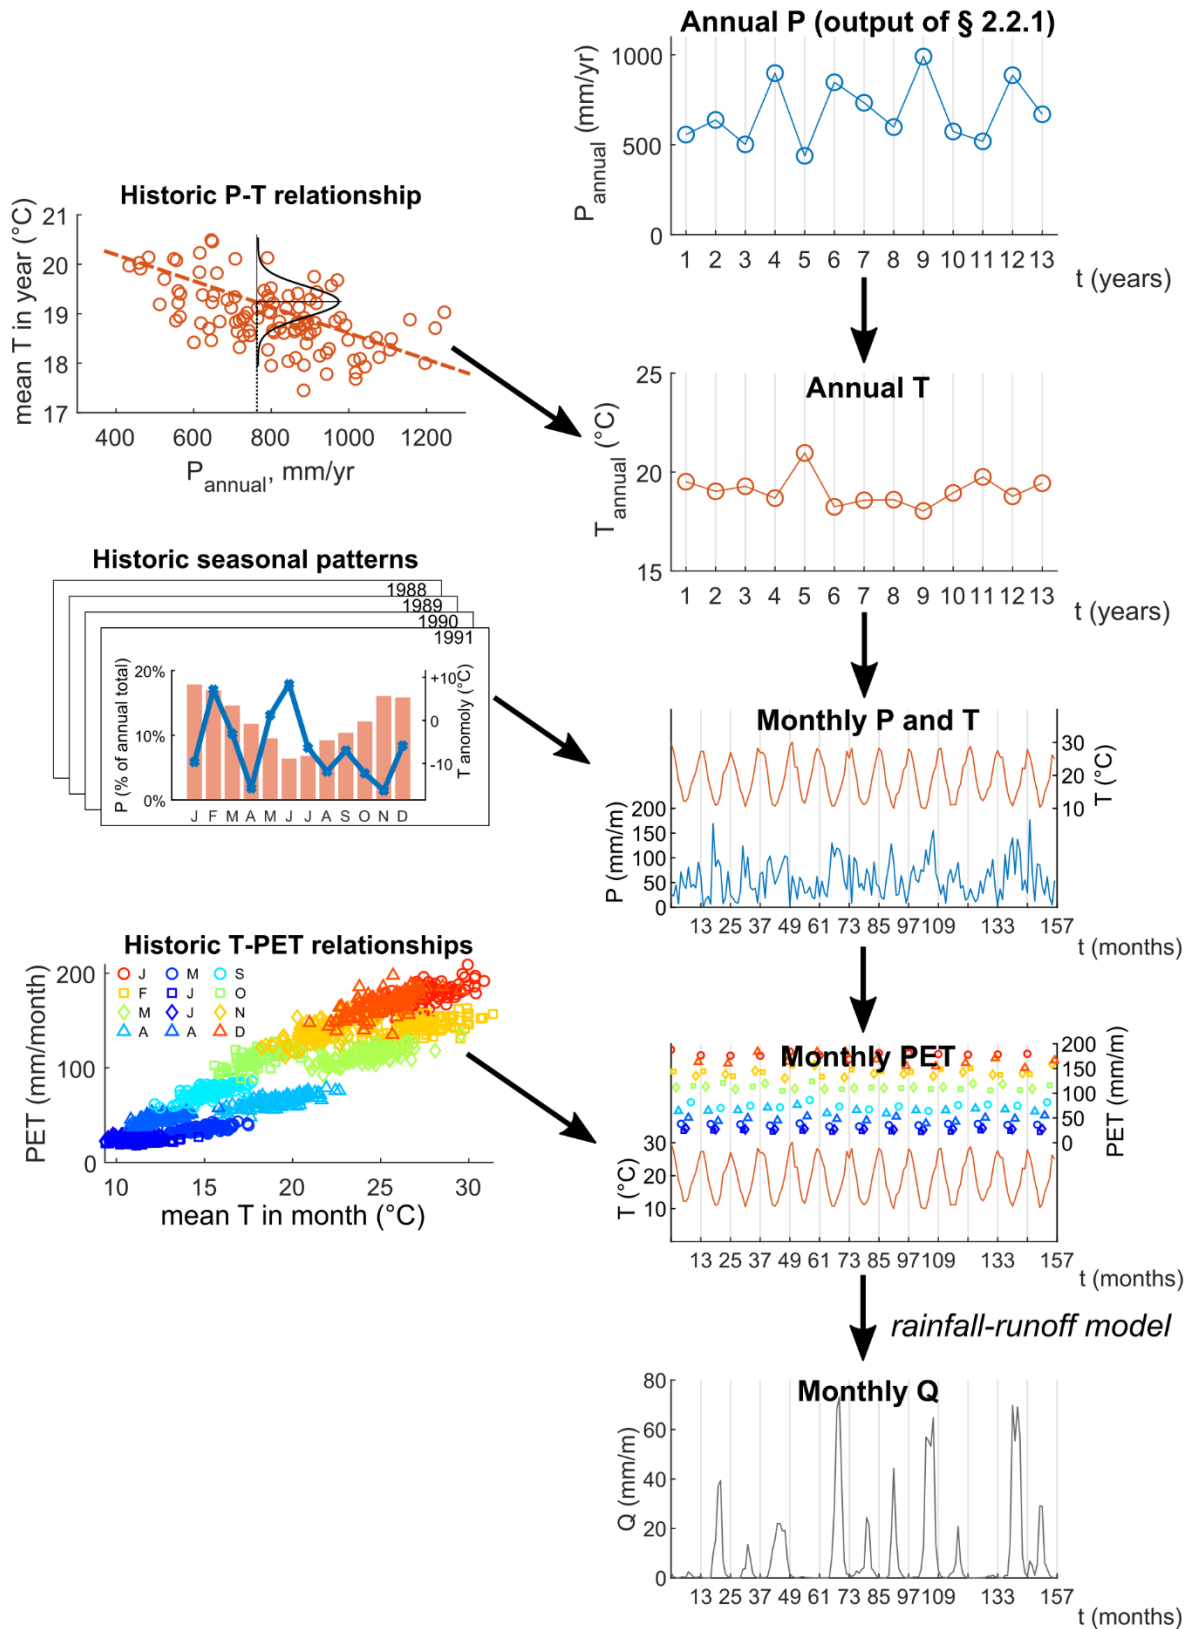
*Figure 5: Diagrammatic representation of framework steps to generate baseline (no perturbation) stochastic monthly data for all data types (P, T, PET and Q), starting from the synthetic annual timeseries of precipitation (ie. output of Section 2.2.3).*

For a given synthetic year, the same historical year is donor across all sub-areas. This maintains historic spatial cross correlations in monthly data.

Realism is increased by restricting potential donors to historical years with similar hydroclimatology. This is achieved by categorizing historical years as 'dry' (driest 33.3% of years, by P), 'wet' (wettest 33.3%) or 'medium' (remainder). Synthetic years can only be assigned a pattern from a historical year in the same category. Note, after calculating threshold values between categories based on historic data, the threshold values remain unchanged throughout. In the case of perturbation (Section 2.4) towards, eg., lower P, more than 33% of synthetic years will be categorized as 'dry', so that the perturbed timeseries will use patterns from historical dry years more often than historical wet years in this example case.

### 2.2.4 Generation of monthly PET timeseries

Monthly PET is generated via linear regression with monthly T. For some locations, it may be acceptable to use a single relationship across all calendar months, but for others (including our test case) there are considerable seasonal differences in the T-PET relationship from month to month – for example, the slope in the coolest month (July) is much lower than the slope in January (Figure 5). Thus, for each of the 12 calendar months, a separate linear regression is used, for each sub-area separately. For our test case, most of the variability in month-to-month PET is explained by the calendar month, so it is not considered necessary to account for scatter in the relationship for a given month, and no special consideration of spatial cross correlation is given, aside from applying these simple deterministic relationships. For applications of the framework in locations with different ET dynamics, users may code a different approach such as by adding an error model like the one described in the previous sub-section (see McMahon et al., 2015, including the supplement, for general discussion).

## 2.3 Simulating streamflow

Monthly Q is generated using the WAPABA model (Wang et al., 2011), which is built into the framework. Other monthly rainfall-runoff models could also be used if the user is willing to add the required code (see eg. Xu et al., 1998 or Topalović et al., 2019). WAPABA was selected from among a set of three tested monthly models because its performance over varied climatic conditions was superior (not shown) for the study area of the project that funded framework development (the case study in Section 3).

WAPABA has five parameters: two govern the partitioning of rainfall between streamflow, evapotranspiration and changes in storage ($\alpha_1$ and $\alpha_2$); one determines the total capacity of the soil moisture store ($S_{max}$); and the final two determine the quick/slow flow split and drainage rate of slow flow ($\beta$ and $K^{-1}$ respectively). Further information is provided in the Supporting Material, Section S4. WAPABA inputs are monthly P and monthly PET. Appropriate model parameters are assumed to be known, which is why they are listed as an input in Figure 1. Further information on model calibration is provided in the Supplementary Material S5.

A key difficulty is that the sub-areas adopted for stochastic generation may span multiple catchments and/or be only partially gauged. For example, a single sub-area may contain parts of multiple river catchments – an example is shown for the case study in this paper (see Section 3 and associated figures). The suggested approach depends upon the context:

- If there are generally many gauged catchments per sub-area, and the problem is that no single gauge captures the whole sub-area, then it is suggested to identify a single gauged catchment within the sub-area that is considered suitable to act as a 'representative' catchment for the whole sub-area. The simulated streamflow from the representative

catchment can be used in the way that best suits the case study, but the simplest method would be that the simulated streamflow outputs, in units of depth (mm/month), could be assumed to apply over the whole sub-area. For an example of a more nuanced approach, see the case study in Section 3. Note: since synthetic Q is generated using sub-area-wide forcing data (P and PET), model calibration should also use sub-area-wide historic forcing data (P and PET) (see Supplementary Material Figure S4).

- If sub-area streamflow is ungauged, the problem becomes a *prediction in ungauged basins* problem and an appropriate regionalisation technique must be applied to decide appropriate parameters (see, eg., Bloschl et al., 2015; Hrachowitz et al., 2013).

## 2.4   Perturbation

To *perturb* means to change an aspect of hydroclimate ("stressor") for the purposes of testing the impact of the change on system performance. For example, the stressor "mean temperature" might increase by varying amounts, or the stressor "mean precipitation" might go up or down.

Detailed guidance on constructing a stress test experiment is beyond the scope of this paper, but in summary, the stressors can each be considered as individual "axes" in a stress testing "space" (eg. Prudhomme et al., 2010). While it is possible to limit the testing regime so that it focusses on one axis at a time, it is more insightful to test combinations of changes. This enables one to assess system performance based on any interactions between the climate stressors. To achieve this, bottom-up studies typically cover the space via uniform sampling, producing a regular grid. The underlying stochastic generation framework should, in principle, be able to generate data corresponding to any point in this space – that is, any combination of stressor perturbation. Thus, the adopted technical method for each of the five stressors must be applicable even if the other stressors are also being simultaneously perturbed, and this is a key technical challenge for this framework.

In climate stress testing studies, it is common to include only two stressors, usually: (1) changes in mean precipitation; and (2) changes in mean temperature (eg. Prudhomme et al., 2011; Turner et al., 2014; Whateley et al., 2014; Culley et al., 2016; Francois et al., 2018; Freeman et al., 2020). However, the rapid nature of the framework (due mainly to its monthly timestep) allows greater possibilities to expand the types of perturbation considered. Thus, we have included an additional three stressors: changes in low frequency behaviour of precipitation; changes in precipitation seasonality; and changes in rainfall-runoff relationship. The rationale for each stressor is given in the subsections below.

In general, we suggest that the primary factor in selection of stressors should be to identify possible vulnerabilities of the system. Although stress testing studies commonly overlay GCM projections to understand future risk, this does not mean that the selection of stressors should be limited to those that can be simulated by GCMs. For example, the hydrological processes relevant to changes in rainfall-runoff relationship are considerably out of scope for GCMs, but this stressor is still system relevant (for our case study – Section 3) and thus should not be omitted.

The five stressors included in the framework have been selected based on relevance to the project that funded the frameworks development (Section 3), but we feel they may also be relevant in many other contexts. However, if they are considered irrelevant in some locations, it is trivial to remove one or more stressor by simply ensuring values of perturbation are always zero for that stressor. Conversely, it is possible in principle to add other stressors to the existing set if this is required for an application of the framework, provided the user is willing to alter the framework code.

Note, in this framework the perturbations are applied to the baseline data without first enforcing that baseline data statistics exactly match historic statistics (eg. averages). This is in line with the expectation that stochastic data should have similar, but not identical, statistics to historic. Thus, strictly speaking, a given perturbation is considered relative to the baseline stochastic series rather than the historic timeseries. In practice, the baseline timeseries generally match historic statistics well, as can be judged by the figures provided with the case study (Section 3). However, this is usually related to the length of the generated timeseries, with longer synthetic timeseries more likely to provide a closer match to historic statistics due to the averaging out of short term fluctuations.

Following the subsections on individual stressors, a summary of the code's order of operations is given (Section 2.4.6), to clarify which perturbations are applied first. Finally, guidance on perturbation axis limits is provided in Section 2.4.7.

### 2.4.1 Stressor 1: precipitation, temporal average

**Rationale:** This is a first order control on hydrologic response in water systems, and therefore on system performance. Thus, Stressor 1 is routinely included in climate stress tests, as noted.

**Method:** Perturbations in temporal average precipitation are applied as a proportional change that is constant across the synthetic timeseries. For example, if the perturbation value is +5%, every value in the precipitation timeseries is multiplied by 1.05; if the value is -20%, every value in the precipitation timeseries is multiplied by 0.8.

