Journal of Hydrology 406 (2011) 97–112

Contents lists available at ScienceDirect

Journal of Hydrology

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

Attribution of Autumn/Winter 2000 flood risk in England to anthropogenic climate change: A catchment-based study

A.L. Kay [a,*], S.M. Crooks [a], P. Pall [b,c], D.A. Stone [b,d,e]

[a] Centre for Ecology and Hydrology, Maclean Building, Crowmarsh Gifford, Wallingford, Oxfordshire OX10 8BB, UK
[b] AOPP, Department of Physics, University of Oxford, Clarendon Laboratory, Parks Road, Oxford OX1 3PU, UK
[c] Institute for Atmospheric and Climate Science, ETH Zurich, 8092 Zurich, Switzerland
[d] Tyndall Centre for Climate Change Research, Oxford University Centre for the Environment, Oxford, UK
[e] CSAG, Shell Environmental and Geographical Science Building, South Lane, Upper Campus, University of Cape Town, Private Bag X3, Rondebosch, Western Cape 7701, South Africa

ARTICLE INFO

Article history:
Received 17 February 2011
Received in revised form 4 May 2011
Accepted 10 June 2011
Available online 25 June 2011
This manuscript was handled by K. Georgakakos, Editor-in-Chief, with the assistance of Ashish Sharma, Associate Editor

Keywords:
Climate change
Hydrological impacts
Flood frequency
Attribution
Uncertainty

SUMMARY

Although no single weather-related event can be directly attributed to climate change, new techniques make it possible to estimate how much the chance of an event has been altered by anthropogenic emissions. This paper looks at the floods that occurred in England in Autumn/Winter 2000, by using large ensembles of 1-year climate model simulations representing April 2000–March 2001. These represent an industrial climate and four estimates of an hypothetical non-industrial climate (without historical greenhouse gas emissions), and are used to drive hydrological models for eight catchments in England. The simulated flows are used to assess the impact of historical emissions on the chance of occurrence of extreme floods in each catchment, through calculation of the fraction of attributable risk (*FAR*). Combining results for the four non-industrial climates, positive median values of *FAR* indicate that, for all but one catchment, emissions are likely to have led to an increased chance of flooding in the October–December period. Definitive conclusions are difficult however, as there are wide bands of uncertainty in *FAR*, with distributions generally spanning no attributable difference in risk (*FAR* = 0). One catchment shows a decreased flood chance (negative median *FAR*), due to its high permeability, but an analysis of the effect of antecedent conditions shows that a longer period of climate data than 1 year is probably required to obtain more representative values of *FAR* for such catchments. The inclusion of snowfall/snowmelt is also shown to be important for floods over the October–March period, as the reduced likelihood of snowmelt-induced floods in the warmer temperatures of the industrial climate moderates the increased flood chance due to other sources of flooding.

© 2011 Elsevier B.V. All rights reserved.

1. Introduction

There is an increasing consensus among the global scientific community that our climate is changing and that this is mostly due to anthropogenic emissions of greenhouse gases (IPCC, 2007). Allied to this is an increasing concern about how the changing climate will impact upon our natural and built environment. Although no single event can be attributed, solely and directly, to climate change, new techniques make it possible to say something about how much the chance of an event has been affected by humanity's emissions of greenhouse gases (Allen, 2003; Stone and Allen, 2005). For instance, Stott et al. (2004) estimated that, with high likelihood, the chance of a European heatwave similar to that of Summer 2003 had been at least doubled by anthropogenic emissions.

This paper looks at the flood events that occurred in England in Autumn 2000. A report was commissioned by the UK Department for Environment Farming and Rural Affairs (Defra) shortly afterwards to consider the possibility that climate change had contributed to the flooding (CEH and Met Office, 2001). It concluded that the floods and rainfall of October/November 2000 were unusual in a historical context, and said that they were consistent with the potential effects of climate change but could not be directly attributed to them as they could not be distinguished from natural variability. A later study (Hall et al., 2005) aimed to consider the evidence presented in the above report more formally, and came to similar conclusions.

The recent "Seasonal Attribution Project" (http://climateprediction.net/content/seasonal-attribution-experiment) of the climate *prediction*.net experiment used public distributed computing to produce a large ensemble of climate model simulations, in order to compare the risk of the record wet autumn in the year 2000 with and without 20th Century anthropogenic greenhouse gas emissions. It concluded (Pall, 2006) that there was some evidence of an increased

* Corresponding author. Tel.: +44 01491 838800; fax: +44 01491 692424.
E-mail address: alkay@ceh.ac.uk (A.L. Kay).

0022-1694/$ - see front matter © 2011 Elsevier B.V. All rights reserved.
doi:10.1016/j.jhydrol.2011.06.006

chance of high autumn precipitation totals over England and Wales due to emissions. Furthermore, after feeding the ensembles of daily precipitation data into a simple statistical rainfall–runoff model to simulate daily river runoff for England and Wales catchments (Pall et al., 2011), it was concluded that the England and Wales flood risk for Autumn 2000 (i.e. high runoff during autumn) had significantly increased: the best estimate (median) of the fraction of risk attributable to anthropogenic emissions was approximately 0.6. That is, the emissions had most likely increased the chance of such flooding by a factor of about 2.5, as the risk would only have been about 40% of its current level had anthropogenic emissions not occurred.