### 2.4.2 Stressor 2: temperature, temporal average

**Rationale:** Since T is regarded as the most reliable GCM variable, and since T directly influences evaporative demand, it is very commonly included in climate stress tests. However, the influence of this stressor on system performance may vary depending on the system and the aspect of performance in view, and on whether the system is water or energy limited.

**Method:** Perturbations in temporal average temperature are applied as an absolute change, in °C. For example, if the perturbation value is +1.5, every value in the synthetic temperature timeseries is increased by 1.5.

Note, perturbing stressor 1 or stressor 2 means historic relationships between P and T will cease to be preserved. As noted in Section 2.2.2, the origin of the unperturbed T timeseries is the unperturbed P timeseries which is then transformed into T via the historic regression relationship (subject to noise). Perturbing stressor 1 and/or stressor 2 causes this relationship to shift vertically, laterally, or both, in the stochastic data relative to the historic relationship.

### 2.4.3 Stressor 3: low frequency behaviour of preciptiation

**Rationale:** This is important to system performance in carryover systems as outlined in Section 2.2 above, particularly where the main considerations are reliability of water supply. Future changes in the characteristics (amplitude, period) of this behaviour may expose system vulnerabilities even in the absence of a change in average precipitation. There is much diversity among GCMs regarding projected future changes in low frequency oscillation such as ENSO (eg. Bellenger et al., 2014), but in general GCMs project an increase in the frequency of strong El Niño and La Nina events, but disagree about the strength of this change (Cai et al., 2018). Few studies have examined projections for longer (interdecadal) cycles (for an exception see Wei et al., 2017), but (like ENSO) these oscillations are sensitive to ocean dynamics (Liu et al., 2012) which in turn respond to climate drivers, making future changes possible.

**Method:** As briefly mentioned in Section 2.2.1.4, perturbations of the low frequency component are quantified as changes in Hurst coefficient. Recapping 2.2.1.4, for baseline stochastic data (no perturbation) the amplitude parameter in the Broken Line process is tuned to ensure the historic Hurst coefficient is matched (higher amplitude means higher Hurst coefficient). For stressor 3, the Hurst value itself is perturbed – that is, the target Hurst is higher or lower than the historic Hurst value, in line with the perturbation. To achieve the perturbed Hurst value, one possible approach would be to retune the amplitude parameter without altering the segment duration parameter. This would result in oscillations of similar period but different magnitude to historic. However, since in reality the Hurst coefficient is affected by both the period and magnitude of low frequency oscillations, we allow both parameters to vary in response to perturbation, according to the method given in Supplementary Material Section S3. Note, stressor 3 perturbation is done early in the process, before T and PET are generated from P (see Section 2.4.6 and Figure 4). Thus, both T and PET exhibit low frequency behaviour in response to the low frequency behaviour of P.

### 2.4.4 Stressor 4: precipitation, seasonality

**Rationale:** The importance of precipitation seasonality depends on context. It may be less important in some carry-over storage systems, as discussed above. For projects concerned with stream ecology (see case study, Section 3) seasonality of flows may be a core concern for species diversity (Tonkin et al., 2017) particularly since some species are sensitive to flows at certain times of year as, eg. breeding triggers (eg. Poff, 1997). Changes in seasonality may therefore have significant impacts on ecology, particularly in unregulated systems. Also, depending on the synchronicity of precipitation and demand, changes to precipitation seasonality may also impact flow volumes (eg. if rain is shifted towards months of higher evaporative demand, less streamflow will ensue).

**Method:** The seasonality stressor is based on two seasons (warm and cold) relevant to the temperate climate in the example case study. This may be easily adapted to more tropical climates with wet and dry seasons. The user needs to decide which months are in which season (to allow for austral and boreal differences in seasonality). Perturbation values are quantified as an increase in the proportion of annual P that falls in the warm season (with negative values denoting decreases). Thus, in the case where the unperturbed P is 900 mm with a warm:cold split of 300:600, a perturbation of 5% means 45 mm will be subtracted from the cold season and added to the warm season, so that the new split is 345:555. Note, when combined with perturbation in other stressors, unexpected results can occur – for example, a small positive perturbation in stressor 4 combined with a large negative perturbation in stressor 1 will give a decrease in warm season P, even though an increase would be expected from the seasonality perturbation acting alone. Perturbation of seasonality is achieved through three steps:

1. *Conversion of perturbation value into a reallocation amount in mm.* For example, if the perturbation value was 0.08, and the perturbation was being applied to a timeseries with mean annual precipitation of 1000 mm/yr, this means 80 mm/yr on average will be reallocated from winter to summer months.

2. *For each calendar month, assign reallocation amounts (in mm) to each calendar month.* The reallocation pattern is defined as described in Supplementary Material S2.1, accounting for the seasonal definitions specified by the user. If the provided pattern does not suit an application, it can be easily redefined as required. Since it is a redistribution, some months will gain and some will lose, with a net change of zero.

*Note:* following this step, one possible option would be to simply add (or subtract as required) the reallocation amounts directly to the timeseries. For example, if the result of step 2 for

January were +17 mm, then we could simply add 17 mm to all the January months in the synthetic timeseries. This would mean that there are no longer any dry January months in the perturbed timeseries, as the minimum possible value would be 17mm – an undesirable outcome. For this reason, step 3 is based on an alternative approach using adjustment factors.

3. *Calculate monthly adjustment factors based on reallocation amounts.* For each calendar month, calculate an adjustment factor such that, when all instances of the month are multiplied by the factor, the average change equals the absolute change from step 2 (see Supplementary Material S2.2). This multiplicative approach preserves dry months in the original timeseries to a much greater degree than an additive approach, but it causes a small water balance error in most years, which is corrected by simply scaling all months in the year up or down by the required amount (rarely more than 2% in our test case).

### 2.4.5 Stressor 5: relationship between rainfall and runoff

**Rationale:** Recent studies in Australia (eg. Saft et al., 2015) and the USA (Avanzi et al., 2020) demonstrate that it is possible for a multi-year dry period to cause a shift in rainfall-runoff response. This means that, even after we take the lower rainfall into account, the streamflow during the multi-year dry period is still lower than expected. A climate-change induced permanent reduction in future precipitation could result in a permanent shift in rainfall-runoff relationship, and because the processes causing the shift are poorly understood, the shift is not possible to predict in advance, which is why it is treated here as an stressor in its own right. It could be argued that such measures are redundant since we should expect that hydrological models can simulate these transitions. Although some hydrological models can do so if calibrated jointly to the 'before' and 'after' (Fowler et al., 2016), studies in Australia indicate that these models could not have anticipated the transition in advance of it occurring – that is, if calibrated to the 'before' only, they do not accurately simulate the 'after' in independent evaluation (Fowler et al., 2018). Until these issues are resolved, a separate stressor may be justified to explicitly test different plausible future shifts in rainfall-runoff relationship.

**Method:** Simplistically, a possible implementation for stressor 5 might be to multiply every streamflow value by a set factor, like for precipitation in stressor 1. However, this is inconsistent with observed behaviour during Australia's Millennium Drought: namely, the observed proportional impact was much greater for dry years than wet years. The framework replicates this behaviour by modifying methods from Saft et al. (2015), as discussed below.

Saft et al. (2015) generated scatter plots between annual P (x axis) and annual Q (y axis), where the latter was transformed to linearise the relationship. A shift in rainfall-runoff relationship means a vertical shift on this P-Q plot, with every year shifted by the same amount in transformed space. However, the effect of reversing the transform means that the actual change varies depending on the flow value.

For the transform, Saft et al. (2015) applied a one-parameter Box-Cox transformation (Box and Cox, 1964), which is adopted here also. However, whereas Saft et al. (2015) allowed a different Box-Cox $\lambda$ value for each catchment, here a single value is assumed to apply across all sub-areas. Perturbation axis limits and gradations are then defined within the Box-Cox transformed space. One way to choose the $\lambda$ value, demonstrated in Supplementary Material S7, is to minimise the skew across all sub-areas simultaneously.

Practically, the steps undertaken to transform a given synthetic sequence, given a negative perturbation 5 value $-p_5$, (ie. corresponding to a decrease rather than an increase) are as follows. These steps are shown graphically using the case study data later in the paper (see Section 3). For each synthetic year:

1. Calculate the total flow for the year (ie. sum the 12 monthly simulated flow values)
2. Apply Box Cox transformation to the output of (1)
3. Perturb by subtracting $p_5$ from the output of (2)
4. If the output of (3) is negative, set it to 0. In the perturbed series it will be a zero flow year.
5. Otherwise, apply reverse Box Cox transformation on the output of (3)
6. Disaggregate the output of (5) to monthly using the original flow pattern from (1)

### 2.4.6 Perturbation order of operations

The order in which the perturbations take place is different to the order in which they are presented above. To avoid potential confusion, we explicitly list the order of operations:

1. Generate high frequency component of P (annual)
2. Undertake method in Section 2.2.1 and 2.4.3 which generates the low frequency component of P in such a way as to match the desired (perturbed) Hurst coefficient. Thus, the low frequency behaviour (stressor 3) is the first perturbation in the order of operations.
3. Generate annual temperature using the output of (2) and historic T-P correlation (Section 2.2.2)
4. Perturb precipitation temporal average (stressor 1)
5. Perturb temperature temporal average (stressor 2)
6. Disaggregate to monthly (Section 2.2.3)
7. Perturb seasonality (stressor 4)
8. Generate monthly PET from monthly T (Section 2.2.4)
9. Rainfall-runoff modelling (Section 2.3)
10. Perturb rainfall-runoff relationship (stressor 5)

### 2.4.7 Guidance for setting perturbation axis limits

As noted, the stressors can each be considered as individual "axes" in a stress testing "space". The axis limits are thus the extrema in each direction by which the stressor in question is perturbed. Setting perturbation axis limits is a subjective choice that may vary by case study. Axis limits are often set to envelop the range of outcomes presented by GCM models (see, eg., Henley et al., 2019), with additional margin on either side reflecting that future conditions may be beyond the range presented by GCMs (Brown and Wilby, 2012). It is expected that this will be an appropriate method for Axis 1 (long term mean P), Axis 2 (long term mean T) and Axis 4 (seasonality of precipitation) in most cases, and an assessment of GCM simulations for the region in question is recommended for these axes.

However, Axis 3 (low frequency behaviour of precipitation) is not an area of strength for GCMs, although work is ongoing and newer generations of models are improving (see references in Section 2.4.3 above and Bayr et al. 2018). For long term oscillations such as the Pacific Decadal Oscillation, quality of simulations varies by continent and the simulated oscillations often have the incorrect period (Wei et al., 2017). For the shorter ENSO, models are often biased towards the La Niña state, while "too weak atmospheric feedbacks can cause quite different ENSO dynamics to observed" (Bayr et al., 2018, p3171). Given these critiques, it seems reasonable to refer to independent information when setting the limits of Axis 3. Given Axis 3 is defined as a change in Hurst coefficient, a useful starting point might be published values for observed Hurst coefficients, for the region in question but also at wider spatial scales (continental or global) to characterise the limits of observed Hurst values.

GCMs provide no information about changes in rainfall-runoff relationships, so Axis 5 limits must be decided based on hydrological factors. If the study region has seen a historic change

in rainfall-runoff relationship, this may provide a useful reference point (see case study, Section 3). In cases where no shift has occurred historically, it is difficult and subjective to decide the range of future values to test. One possible route would be to adopt identical limits to the case study herein (Section 3), but since Axis 5 limits are contingent on the adopted Box-Cox $\lambda$ value, adopting the case study limits assumes the $\lambda$ is a suitable choice. This may be the case if the skew of the transformed streamflow is not significantly different from zero at any sub-area (see testing methods presented in Supplementary Material S7).
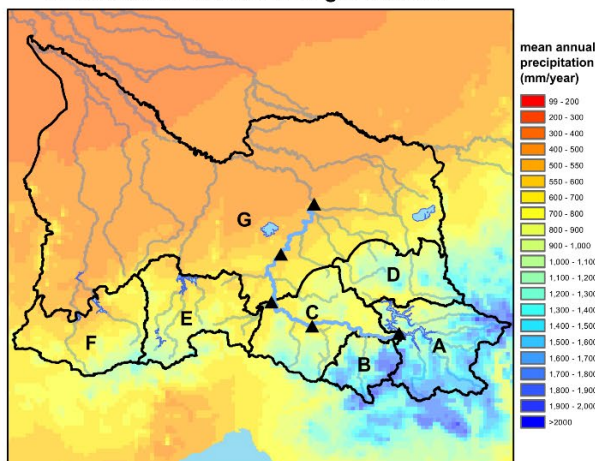
# 3 Case study application

## 3.1 The Goulburn-Broken-Campaspe-Loddon (GBCL) system

We demonstrate the framework on south-east Australia's 40,000 km² Goulburn-Broken-Campaspe-Loddon (GBCL) system (Figure 6a). This region is made of up the traditional lands of the Yorta Yorta, Taunarung and Dja Dja Wurrung Clans. Part of the Murray Darling Basin,
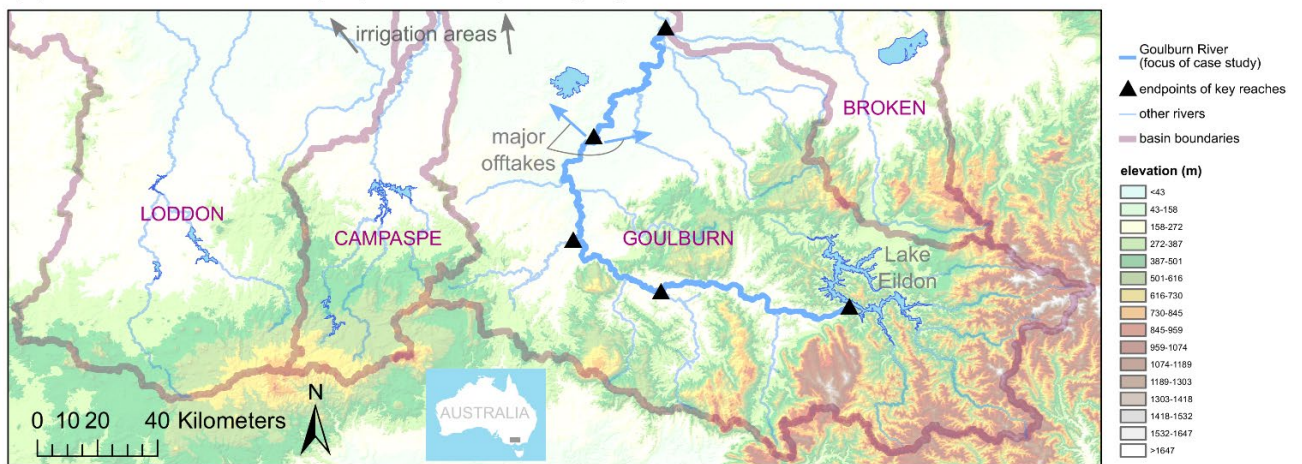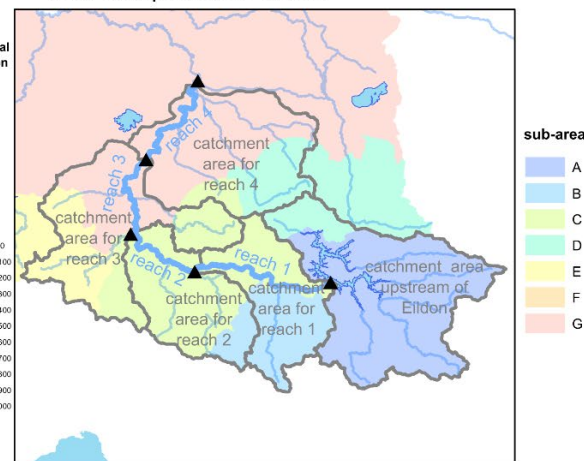


*Figure 6: Maps of study area. (a) Topography, basin boundaries, and the key river in focus in this article (Goulburn). (b) Division of study area into sub-areas chosen considering catchment areas of key reservoirs, hydroclimate, and known hydrologic response. (c) overlay of regions from (b) with tributary catchments areas for the Goulburn River, showing how each reach accepts inflows from multiple stochastic sub-areas.*

Australia's most important agricultural region, GBCL rivers supply approximately 12% of the basin's water from approximately 2% of the area. The system comprises a main regulated waterway (the Goulburn River) impounded by a large carry-over reservoir (Lake Eildon) and supplemented by smaller impounded waterways (Campaspe, Loddon and Broken Rivers) connected together by transfer channels (not shown in Figure 6). Due to the Murray Darling Basin's water trading scheme, water impounded within the Goulburn system is used over a wide area and is often traded as far as 700 km away. The main water use is agricultural irrigation, but the system also provides municipal supply for numerous towns. GBCL rivers support significant aquatic habitats including for turtles, platypus, fish such as the critically endangered Murray Cod, and ecologically significant riparian forests (Koster et al., 2012; Horne et al, 2020). Federal and state governments have invested significantly in water entitlements purchased for environmental use (Skinner and Langford, 2013), and this water is actively managed and released for environmental benefit (Doolan et al 2017; Stewardson and Guarino, 2018; Gawne et al., 2020), which may include objectives to inundate key wetlands or provide flow events that trigger fish spawning at specific times of year. Often this is achieved by supplementing ecologically important flow events originating in unregulated tributaries (Watts et al., 2011; Docker and Johnson, 2017).

The stress test arises from a study of the vulnerability of environmental water management to climate variability and change. The full study encompasses all four rivers (Goulburn, Broken, Campaspe and Loddon) so this is the extent of the climate data generation (Figure 6b). For the hydrological components of this demonstration, we focus on the Goulburn River only (Figure 6a) as an example of a regulated carry-over system. Lake Eildon's capacity, $3 \times 10^6$ m$^3$ is more than double the historic median annual inflow. Downstream of Lake Eildon, the regulated portion of the Goulburn River is divided into four reaches (see triangles in Figure 6) by environmental flow managers, which require consideration in the setup of the stress test. Note, the reach division is an arbitrary feature of our case study; users are encouraged to sub-divide their study area in whichever way best suits their study area and serves their application (see also Section 3.4.1).

As described in John et al. (2021c), the streamflow outputs of the framework are passed to a monthly timestep river systems model which represents the operation of the dams and transfer channels. The vulnerability of the system to climate change and variability is assessed via two criteria: supply reliability (an output of the river systems model); and ecological health (an output of associated ecological models forced with outputs of the river systems model). Here, we briefly summarise the results of this vulnerability assessment, with further detail provided in John et al. (2021c).

## 3.2   Data

Gridded historic precipitation and temperature data are from the *Australian Water Availability Project* (AWAP; Jones et al. 2009). Note, the adopted formulation for temperature is $T_{max}$, the maximum daily temperature. Gridded potential evapotranspiration data are from the *Scientific Information for Land Owners* project (SILO, Jeffrey et al., 2001). Of the PET options provided by SILO we adopt Morton's Wet Environment Evaporation Over Land (Morton, 1983) because the main purpose of PET for this study is to force the rainfall-runoff model, and this formulation has been used in prior rainfall-runoff modelling studies in this region (eg. Fowler et al., 2020a). Streamflow data are from the Australian Bureau of Meteorology's *Hydrologic Reference Stations* project (Turner et al., 2012) as compiled in the *CAMELS-AUS* dataset (Fowler et al., 2020b). Information on storage capacity is from https://www.g-mwater.com.au/ (last accessed 22nd October 2021).

## 3.3 Stress test set up: selection of stressors

Earlier (Section 2.4), a rationale for inclusion of each of the five stressors was given, in general terms. The specific reasons cited for each stressor are all true for the GBCL case study. For example, stressor 3 (low frequency behaviour) is important to system performance because local climate variability is relatively high on a global scale, leading to spells of relatively higher cumulative deficit below the median (eg. Peel et al., 2004). While the large size of Lake Eildon buffers the system against short term (up to five years) variation, longer-scale oscillations may go beyond buffering capacity and so reveal system vulnerabilities.

Regarding stressor 4 (precipitation seasonality), many regulated systems may be insensitive to this stressor because the storage mitigates against seasonal lows. In the GBCL system, seasonality is important regardless because unregulated tributaries generate ecologically important flow events at specific times of year, as mentioned above. In addition, seasonality of flows affects announcement of seasonal volumes of water allocation, which can in turn influence system operation and water trade.

Lastly, regarding stressor 5 (changes in rainfall-runoff relationship), the 13-year "Millennium" Drought (1997-2009) saw changes in rainfall-runoff relationship for many GBCL waterways including the Campaspe River, Loddon River and unregulated tributaries of the Goulburn River (eg. Saft et al., 2015). The streamflow in some waterways was approximately 70% below the expected flow given the observed rainfall. Some catchments have recovered, returning to their earlier relationship between P and Q, whereas others seemingly remain in a low flow state (Peterson et al., 2021). These recent flow reductions are an important issue for project partners and the method needs to allow for the possibility that further future shifts in rainfall-runoff relationship may occur.

In line with the above, all five of the stressors in the framework are adopted in the case study. Table 1 outlines how the axis limits are set, with various factors taken into consideration.

*Table 1: Adopted axis limits for the stress testing space*

| Stressor / axis | Axis units, limits & gradations | Considerations |
|---|---|---|
| 1. P, long term mean | Unit: % change, mean annual P<br>Highest risk considered: -40 %<br>Lowest risk considered: +15 %<br>Number and interval of gradations: 12 @ 5% | Axis limits are set with reference to GCMs, using CMIP5 projections collated by the 'Climate Futures' tool (Whetton et al., 2012, see (www.climatechangeinaustralia.gov.au/en/projections-tools, last accessed 22nd October 2021). Project partners are interested in planning timescales out to 20 years so the year 2040 is in focus. |
| 2. T, long term mean | Unit: absolute change in °C<br>Highest risk considered: +4°C<br>Lowest risk considered: +0°C<br>Number and interval of gradations: 9 @ 0.5°C | For temperature, the collated projections indicate local warming of less than 3°C by 2040; we envelop this range by extending to +4°C. For precipitation, most models lie in the range -15% to +5%. We envelop this range by extending to [-40%, +15%]. The -40% limit is considerably beyond the GCM projections which reflects the desire of the project partners to explore lower rainfall futures well beyond the ~15% reductions experienced during the recent 13-year drought. |
| 3. P, low frequency behaviour | Unit: change in Hurst exponent<br>Highest risk considered: +0.075 in Hurst exponent<br>Lowest risk considered: -0.030 in Hurst exponent | Historically, the seven sub-areas adopted for this study exhibit Hurst values between 0.66 (Sub-area G) and 0.72 (Sub-area B), a range of 0.06. Higher values have been reported for Australian rainfall (eg. 0.76 in Whiting et al., 2003). Iliopoulou et al., (2018) explore a global dataset and report that very few locations have values > 0.8. Thus, +0.075 is adopted as a top value for the axis |

| | Number and interval of gradations: 8 @ 0.015 | since 0.72 + 0.075 is approximately 0.8. Decreases in Hurst coefficient out to -0.03 are also considered. |
|---|---|---|
| 4. P, seasonality | Unit: % of yearly P transferred from cold to warm months<br>Highest risk considered: +15 %<br>Lowest risk considered: -6 %<br>Number and interval of gradations: 8 @ 3 % | Like Axes 1 and 2, GCM projections collated in the 'Climate Futures' tool are the key reference. The tool provides % changes using the same seasons as adopted here (warm NDJFMA, cool MJJASO), and when converted to the measure adopted here (see 'unit') the 2040 projections translate to approximately -4% to +15%. We envelop this range by adopting limits of -6% & +15%. |
| 5. rainfall-runoff relationship | Unit: vertical shift in Q (mm/yr) but in Box-Cox transformed space ($\lambda=0.79$)<br>Highest risk considered: -50 units<br>Lowest risk considered: +12.5 units<br>Number and interval of gradations: 9 @ 6.25 units. | Parts of the study region have seen historic shifts in rainfall-runoff relationship, and the intention is to set the axis limits relative to the worst historic change. The adopted Box-Cox parameter is 0.79, derived in Supplementary Material S7. In Box-Cox transformed space, the largest observed decline during the Millennium Drought was 25.2 units, observed in Sub-area E (Section 3.7). The axis limits were therefore set from 12.5 (ie. a moderate shift towards more flow for a given rainfall) to -50 (ie. a decline twice as severe as the most severe historic shift). |

## 3.4   Stress test set-up: spatial and rainfall-runoff modelling considerations

### 3.4.1   Division into sub-areas

As noted in Section 2.1, there are at least three reasons why sub-division of the study area is useful in a given application. Below we discuss how each reason relates to this case study. Users should note that the following comments are case specific and do not imply a general rule; other applications may be more straightforward.

- *Reason 1: create demarcations that are necessary for a given application (eg. upstream versus downstream of a reservoir).* Since the unregulated tributaries of the Goulburn River are instrumental in providing environmental flows, it is not acceptable (within this application) to lump them together with the area upstream of Lake Eildon. Similar considerations apply to the other reservoirs (although they are not in focus here). Thus, the downstream boundaries of some sub-areas have been drawn at the major reservoirs on the Goulburn, Campaspe and Loddon Rivers.
- *Reason 2: divide a heterogeneous study area into more homogeneous sub-areas.* An example of applying this principle based on climate data is the distinction between sub-areas B and C. This distinction makes spatially lumped averages meaningful because they are more representative of underlying conditions; further, it improves confidence in rainfall-runoff modelling by avoiding unreasonable requirements of models (ie. the requirement to represent many contrasting areas simultaneously). We also apply this principle based on similarity of hydrological response; for example, Figure 6c shows how a small portion (marked in light green) of Lake Eildon's catchment area is grouped with areas outside – this is because it is known that this region (the Strathbogie Ranges) behaves differently hydrologically to the rest of the catchment (similar considerations apply to the boundary between sub-areas C and E). Because boundaries based on similarity do not follow reach inflow catchment area boundaries (Figure 5c), it is necessary to apply a post-

hoc step where flows from sub-areas are added together in appropriate weighted averages to provide total reach inflows (see Supplementary Material S4). While the last step may be inconvenient, the workflow overall merely reflects the varied hydroclimate of the study area (Figure 6b). Lastly, the large flat and dry area in the north of the study area, called the Riverine Plain (Sub-area G), requires special consideration. Despite its large size, this area contributes very little streamflow and is lumped together as a single unit, reflecting its relative homogeneity.

- *Reason 3: capture distinct but correlated behaviour, eg. in different tributaries.* This principle has already been touched on above when it was noted that the unregulated and regulated tribuaries need to be distinct rather than lumped, for this case study. The principle is also important among the unregulated tributaries: for example, inflows to Reaches 1 and 4 are largely determined by sub-areas B and D respectively; keeping these sub-areas distinct allows for a more nuanced representation, capturing how these tributary inflows combine to provide flows for the ecologically important downstream reaches.

### 3.4.2   Rainfall-runoff modelling and representative catchments

The inbuilt WAPABA rainfall-runoff model is adopted for streamflow generation (Wang et al., 2011). A potential difficulty is that none of the sub-areas are gauged in their entirety, although all of them have multiple gauged subcatchments in the CAMELS-AUS dataset (Fowler et al., 2020b). The solution is to select one of these subcatchments – where possible one whose hydroclimate is close to average for the sub-area – and to assume it is "representative" of the whole sub-area, scaling simulated flows as required to represent the whole sub-area. Since the sub-area boundaries are intentionally set to enclose similar regions, this method provides an adequate solution. Normally an approach like this might suffer from issues with flow routing, but these issues are moot due to the monthly (not daily) timestep. Selected representative catchments for each sub-area are listed in Supplementary Material S4. The models are calibrated using a method (Fowler et al., 2016) which equally weights the performance in wet and dry periods – see Supplementary Material S5 for more information. This method uses the Kling Gupta Efficiency (KGE; Gupta et al., 2009) as the measure of model performance. The KGE considers the model's ability to match three relevant aspects for water resource and ecological modelling, namely: (i) the long term average flow; (ii) the variability in flows; and (iii) flow timing (as measured by temporal correlation). Although the KGE's focus on high flows has led some studies to adopt it blended together with a low flow metric (eg. Knoben et al., 2020; Trotter et al., 2021), these studies used daily models. Given monthly streamflows are less skewed than daily, the tendency for undue focus on high flows is diminished.

### 3.5   Results, 1: separation into high and low frequency components

For brevity, we show only highlights of the case study application, but the user can interrogate the results set by downloading the framework and running the example. The first such highlight is the separation of annual P into high and low frequency components (Figure 7). Rather than showing each Intrinsic Mode Function (IMF, see Figure 3), only the combined high and combined low frequency sequences are shown. In this case the distinction between high and low frequency is made between IMFs 2 and 3. Broadly, this gives a separation into inter-decadal and intra-decadal variation, which is considered appropriate since the capacity of Lake Eildon ($3{\times}10^6$ ML) is sufficient to buffer the system against short (<5 year) variations but is less robust to interdecadal fluctuations. The resulting low frequency timeseries has 11 zero crossings in 107 years (as noted), giving an average period of approximately 20 years.

Despite the varied hydroclimate between sub-areas, the low frequency sequences tend to enter multi-year dry periods and wet periods at the same time. This means that the area-weighted
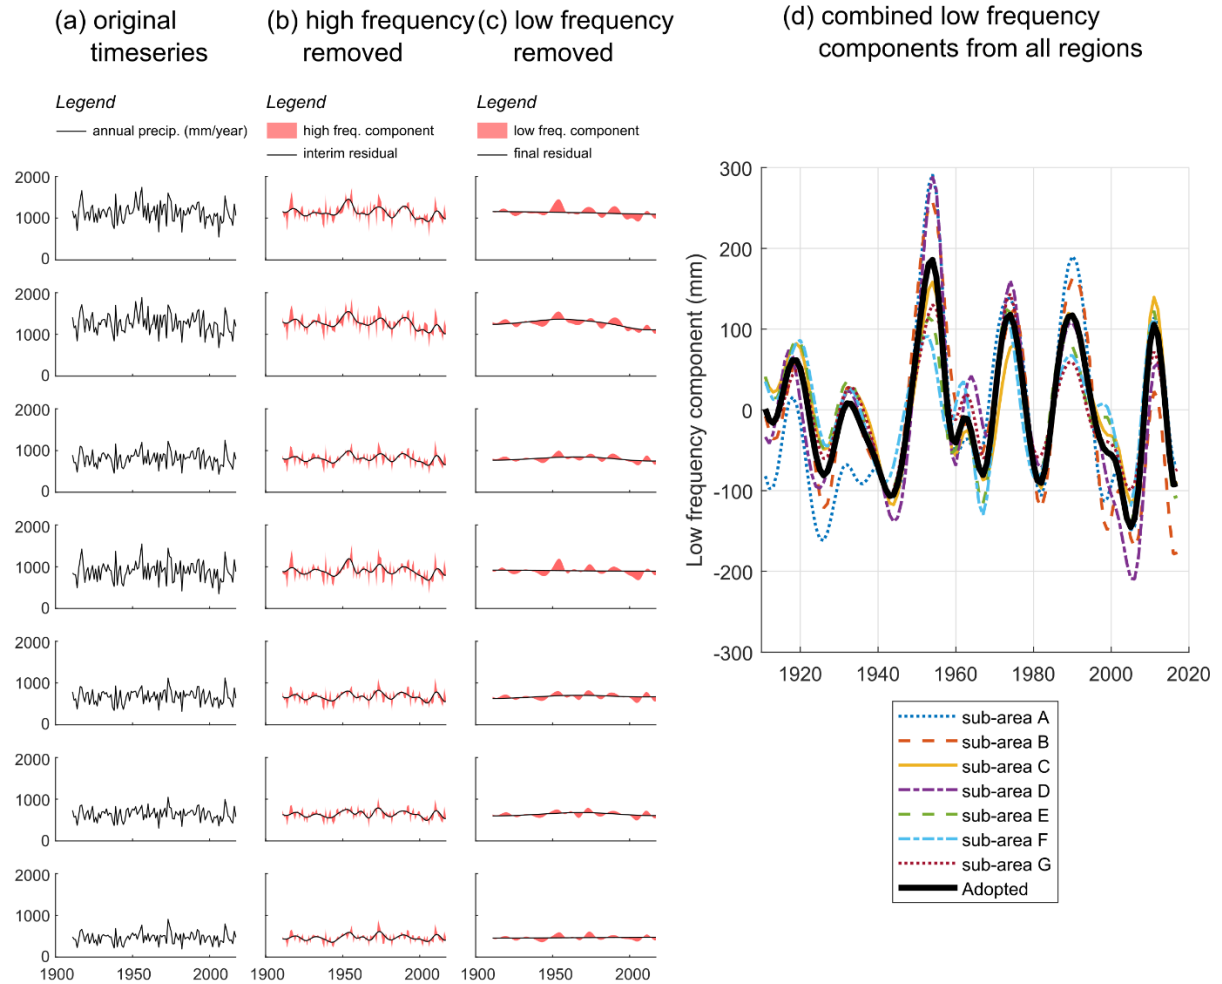
*Figure 7: Output of CEEMDAN applied to the 7 sub-areas of the case study, (a-c) for each sub-area individually, showing the low and high frequency components that result when the first two IMFs are combined to give high frequency and the rest for low frequency; (d) the low frequency components for each sub-area, shown with the adopted timeseries). Note, although the 'adopted' line amplitude underestimates reality in some sub-areas and overestimates in others, the way it is subsequently used (see Sections 2.2.1.4 and 2.4.3) inherently corrects this so that the historic Hurst exponent is matched for each sub-area. Note, code to produce a variant of this figure is provided with the framework.*

average (Figure 7d, note sub-area G is excluded from the weighting because it is so dry as to contribute little streamflow yet so large as to dominate the weightings) is broadly representative of each of the underlying sequences. Recall that the mismatches in amplitude between the adopted combined timeseries and a given sub-area timeseries are later resolved by the Hurst coefficient matching procedure (Section 2.2.1.4 and 2.4.3). The mismatches in timing with which the different sub-areas undergo multi-year dry and wet periods does cause some issues with generated data, as discussed in the Limitations section (Section 4.1).

## 3.6 Results, 2: stochastic generation for baseline (no perturbation)

We now proceed to evaluate the quality of baseline stochastic outputs (no perturbation), based on a replicate length of 3000 years. The evaluation is done in two parts: an evaluation of seasonality (Figure 8); and an evaluation of the frequency distribution of annual flows and multi-annual flow sums (Figure 9). Note, to avoid conflating the comparison of seasonality with the assessment of bias, Figure 8 only shows percentages in each month rather than absolute values.

Figure 8 reveals an excellent match in seasonality across P, $T_{max}$ and PET (a-c) and a good match for Q (d). Note that these comparisons are for Sub-area B as an example, but the results are broadly consistent with other sub-areas. The match in seasonality in P and $T_{max}$ is unsurprising given the method focusses on directly adopting historic within-year patterns. While the seasonal pattern for PET is also well matched, the variability within each month is underestimated, which is consistent with the decision not to explicitly represent the variability around the T-PET relationship for each calendar month (Section 2.2.4). Plot (d) relates to the representative catchment for Sub-area B; this plot thus reflects the ability of the WAPABA model to represent seasonal dynamics in this example catchment. The comparison reveals underestimation during the driest part of the year and over-variable November and December flows, but these issues are relatively minor, and overall the model seasonal performance is good. Further information on WAPABA calibration performance across all sub-areas, including a comparison during the historic period, is given in Supplementary Material S6.

Changing focus now to assess the match in annual streamflow volumes, Figure 9a indicates the stochastic replicates show very little bias and a similar frequency distribution to historic. Whereas Figure 8 focused on performance in a single example sub-area, Figure 9a assesses performance over a wide area because it compares aggregated flows that include inflows to the four key reaches of the Goulburn River downstream of Lake Eildon. Thus, the pleasing match reflects performance across the five sub-areas that provide reach tributary inflows, which represents significant climatic variability (see Figure 6b).

The next consideration is whether multi-year dry periods and wet periods are well matched. Figure 9b is based on the same information as Figure 9a, but the streamflow sequences are accumulated into 5-year rolling sums prior to the analysis. Overall, the stochastic data matches
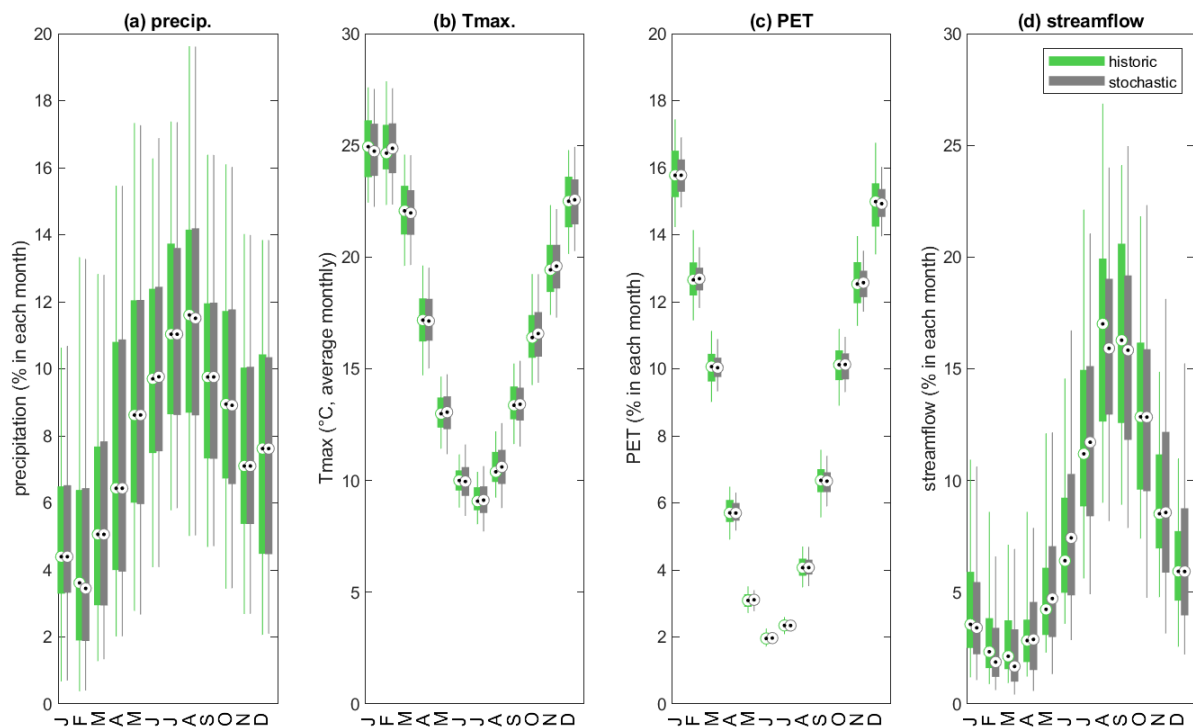


*Figure 8: Performance assessment of the seasonality of stochastic replicates, for the baseline (no perturbation) for an example sub-area (sub-area B). For a, c and d the boxplots show the variability in the proportion of the annual value reported in that month. Plot d relates to the representative catchment of sub-area B, station 405209 (Acheron River at Taggerty). Box plots show 5th, 25th, 50th (filled circle), 75th and 95th percentiles. Note, code to produce this figure is provided with the framework.*
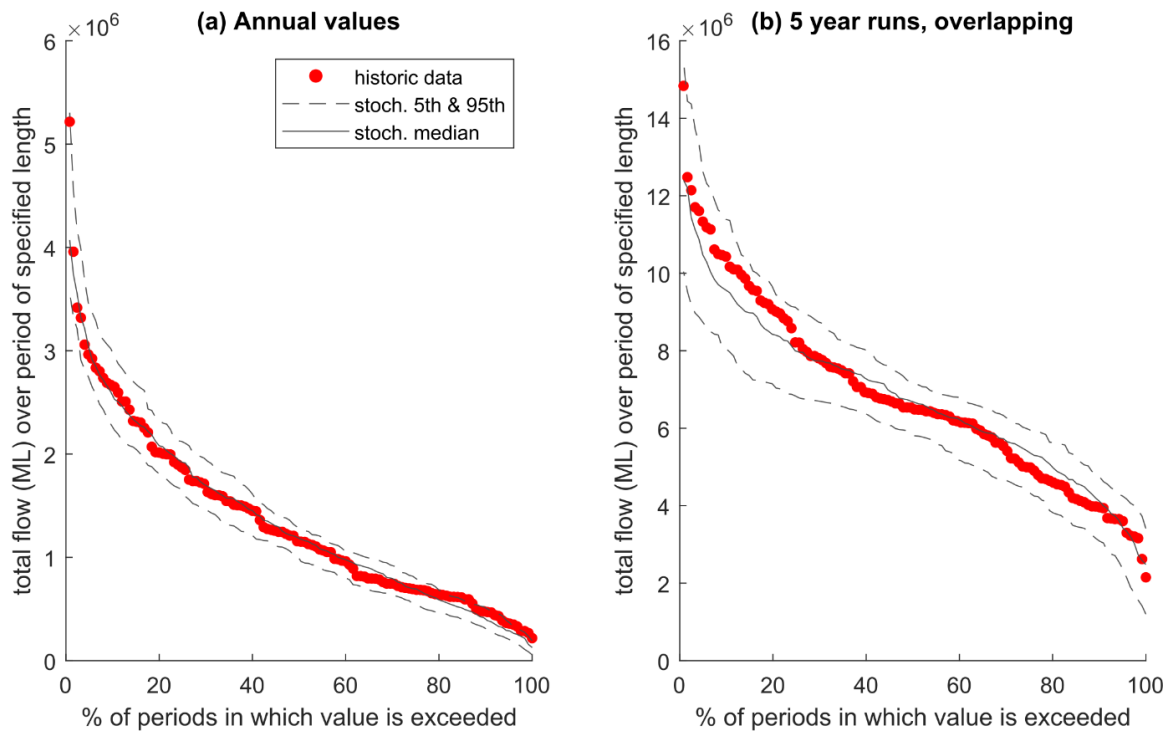
*Figure 9: Performance assessment of frequency distribution of annual flows and multi-annual flow sums for the stochastic baseline (no perturbation). Both figures relate to the sum of inflows to the four reaches of the Goulburn River downstream of Lake Eildon (see Figure 6). (a) frequency curve of annual flows; (b) frequency curve of 5-year overlapping runs. The stochastic median and $5^{th}/95^{th}$ percentiles are constructed by splitting the 10,000 year stochastic sequence into segments of the same length as the historic series (107 years), then analyzing each segment in an identical fashion to how the historic timeseries is processed, to ensure a fair comparison. This leads to a set of many lines, and for each ordinate on the x axis, the median and $5^{th}/95^{th}$ percentile y values among the lines are calculated. Note, code to produce this figure is provided with the framework.*

the historic 5-year rolling sums well. The lower tail of the distribution is particularly well matched, with the median line among the replicates following the historic relatively closely, and the lower limit indicating that the stochastic data has five-year droughts more severe than those in the historic record, as we would expect for a 3,000 year sequence. The match at the upper end is less close, with a persistent 5-10% underestimation of five-year wet sequences in the $10^{th}$ to $30^{th}$ percentile range. This may also reflect the limits of the stochastic workflow (ie. CEEMDAN-based split followed by Broken Line Process) to replicate the way the years are organised into sequences. However, considering the many aspects of the stochastic method and the varied hydroclimatic conditions represented, the overall result in Figure 9b is very favourable.

## 3.7   Results, 3: changes in rainfall-runoff relationship

The fifth perturbation stressor, "change in rainfall-runoff relationship" is an innovative part of this study, so we briefly focus on this aspect of the case study. Parts of the case study region have experienced historic shifts in rainfall-runoff relationship, and Figure 10a shows two such catchments (specifically, the representative catchments for Sub-areas C and E, the latter being the one with the greatest shift among all the sub-areas) and one catchment that has remained stable (Sub-area A). Figure 10b shows the practical workflow to shift an existing synthetic monthly sequence down by 15 units in the Box-Cox transformed space (recall that the axis varies from +15 units down to -50 units; Table 1). It is hoped that this visual demonstration of the steps helps to clarify the process.
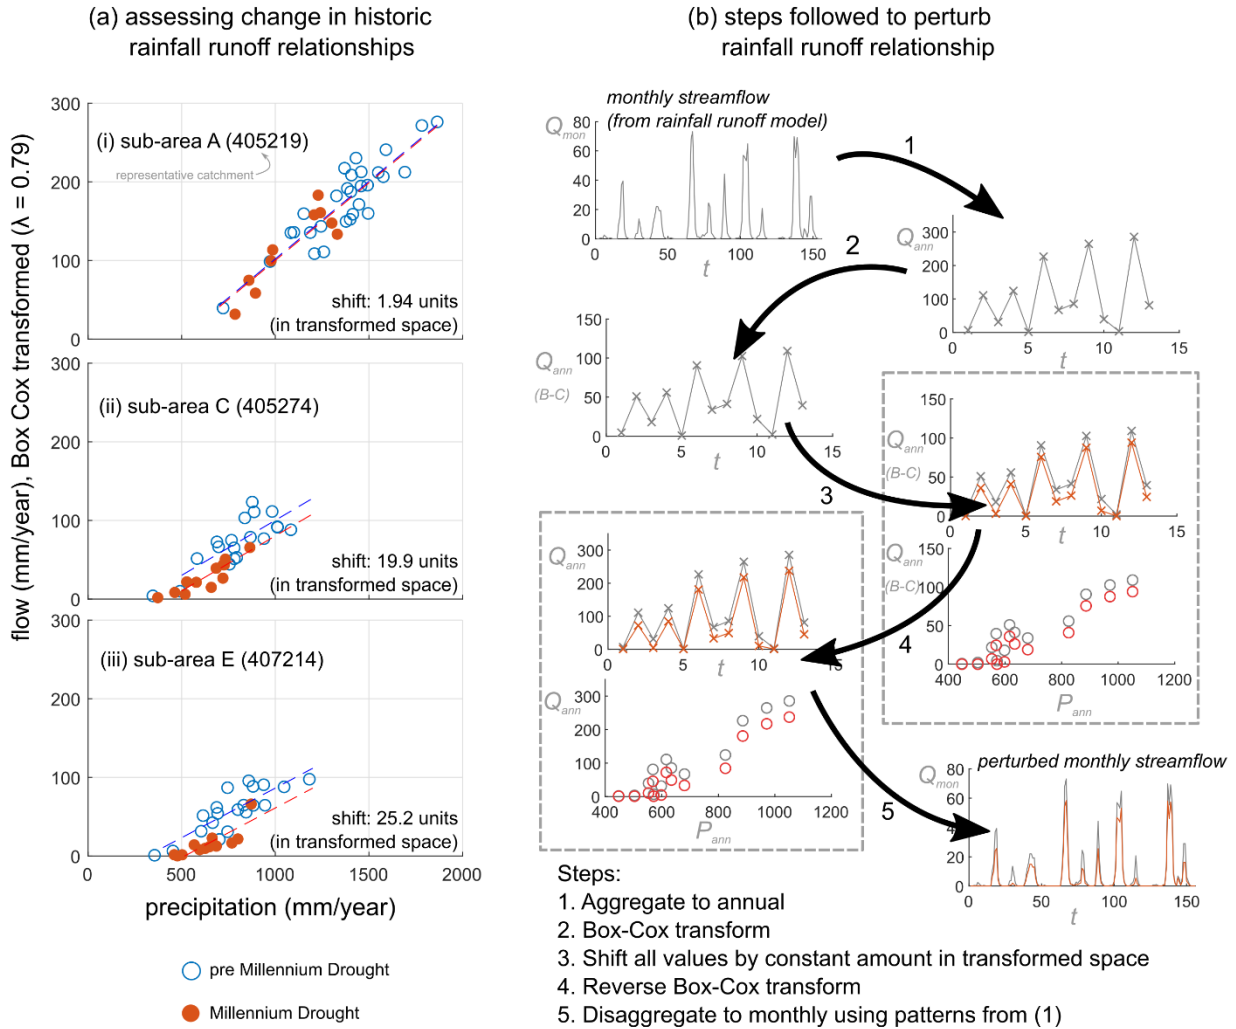
Figure 10: (a) assessment of changed rainfall-runoff relationships for three of the sub-areas in the case study, including one example of no change (sub-area A) and two examples of change (C and E). (b) steps required to perturb the rainfall-runoff relationship for an example timeseries of 13 years, for an example shift of 15 units in transformed space and the same Box Cox parameter as in (a). Note: in (b), P-Q plots do not constitute a separate step; they are shown only to emphasise the link with (a).

## 3.8 Results, 4: perturbation of stochastic replicates

This subsection examines changes in stochastic data after perturbation. Figure 11a to 11e repeat the seasonality and streamflow volume frequency curves from Figures 8 and 9, for each perturbation stressor. In each case a 'higher risk' (red) and 'lower risk' (blue) perturbation is performed, and the result is compared with the baseline, except for temperature which is not expected to decrease so a 'lower risk' scenario is not considered. Figure 10f shows a combination, with each stressor exhibiting a small perturbation towards higher risk.

Perhaps unsurprisingly, the highest impact scenario (in terms of streamflow volumes) is the Stressor 1 30% reduction in mean annual precipitation, which reduces streamflow by 80-90%. The next highest is the Stressor 5 reduction in rainfall-runoff relationship by 25 units, which reduced streamflow volumes by 30% (and by as much as 50% in drier years). For Stressor 4 (seasonality of precipitation), the shift from cold month precipitation to warm month precipitation is reflected in both the seasonal streamflow plots but also in a decrease in average streamflow, because more precipitation is falling concurrently with high evaporative demand

and thus a larger proportion goes to evapotranspiration. The Stressor 3 results (low frequency behaviour) indicates minimal impact on long term average streamflow (as expected), but the steeper gradient (Figure 1c-iii) indicates a more variable climate with larger multi-year oscillations around the mean.

The final test showing a combination of small perturbations in all axes (Figure 11f) is a thought provoking result, because none of the axes are very far perturbed yet the streamflow volumes reduce by approximately half. Thus, interactions between the axes are important and may dominate how climate change impacts on water resources, a theme that is explored elsewhere (John et al., 2021b). It suggests the value of frameworks, such as this one, that can explore multiple aspects of climatic change simultaneously and in a systematic way.

## 3.9   Results, 5: Vulnerability assessment for case study

For completeness, we also summarise the results of the vulnerability assessment for this case study, but is noted that the system and ecological modelling required is out of scope of the framework's publicly released code. Figure 12 shows the system response to climatic changes in terms of changes in allocation reliability (for high reliability shares; Figure 12a) and ecological health (for large bodied fish; Figure 12b). Note that John et al. (2021c) also assessed a range of other aspects (eg. low reliability water shares; ecological endpoints such as river-dependent vegetation and small bodied fish).

The colour gradients in Figure 12 indicate greatest sensitivity to two of the five axes: Axis 1 ($P_{long\ term\ average}$) and Axis 5 (rainfall-runoff relationship). Axis 2 (T) also has a clear influence on outcomes, but is secondary to the first two. The insensitivity to Axis 4 (seasonality) is because regulated environmental releases can offset changes in seasonality of inflows – thus, when repeated for unregulated locations, the vulnerability assessment shows an increased sensitivity to seasonality (not shown). The insensitivity to Axis 3 ($P_{low\ frequency}$) arises partly because Figure 12 reports on the median outcome, whereas this axis causes a widening in the distribution of outcomes without changing the central tendency (eg. more extreme multi-year droughts/floods with unchanged median) as mentioned above.

Figure 12 also shows interactions between any two of the axes. For example, a reduction in precipitation of 15% causes the years fully allocated to reduce from 75% to 40%. However, if this occurs simultaneously with a temperature increase of 3°C, the outcome lowers further to 20%. This is consistent with the comments above about the significance of interactions. While the system is vulnerable to changes in average precipitation, most GCMs project modest changes in this stressor within the planning window of 2020-2040. In contrast, shifts in rainfall-runoff relationship have occurred rapidly in the past (eg. within 1-3 years; Peterson et al., 2021), so the potential for this stressor to impact the system during planning timeframes is relatively acute, underscoring the importance of hydrological research on this topic.

# 4   Discussion

## 4.1   Limitations

The framework is complex and its output multifaceted, so we supplement the above figures assessing the outputs (ie. Figures 7-11) with fourteen further figures in the Supplementary Material, Section S8. While most of these figures confirm the ability of the framework to replicate the features of historical data, there is one category with room for improvement: namely, spatial correlation. Figures S8.2 to S8.7 demonstrate that the historic data show a lower spatial correlation between sub-areas, associated with a greater propensity for a given
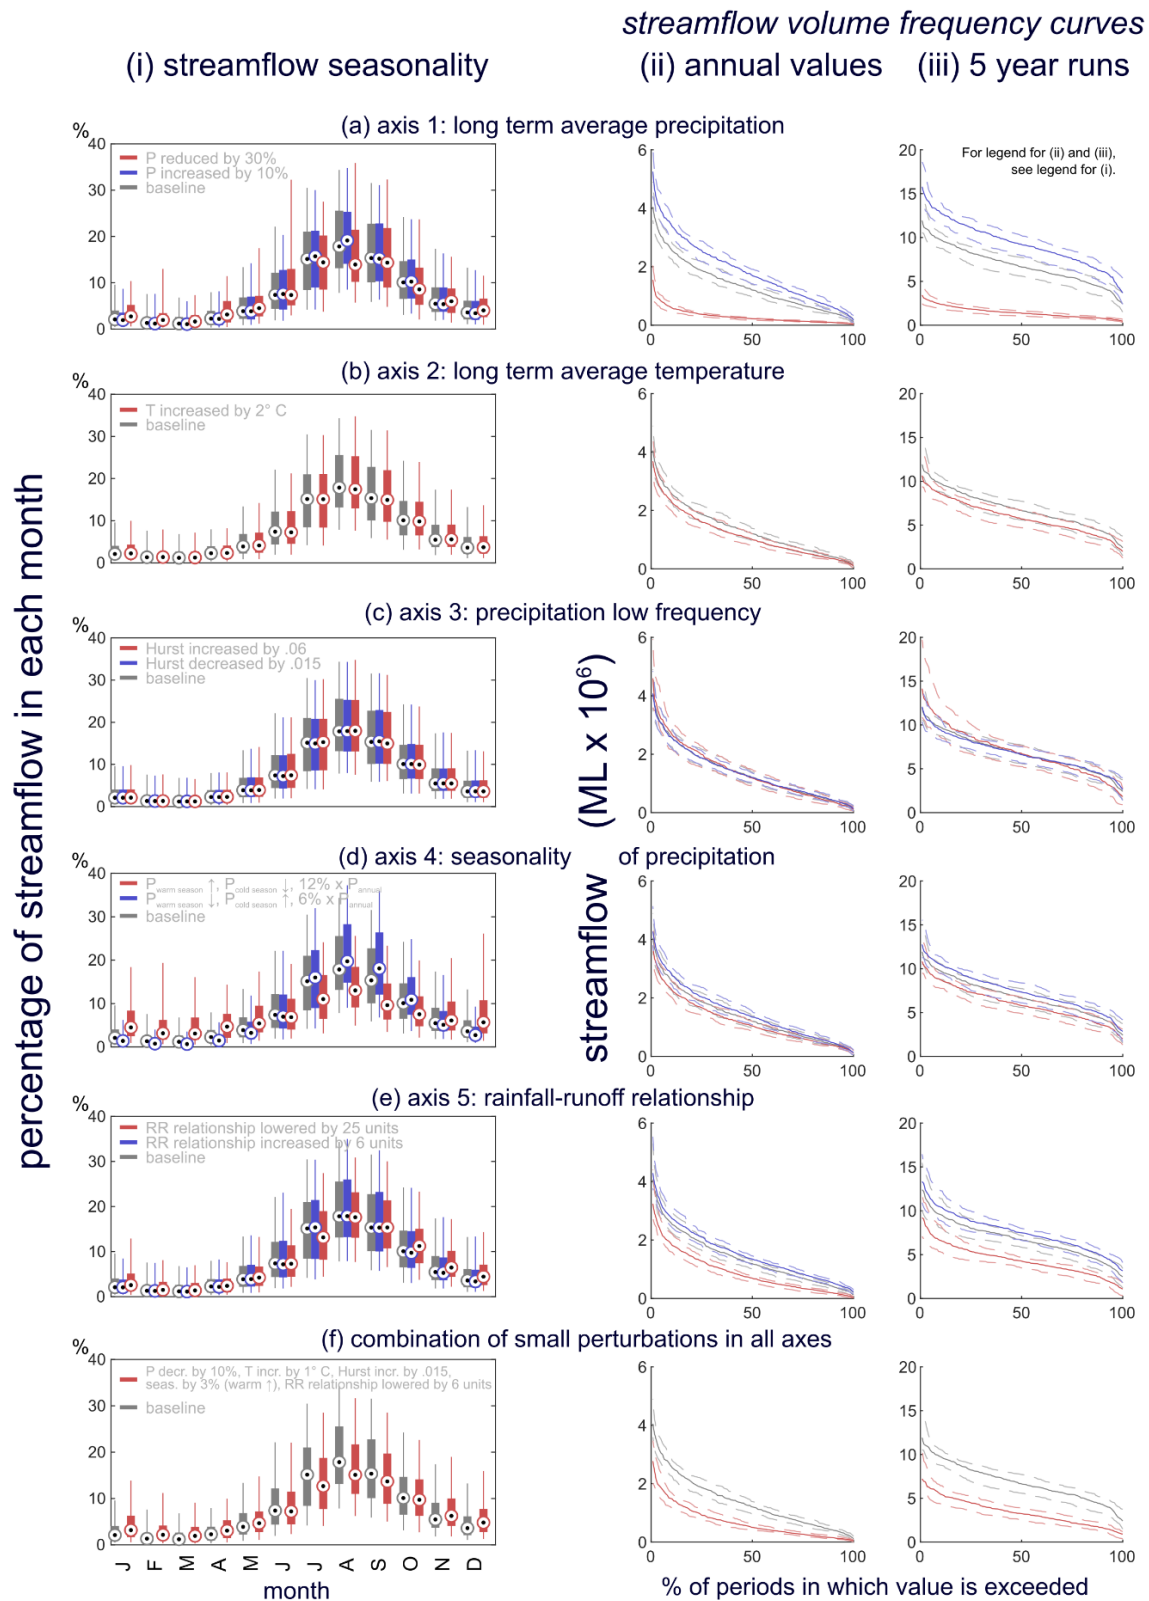
Figure 11: Evaluation of perturbed stochastic sequences for the to the sum of inflows to the four reaches of the Goulburn River downstream of Lake Eildon (see Figure 6). (a-e) each of the 5 stressor axes perturbed individually; and (f) an example where all axes are perturbed simultaneously. Column (i) follows the format of Figure 8d; columns ii and iii follow the format of Figure 9a and 9b, respectively. Box plots show 5th, 25th, 50th, 75th & 95th percentiles. Note, code to produce this figure is provided in the framework.
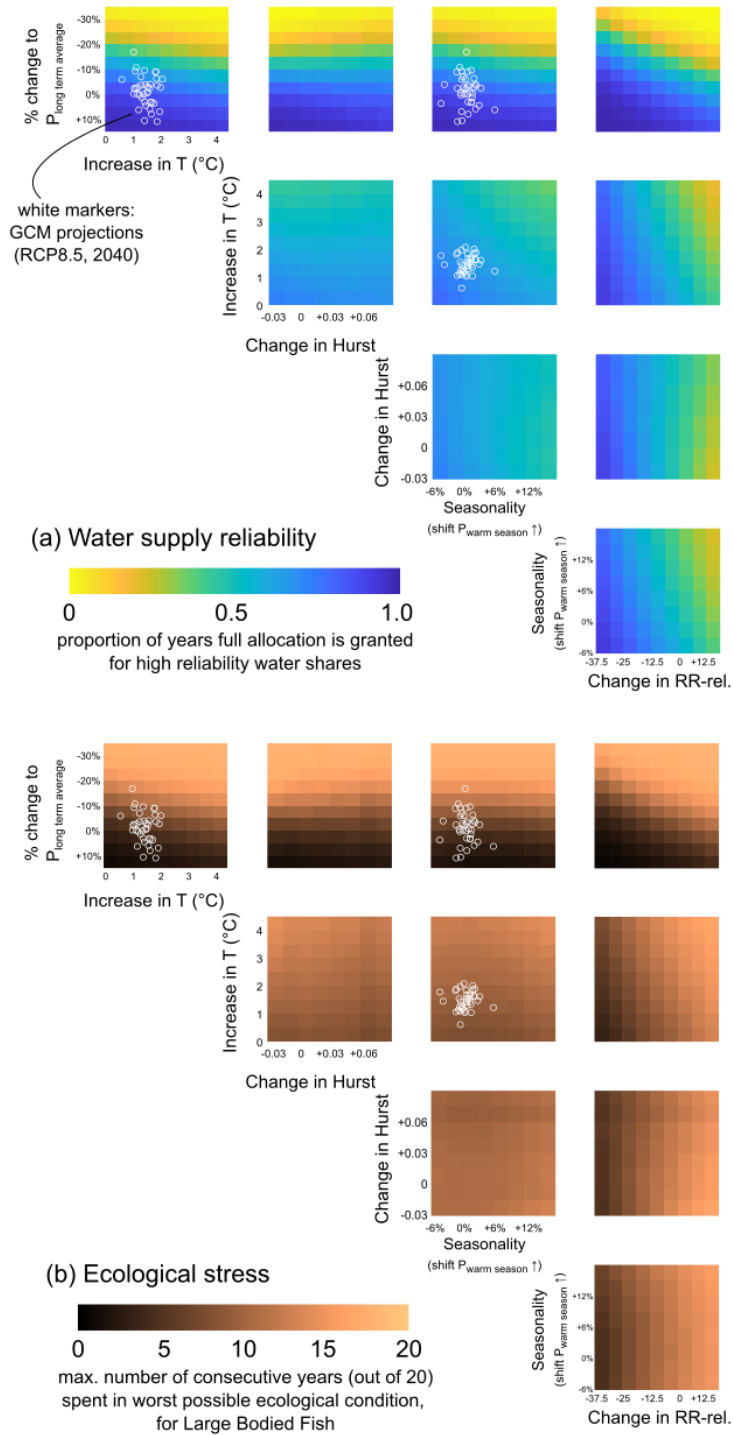
*Figure 12: Stress testing vulnerability assessment for the Goulburn River for two selected measures. Each large square shows the response surface for two stressors at a time. Within these squares, each grid cell shows the median of the distribution of outcomes arising when two variables are fixed (as specified by axis ordinates) and the other three are allowed to vary according to their limits and gradations (Table 1). As outlined by John et al. (2021c), the results consider 100 × 20 year sequences for each combination of stressors. GCM results are considered unreliable for Axis 3 (Hurst / $P_{low\ frequency}$) and are irrelevant for Axis 5 (rainfall-runoff relationship), so GCM results are not shown on plots involving these axes. Ecological results are for the Lower Goulburn River and are thus sensitive to both unregulated and regulated components of the system.*

drought or flood to be more severe in one sub-area compared to another. One possible contributor could be the annual-monthly disaggregation step, since the existing mechanism to allocate patterns from wet, medium and dry years may not be nuanced enough to capture the complexity of monthly dynamics. However, given the problem also exists at annual and multi-annual timescales, it may be associated with the assumption (Figure 7d) that a single combined low-frequency sequence is adequate across all sub-areas, when in fact there were regional differences in timing of the historic low-frequency sequence. This assumption synchronises the stochastic data across sub-areas to a greater degree than in reality. Given the pleasing results (Figure 9) at system-scale (ie. after aggregating the sub-areas), we suggest this limitation is not critical to this case study nor is it likely to be critical to similar cases studies. Nonetheless, future research may examine options to remedy this limitation, with a key challenge being defining dynamics across sub-areas. That is, we know each sub-area should diverge from the 'combined' sequence, but how should these temporal divergences be arrange spatially across different sub-areas? We leave such questions to future research.

A higher-level limitation may exist regarding the generality of the framework. It is unclear whether the order of operations formulated here (Section 2.4.6) is generalisable to *any* desired set of stressors. It is possible that an additional stressor may exist for which there is no sequential order of operations that allows all stressors to behave independently. Such cases may require a methodology that examines tradeoffs on-the-fly between different desired traits of stochastic data (eg. Guo et al., 2018), outputting the *closest possible* stochastic sequence for a given set of perturbations. This is very different to the adopted methodology. Nonetheless, this may be somewhat of a hypothetical limitation, since the authors have not yet come across a stressor type that would "break" the current methodology, suggesting it has significant potential.

## 4.2   Computational effort and appropriate model complexity

The purpose of this subsection is threefold: (i) to report on the computation effort required; (ii) to discuss issues of appropriate model complexity related to the framework and/or raised while working on the case study; and (iii) to briefly discuss possible future directions regarding sampling of the climate stress testing space.

In the introduction, it was stated that a key benefit of a monthly-timestep stochastic framework is improved run times, so here we report the computational time taken in the case study. Working in Matlab on a 2016-era laptop (HP Elitebook with Intel i7-7500U 2.7 GHz processor), the generation of 3,000 years of stochastic data takes approximately 1.1 seconds. To cover the stress testing space, this 1.1 second task is repeated for each possible combination of gradations mentioned in Table 1, so there are ($12 \times 9 \times 8 \times 8 \times 9$) 62,208 combinations (for a total of $186 \times 10^6$ simulated years). The total run time is approximately 19 hours. Thus, a laptop can complete all the stochastic generation for the case study in less than one day continuous computing time.

However, runtime considerations are not limited to stochastic data generation – the run time of subsequent modelling steps is usually also important. In our case study, a pre-existing detailed daily river systems model was available, but it took ten minutes to run the 120-year historic sequence, which would translate to approximately 30 years in run time for $186 \times 10^6$ simulated years. It is hard to imagine an extensive stress test when each model run is so time consuming. Reframing the model as a monthly model, and ensuring the model included only those details relevant for the study (as described by John et al., 2021c), reduced run time by between three and four orders of magnitude (to a similar total time as the data generation). These design choices, including but not limited to timestep, enabled a much more thorough exploration of

the stress testing space within a modest computational budget. In cases where a system model already exists, either from previous research or from system operators, it would seem straightforward (and tempting) to adopt such models rather than build a model for each new context. However, our experience suggests that such "inherited" models should be carefully assessed for appropriateness for each new research question. Time spent up-front to gain the "right" model for the question pays dividends later. The challenges here are not only technical: in our case, additional (and ongoing) effort promoting trust in the new model among project partners was (and is) key to acceptance among stakeholders.

Our experience designing the framework and case study suggests that issues of appropriate model complexity are important within stress testing studies and need to be considered up-front ie. prior to investment in more detailed experimental design. For example, the choice to implement a monthly stochastic framework went hand in hand with the decision to simplify the river systems model (including but not limited to the decision on timestep). This underscores the importance of an initial holistic planning phase in decision scaling and scenario neutral studies, during which model complexity is agreed. Although issues of appropriate model complexity have been discussed for decades within hydrology (eg. Grayson et al., 2002), they are commonly framed in terms of a trade-off between realism and identifiability/predictive accuracy (eg. Akaike, 1973). The notion that there is a tradeoff between complexity/realism and the assessment of uncertainty is relatively new within the climate sciences (eg. Helgeson et al., 2020). The quality of uncertainty assessment in the context of climate stress testing may depend on such details as how many stressors and gradations can be tested within a given computational budget.

A final issue concerns the appropriate exploration of highly dimensional stress testing spaces. Sampling may seem straightforward for the common situation of a 2- or 3-dimensional climate stress testing study, but a 5-stressor study (as here) entails orders of magnitudes more computational effort to achieve the same sampling density. In response, some authors have proposed schemes for initial screening of a larger number of stressors, in order to determine those most relevant and produce a lower dimensional space more amenable to high-density sampling (Culley et al., 2021). A further question is whether all areas of the sampling space are equally important – for example, areas suitable for lower sampling density could include areas where system failure is rare or where stressor combinations are physically implausible or considered unlikely. More broadly, it raises the question of whether a uniform sampling approach remains appropriate. In different contexts, such as calibration of multi-parameter models, uniform sampling is usually avoided in favour of directional searches towards some "optimum" value (eg. Duan et al., 1992) or to explore the shape of some "behavioural" region (eg. Vrugt et al., 2008; 2013). It is possible that future researchers may adapt these techniques to direct available computing resources to resolve the parts of the space that matter most, such as better exploring the form of the failure surface and its boundaries.

In summary, the issues discussed above are: (i) whether to accept an "inherited" model or build one suited to the context; (ii) if building a model, how to tailor it to obtain the best tradeoff between model complexity and uncertainty estimation; and (iii) appropriate sampling of highly dimensional climate stress testing spaces. As the science of bottom-up climate risk assessment advances and matures, we suggest that explicit consideration and discussion of these issues should become more commonplace.

# 5   Conclusions

Although bottom-up climate stress tests are becoming more common, there are still relatively few frameworks for generating climate and streamflow data for multi-site many-dimensional

(more than two) stress tests. In this paper we have presented a freely available monthly-timestep climate stress testing framework. While most climate stress testing is focused on the daily timestep, monthly timesteps are suitable for systems whose dominant dynamic is seasonal or longer, which is the case for many water resource systems. Furthermore, the longer timestep means the framework can generate data rapidly, making extensive stress testing possible with a limited computational budget. This allows for the inclusion of stressor types such as modulating low frequency climatic behaviour, precipitation seasonality and allowing for future changes in rainfall-runoff relationship.

The case study demonstrated how the framework can be applied to a relatively complicated case study with varied climatic conditions across a relatively large spatial scale. The requirement for eight different sub-areas and four reach inflows was a suitable test to demonstrate the flexibility of the framework. The results indicate a good match between historic data and the baseline stochastic output, and in particular a pleasing replication of multi-year runs. Testing confirmed that the five stressors can be perturbed plausibly and simultaneously to cover the stress testing space, within a modest computational budget (less than 24 hours computing time on a 2016-era laptop).

The Matlab / Octave code is freely downloadable from https://doi.org/10.5281/zenodo.5617008 and is subject to a GNU General Public License v3.0.

# 6    Acknowledgements

# 7    References

Avanzi, F., J. Rungee, T. Maurer, R. Bales, O. Ma, S. Glaser, S., & M. Conklin, 2020. Climate elasticity of evapotranspiration shifts the water balance of Mediterranean climates during multi-year droughts. Hydrology and Earth System Sciences, 24, 4317–4337, https://doi.org/10.5194/hess-24-4317-2020

Bayr, T., Latif, M., Dommenget, D., Wengel, C., Harlaß, J., & Park, W., 2018. Mean-state dependence of ENSO atmospheric feedbacks in climate models. Climate Dynamics, 50(9-10), 3171-3194. Clim Dyn (2018) 50:3171–3194 https://doi.org/10.1007/s00382-017-3799-2

Bellenger, H., Guilyardi, E., Leloup, J., Lengaigne, M., & Vialard, J., 2014. ENSO representation in climate models: From CMIP3 to CMIP5. Climate Dynamics, 42(7-8), 1999–2018. https://doi.org/10.1007/s00382-013-1783-z

Bennett, B., Devanand, A., Culley, S., Westra, S., Guo, D., Maier, H., 2021. A modelling framework and R-package for evaluating system performance under hydroclimate variability and change, Environmental Modelling & Software, https://doi.org/10.1016/j.envsoft.2021.104999

Blöschl, G., Sivapalan, M., Savenije, H., Wagener, T., & Viglione, A. (Eds.), 2013. Runoff prediction in ungauged basins: synthesis across processes, places and scales. Cambridge University Press.

Box, G. E., & Cox, D. R., 1964. An analysis of transformations. Journal of the Royal Statistical Society: Series B (Methodological), 26(2), 211-243. https://doi.org/10.1111/j.2517-6161.1964.tb00553.x

Brown C., Steinschneider S., Ray P., Wi S., Basdekas L., Yates D., 2019. Decision Scaling (DS): Decision Support for Climate Change. In: Marchau V., Walker W., Bloemen P., Popper S. (eds) Decision Making under Deep Uncertainty. Springer, Cham. https://doi.org/10.1007/978-3-030-05252-2_12

Brown, C., & Wilby, R. L., 2012. An alternate approach to assessing climate risks. Eos, Transactions American Geophysical Union, 93(41), 401-402. https://doi.org/10.1029/2012EO410001

Cai, W., Wang, G., Dewitte, B., Wu, L., Santoso, A., Takahashi, K., et al., 2018. Increased variability of eastern Pacific El Niño under greenhouse warming. Nature, 564(7735), 201-206. https://doi.org/10.1038/s41586-018-0776-9

Chiew, F. H. S., J. Teng, J. Vaze, D. A. Post, J. M. Perraud, D. G. C. Kirono, and N. R. Viney, 2009, Estimating climate change impact on runoff across southeast Australia: Method, results, and implications of the modeling method, Water Resour. Res., 45, W10414, https://doi.org/10.1029/2008WR007338.

Christensen, N.S., Wood, A.W., Voisin, N. et al., 2004 The Effects of Climate Change on the Hydrology and Water Resources of the Colorado River Basin. Climatic Change 62, 337–363. https://doi.org/10.1023/B:CLIM.0000013684.13621.1f

Colominas, M.A., Schlotthauer, G., Torres, M.E. and Flandrin, P., 2012. Noise-assisted EMD methods in action. Advances in Adaptive Data Analysis, 4(04), p.1250025. https://doi.org/10.1142/S1793536912500252

Colominas, M. A., Schlotthauer, G., & Torres, M. E., 2014. Improved complete ensemble EMD: A suitable tool for biomedical signal processing. Biomedical Signal Processing and Control, 14, 19-29. https://doi.org/10.1016/j.bspc.2014.06.009. Code at: http://bioingenieria.edu.ar/grupos/ldnlys/metorres/metorres_files/ceemdan_v2014.m

Culley, S., Bennett, B., Westra, S. and Maier, H.R., 2019. Generating realistic perturbed hydrometeorological time series to inform scenario-neutral climate impact assessments. Journal of Hydrology, 576, pp.111-122. https://doi.org/10.1016/j.jhydrol.2019.06.005

Culley, S., Noble, S., Yates, A., Timbs, M., Westra, S., Maier, H. R., et al., 2016. A bottom-up approach to identifying the maximum operational adaptive capacity of water resource systems to a changing climate. Water Resources Research, 52(9), 6751-6768. https://doi.org/10.1002/2015WR018253

Culley, S., Maier, H. R., Westra, S., & Bennett, B. (2021). Identifying critical climate conditions for use in scenario-neutral climate impact assessments. Env. Modelling & Software, 136, 104948. https://doi.org/10.1016/j.envsoft.2020.104948

Docker, B. B., & Johnson, H. L., 2017. Environmental water delivery: Maximizing ecological outcomes in a constrained operating environment. Chapter 24 (pp. 563-598) in Horne et al., (eds.), Water For The Environment. Academic Press. https://doi.org/10.1016/B978-0-12-803907-6.00024-3

Doolan, J., Ashworth, B., Swirepik, J., 2017. Environmental water delivery: Maximizing ecological outcomes in a constrained operating environment. Chapter 23 (pp. 539-561) in Horne et al., (eds.), Water For The Environment. Academic Press. https://doi.org/10.1016/B978-0-12-803907-6.00023-1

Duan, Q., Sorooshian, S., & Gupta, V., 1992. Effective and efficient global optimization for conceptual rainfall-runoff models. Water resources research, 28(4), 1015-1031. https://doi.org/10.1029/91WR02985

Fowler, K., Peel, M., Western, A., Zhang, L. and Peterson, T., 2016. Simulating runoff under changing climatic conditions: Revisiting an apparent deficiency of conceptual rainfall-runoff models. Water Resources Research, 52, 1820–1846. https://doi.org/10.1002/2015WR018068

Fowler, K., Peel, M., Western, A., Zhang, L., 2018. Improved rainfall-runoff calibration for drying climate: Choice of objective function. Water Resources Research, 54. https://doi.org/10.1029/2017WR022466

Fowler, K., Knoben, W., Peel, M., Peterson, T., Ryu, D., Saft, M., Seo, K., Western, A., 2020a. Many commonly used rainfall-runoff models lack long, slow dynamics: Implications for runoff projections. Water Resources Research, 56, e2019WR025286. https://doi.org/10.1029/2019WR025286

Fowler, K.; Acharya, S.; Addor, N.; Chou, C.; Peel, M., 2020b. CAMELS-AUS: Hydrometeorological time series and landscape attributes for 222 catchments in Australia. Manuscript submitted to Environmental and Earth Systems Data, September 2020. Dataset freely downloadable from https://doi.pangaea.de/10.1594/PANGAEA.921850

François, B., Hingray, B., Borga, M., Zoccatelli, D., Brown, C., & Creutin, J. D., 2018. Impact of climate change on combined solar and run-of-river power in northern Italy. Energies, 11(2), 290.

Freeman, S. S. G., Brown, C., Cañada, H., Martinez, V., Nava, A. P., Ray, P., et al., 2020. Resilience by design in Mexico City: A participatory human-hydrologic systems approach. Water Security, 9, 100053. https://doi.org/10.1016/j.wasec.2019.100053

Garcia, L. E., Dawdy, D. R., & Mejia, J. M., 1972. Long memory monthly streamflow simulation by a broken line model. Water Resources Research, 8(4), 1100-1105. https://doi.org/10.1029/WR008i004p01100

García-Ruiz, J. M., López-Moreno, J. I., Vicente-Serrano, S. M., Lasanta–Martínez, T., & Beguería, S., 2011. Mediterranean water resources in a global change scenario. Earth-Science Reviews, 105(3-4), 121-139.

Ghile, Y. B., Taner, M. Ü., Brown, C., Grijsen, J. G., & Talbi, A., 2014. Bottom-up climate risk assessment of infrastructure investment in the Niger River Basin. Climatic change, 122(1-2), 97-110. https://doi.org/10.1007/s10584-013-1008-9

Grayson, R. B., Blöschl, G., Western, A. W., & McMahon, T. A., 2002. Advances in the use of observed spatial patterns of catchment hydrological response. Advances in Water Resources, 25(8-12), 1313-1334. https://doi.org/10.1016/S0309-1708(02)00060-X

Giudici, F., Castelletti, A., Giuliani, M., & Maier, H. R., 2020. An active learning approach for identifying the smallest subset of informative scenarios for robust planning under deep uncertainty. Environmental Modelling & Software, 127, 104681. https://doi.org/10.1016/j.envsoft.2020.104681

Guo, D., Westra, S., & Maier, H. R., 2018. An inverse approach to perturb historical rainfall data for scenario-neutral climate impact studies. Journal of Hydrology, 556, 877-890. https://doi.org/10.1016/j.jhydrol.2016.03.025

Gupta, H. V., Kling, H., Yilmaz, K. K., & Martinez, G. F., 2009. Decomposition of the mean squared error and NSE performance criteria: Implications for improving hydrological modelling. Journal of hydrology, 377(1-2), 80-91. https://doi.org/10.1016/j.jhydrol.2009.08.003

Helgeson, C., Srikrishnan, V., Keller, K., & Tuana, N., 2020. Why simpler computer simulation models can be epistemically better for informing decisions. Philosophy of Science, accepted article, https://doi.org/10.1086/711501

Henley, B. J., Peel, M. C., Nathan, R., King, A. D., Ukkola, A. M., Karoly, D. J., & Tan, K. S., 2019. Amplification of risks to water supply at 1.5° C and 2° C in drying climates: a case study for Melbourne, Australia. Environmental Research Letters, 14(8), 084028. https://doi.org/10.1088/1748-9326/ab26ef

Horne, A., Webb, A., Rumpff, L., Mussehl, M., Fowler, K & John, A., 2020. Kaiela (Lower Goulburn River) Environmental Flows Study. Report prepared by the University of Melbourne, Australia, for the Goulburn-Broken Catchment Management Authority, Australia. 24th November 2020.

Hrachowitz, M., H.H.G. Savenije, G. Blöschl, J.J. McDonnell, M. Sivapalan, J.W. Pomeroy, B. Arheimer, T. Blume, M.P. Clark, U. Ehret, F. Fenicia, J.E. Freer, A. Gelfan, H.V. Gupta, D.A. Hughes, R.W. Hut, A. Montanari, S. Pande, D. Tetzlaff, P.A. Troch, S. Uhlenbrook, T. Wagener, H.C. Winsemius, R.A. Woods, E. Zehe & C. Cudennec, 2013. A decade of Predictions in Ungauged Basins (PUB)—a review, Hydrological Sciences Journal, 58:6, 1198-1255, https://doi.org/10.1080/02626667.2013.803183

Huang, N. E., Shen, Z., Long, S. R., Wu, M. C., Shih, H. H., Zheng, Q., Yen, N., Tung, C. C. & Liu, H. H., 1998. The empirical mode decomposition and the Hilbert spectrum for nonlinear and non-stationary time series analysis Proc. R. Soc. Lond. A, 454 , pp. 903-995

Hurst, H. E., 1951. Long term storage capacities of reservoirs. Trans. ASCE 116: 776– 808.

Iliopoulou, T., Papalexiou, S. M., Markonis, Y., & Koutsoyiannis, D., 2018. Revisiting long-range dependence in annual precipitation. Journal of Hydrology, 556, 891-900. http://dx.doi.org/10.1016/j.jhydrol.2016.04.015

Jeffrey, S. J., J. O. Carter, K. B. Moodie, and A. R. Beswick, 2001, Using spatial interpolation to construct a comprehensive archive of Australian climate data, Environmental Modelling and Software, 16 (4), 309-330, https://doi.org/10.1016/S1364-8152(01)00008-1

John, A., Fowler K., Nathan R., Horne A. & Stewardson M., 2021a, Disaggregated monthly hydrological models can outperform daily models in providing daily flow statistics and extrapolate well to a drying climate Journal of Hydrology, 598, https://doi.org/10.1016/j.jhydrol.2021.126471.

John, A., Nathan, R., Horne, A., Stewardson, M., & Webb, J. A., 2021b. How to incorporate climate change into modelling environmental water outcomes: a review. Journal of Water and Climate Change, 11(2), 327-340.

John, A., Nathan R., Horne A., Fowler K., & Stewardson M., 2021c. Assessing sensitivities and interactions among many dimensions of climate change stressors on freshwater ecosystems. Submitted to Water Resources Research, January 2021.

Jones, D. A., Wang, W., & Fawcett, R., 2009. High-quality spatial climate datasets for Australia. Australian Meteorological and Oceanographic Journal, 58(4), 233–248. https://doi.org/10.22499/2.5804.003

Kiem, A. S., 2004. Multi-decadal variability of drought risk, eastern Australia, Hydrological Processes, 18 (11), 2039_2050, https://doi.org/ 10.1002/hyp.1460.Knoben, W. J., Freer, J. E., Peel, M. C., Fowler, K. J. A., & Woods, R. A. (2020). A brief analysis of conceptual model structure uncertainty using 36 models and 559 catchments. Water Resources Research, 56(9), e2019WR025975. https://doi.org/10.1029/2019WR025975

Koster, W., Crook, D., Dawson, D., and Moloney, P., 2012. Status of fish populations in the lower Goulburn River (2003-2012). Report prepared by the Arthur Rylah Institute of the Department of Sustainability and Environment (Victoria), https://www.gbcma.vic.gov.au/downloads/RH_Publications/Assessment_of_the_status_of_fish_stocks_in_the_lower_Goulburn_River_December_2012.pdf

Koutsoyiannis, D., 2005. Hydrologic persistence and the Hurst phenomenon. Water encyclopedia, 3, 210-221. https://doi.org/10.1002/047147844X.sw434

Lall, U., & Sharma, A., 1996. A nearest neighbor bootstrap for resampling hydrologic time series. Water Resources Research, 32(3), 679-693. https://doi.org/10.1029/95WR02966

Liu, Z., 2012. Dynamics of Interdecadal Climate Variability: A Historical Perspective, American Meteorological Society, 25, 1963-1995, https://doi.org/10.1175/2011JCLI3980.1

Matalas, N., 1967: Mathematical assessment of synthetic hydrology, Water Resour. Res., 3(4), 937–945, https://doi.org/10.1029/WR003i004p00937

McMahon TA, Pegram GGS, Vogel RM & Peel MC, 2007. Review of Gould-Dincer reservoir storage-yield-reliability estimates. Advances in Water Resources. 30, 1873-1882. https://doi.org/10.1016/j.advwatres.2007.02.004

McMahon, T. A., Kiem, A. S., Peel, M. C., Jordan, P. W., & Pegram, G. G., 2008. A new approach to stochastically generating six-monthly rainfall sequences based on empirical mode decomposition. Journal of Hydrometeorology, 9(6), 1377-1389. https:/doi.org/10.1175/2008JHM991.1

McMahon, T. A., Peel, M. C., Lowe, L., Srikanthan, R., and McVicar, T. R., 2013. Estimating actual, potential, reference crop and pan evaporation using standard meteorological data: a pragmatic synthesis, Hydrol. Earth Syst. Sci., 17, 1331–1363, https://doi.org/10.5194/hess-17-1331-2013

McMahon, T. A., Peel, M. C., & Karoly, D. J., 2015. Assessment of precipitation and temperature data from CMIP3 global climate models for hydrologic simulation. Hydrology and Earth System Sciences, 19(1), 361-377. http://www.hydrol-earth-syst-sci.net/19/361/2015/

McMahon, T. A., Finlayson, B. L. and Peel, M. C., 2016. Historical developments of models for estimating evaporation using standard meteorological data. WIRES - Water, 3(6), 788-818.

McPhail, C., Maier, H. R., Westra, S., Kwakkel, J. H., & van der Linden, L. (2020). Impact of scenario selection on robustness. Water Resources Research, 56(9), e2019WR026515. https://doi.org/10.1029/2019WR026515

Mejia, J. M., Rodriguez-Iturbe, I. and Dawdy, D. R., 1972. Streamflow simulation: 2. The broken line process as a potential model for hydrologic simulation. Water Resources Research 8 (4), 931-941. https://doi.org/10.1029/WR008i004p00931

Morton, F. I., 1983. Operational estimates of areal evapotranspiration and their significance to the science and practice of hydrology. Journal of Hydrology, 66(1-4), 1–76. https://doi.org/10.1016/0022-1694(83)90177-4

Mukundan, R., Acharya, N., Gelda, R. K., Frei, A., & Owens, E. M., 2019. Modeling streamflow sensitivity to climate change in New York City water supply streams using a stochastic weather generator. Journal of Hydrology: Regional Studies, 21, 147-158. https://doi.org/10.1016/j.ejrh.2019.01.001

Nathan, R. J., McMahon, T. A., Peel, M. C., & Horne, A., 2019. Assessing the degree of hydrologic stress due to climate change. Climatic Change, 156(1), 87-104. https://doi.org/10.1007/s10584-019-02497-4

Peel, M. C., G. G. S. Pegram, and T. A. McMahon, 2005. Global analysis of runs of annual precipitation and runoff equal to or below the median: Run magnitude and severity, International Journal of Climatology, 24 (7), 549_568, https://doi.org/10.1002/joc.1147.

Peterson, T. J., Saft, M., Peel, M. C., & John, A. (2021). Watersheds may not recover from drought. Science, 372(6543), 745-749. https://doi.org/ 10.1126/science.abd5085

Poff, N. L., Allan, J. D., Bain, M. B., Karr, J. R., Prestegaard, K. L., Richter, B. D., et al., 1997. The natural flow regime. BioScience, 47(11), 769-784. https://doi.org/10.2307/1313099

Poff, N. L., Brown, C. M., Grantham, T. E., Matthews, J. H., Palmer, M. A., Spence, C. M., et al., 2016. Sustainable water management under future uncertainty with eco-engineering decision scaling. Nature Climate Change, 6(1), 25-34. https://doi.org/10.1038/nclimate2765

Prudhomme, C., Wilby, R. L., Crooks, S., Kay, A. L., & Reynard, N. S., 2010. Scenario-neutral approach to climate change impact studies: application to flood risk. Journal of Hydrology, 390(3-4), 198-209. http://doi.org/10.1016/j.jhydrol.2010.06.043

Saft, M., Western, A. W., Zhang, L., Peel, M. C., & Potter, N. J., 2015. The influence of multiyear drought on the annual rainfallrunoff relationship: An Australian perspective. Water Resources Research, 51(4), 2444–2463. https://doi.org/10.1002/2014WR015348

Skinner, D., & Langford, J., 2013. Legislating for sustainable basin management: the story of Australia's Water Act (2007). Water Policy, 15(6), 871-894.

Srikanthan, R. and McMahon, T. A., 2001. Stochastic generation of annual, monthly and daily climate data: A review, Hydrol. Earth Syst. Sci., 5, 653–670, https://doi.org/10.5194/hess-5-653-2001, 2001.

Steinschneider, S., & Brown, C., 2013. A semiparametric multivariate, multisite weather generator with low-frequency variability for use in climate risk assessments. Water resources research, 49(11), 7205-7220. https://doi.org/10.1002/wrcr.20528

Thyer, M., & Kuczera, G., 2000. Modeling long-term persistence in hydroclimatic time series using a hidden state Markov Model. Water resources research, 36(11), 3301-3310. https://doi.org/10.1029/2000WR900157

Tonkin, J. D., Bogan, M. T., Bonada, N., Rios-Touma, B., & Lytle, D. A., 2017. Seasonality and predictability shape temporal species diversity. Ecology, 98(5), 1201-1216. https://doi.org/10.1002/ecy.1761

Torres ME, Colominas MA, Schlotthauer G, Flandrin P, 2011. A Complete Ensemble Empirical Mode Decomposition with Adaptive Noise" Proc. 36th Int. Conf. on Acoustics, Speech and Signa Processing ICASSP (May 22-27, Prague, Czech Republic)

Topalović, Ž., A. Todorović & J. Plavšić, 2020. Evaluating the transferability of monthly water balance models under changing climate conditions, Hydrological Sciences Journal, 65:6, 928-950, https://doi.org/10.1080/02626667.2020.1725238

Trotter, L., Saft, M., Peel, M., Fowler, K., 2021: "Naïve" inclusion of diverse climates in calibration is not sufficient to improve model reliability under future climate uncertainty. 24th International Congress on Modelling and Simulation (MODSIM2021), Sydney, Australia, December 2021.

Turner, M., Bari, M., Amirthanathan, G., & Ahmad, Z., 2012. Australian network of hydrologic reference stations-advances in design, development and implementation. In Hydrology and Water Resources Symposium 2012 (p. 1555). Engineers Australia, http://www.bom.gov.au/water/hrs/media/static/papers/Turner2012.pdf

Turner, S. W., Marlow, D., Ekström, M., Rhodes, B. G., Kularathna, U., & Jeffrey, P. J., 2014. Linking climate projections to performance: A yield-based decision scaling assessment of a large urban water resources system. Water Resources Research, 50(4), 3553-3567. https://doi.org/10.1002/2013WR015156

Wang, Q. J., Pagano, T., Zhou, S., Hapuarachchi, H., Zhang, L. and D. Robertson, 2011. Monthly versus daily water balance models in simulating monthly runoff. Journal of Hydrology 404 (2011) 166–175, https://doi.org/10.1016/j.jhydrol.2011.04.027.

Watts, R. J., Richter, B. D., Opperman, J. J., & Bowmer, K. H., 2011. Dam reoperation in an era of climate change. Marine and Freshwater Research, 62(3), 321-327. https://doi.org/10.1071/MF10047

Weaver, C. P., Lempert, R. J., Brown, C., Hall, J. A., Revell, D., & Sarewitz, D., 2013. Improving the contribution of climate model information to decision making: the value and demands of robust decision frameworks. Wiley Interdisciplinary Reviews: Climate Change, 4(1), 39-60. https://doi.org/10.1002/wcc.202

Wei, Y., Yu, H., Huang, J., He, Y., Yang, B., Guan, X., & Liu, X., 2018. Comparison of the Pacific Decadal Oscillation in climate model simulations and observations. International Journal of Climatology, 38, e99-e118. https://doi.org/10.1002/joc.5355

Whateley, S., Steinschneider, S., & Brown, C., 2014. A climate change range-based method for estimating robustness for water resources supply. Water Resources Research, 50(11), 8944-8961. https://doi.org/10.1002/2014WR015956

Whetton, P., Hennessy, K., Clarke, J. et al., 2012. Use of Representative Climate Futures in impact and adaptation assessment. Climatic Change 115, 433–442 (2012). https://doi.org/10.1007/s10584-012-0471-z

Whiting, J. P., Lambert, M. F., & Metcalfe, A. V., 2003. Modelling persistence in annual Australia point rainfall. Hydrology and Earth System Sciences, 7(2), 197-211. https://doi.org/10.5194/hess-7-197-2003

Wilby, R., & Murphy, C., 2019. Decision-making by water managers despite climate uncertainty. In Pfeffer, W., Smith, J. et al. (ed.), The Oxford Handbook of Planning for Climate Change Hazards. http://doi.org/10.1093/oxfordhb/9780190455811.013.52

Wilks, D. S., & Wilby, R. L., 1999. The weather generation game: a review of stochastic weather models. Progress in physical geography, 23(3), 329-357. https://doi.org/10.1177/030913339902300302

Xu, C. Y., & Singh, V. P., 1998. A review on monthly water balance models for water resources investigations. Water resources management, 12(1), 20-50. https://doi.org/10.1023/A:1007916816469

Yang, Y. E., Wi, S., Ray, P. A., Brown, C. M., & Khalil, A. F., 2016. The future nexus of the Brahmaputra River Basin: Climate, water, energy and food trajectories. Global Environmental Change, 37, 16-30. https://doi.org/10.1016/j.gloenvcha.2016.01.002