In this paper, data resulting from the ensemble experiments of Pall et al. (2011) are fed into continuous simulation rainfall–runoff models calibrated for eight catchments in England, four in the South–East and four in the North–East, each affected by flooding in Autumn 2000. The aim is to assess how much the chance of such flooding in each of the catchments has been influenced by anthropogenic emissions, and whether the results are dependent upon catchment characteristics. The use of continuous simulation models makes the results more robust to temporal and spatial variation of rainfall inputs and to antecedent conditions, so that differences due to catchment characteristics and location are better accounted for. In addition, a snowmelt module is applied, as the increased temperatures due to climate change are likely to mean a decreased chance of large snowmelt-induced flood events, and the effect of initial conditions is investigated.

A summary of the key features of the autumn 2000 flooding in England, and the conditions that led up to it, is given in Section 2, along with details on the catchments modelled. Section 3 describes the climate model ensemble data, the hydrological models, and the methodology followed. The results are presented and discussed in Section 4, with a summary and conclusions in Section 5.

2. The Autumn 2000 floods in England

2.1. Key features

This information is taken from CEH and Met Office (2001) and Marsh and Dale (2002).

2.1.1. Antecedent conditions

The winter of 1999–2000 (December–February) was wetter than average, as was Spring 2000 (March–May). But Summer 2000 (June–August) was drier than average, so late summer soil moisture deficits (SMDs; calculated using MORECS, the Met Office Rainfall and Evaporation Calculation System, Thompson et al., 1982) were close to the seasonal average in most lowland (southern and eastern) areas, and mostly well below average in western and northern areas, where only small deficits were carried into the autumn. Measured groundwater levels were above average throughout the period, due to several recent wet winters.

2.1.2. September 2000

Rainfall in September quickly reduced the lowland SMDs.

2.1.3. October and November 2000

There were a number of heavy rainfall events in early Autumn, with totals for the months of October and November generally amounting to two to three times the 1961–1990 average for those months over England and Wales. The rainfall totals for these two months combined were estimated to have a return period of 200 years or more over large parts of the country, with a number of individual rainfall events within the period also having return periods in excess of 100 years for certain parts of the country. This led to flooding affecting many different types of catchment: Larger basins with a groundwater component to their flows experienced extended periods of high flows, whilst smaller (impervious) catchments experienced sequences of high flow events. However, overall the floods were "notable more for their extent and duration ... than their magnitude" (CEH and Met Office, 2001). Flooding affected much of the country at times during the last 3 months of 2000—the most extensive flooding for England and Wales since the snowmelt-generated floods of March 1947. The first significant flooding occurred in the second week of October, with some catchments still reaching peak flows in December (and well into 2001 in restricted cases).

2.2. Catchment details

Modelling is performed for eight catchments in England, four in the South–East and four in the North–East (Fig. 1 and Table 1), each of which was affected to some extent by flooding in Autumn 2000. Four physical properties for each catchment are given in Table 1; area, standard average annual rainfall (SAAR), elevation (mean and range: minimum–maximum) and percentage of the catchment with high permeability bedrock (BHP). Each region includes catchments with a range of areas and permeabilities, while those in the North–East have higher rainfall and higher mean and maximum altitude than those in the South–East. The range of elevation provides an indicator of slope, related to catchment response time. BHP is an indicator of how much of the river flow comes from stored groundwater sources: catchments with higher BHP typically have a longer response time and longer hydrological memory – flows are affected by rainfall and evaporation over preceding seasons and, in extreme situations, preceding years. How the climatic attributes combine with physical properties, unique to each catchment, determines how the flow regime responds to extreme events and changes in climate.

Fig. 2 shows, for each catchment, the observed flows over the period April 2000–March 2001 together with the mean, maximum and minimum flows from any observed data available from 1961 (see Table 1 for start year of flow record) up to April 2000, at three durations (1-day, 10-day and 30-day). These plots demonstrate the extremity of the flows in Autumn/Winter 2000 relative to flows in the recent past, particularly at longer durations. Fig. 2 also shows, for catchment 39001 (Thames @ Kingston), maximum and minimum flows for observed data from 1883. These demonstrate that, although Autumn 2000 was extreme relative to the period 1961–1999, it was not unprecedented in the longer historical context. Note that naturalised flows have been used for catchment 39001; this is the gauged river flow adjusted to take account of net abstractions and discharges upstream of the gauging station. Gauged flows have been used for all other catchments, where naturalised flows are unavailable.

3. Models, data and methods

3.1. The hydrological models and the snowmelt module

Two hydrological models are used; the Probability Distributed Model (PDM; Moore 1985, 2007), generally used for smaller catchments, and the Climate and LAnd-use Scenario Simulation In Catchments model (CLASSIC; Crooks and Naden, 2007), generally used for larger catchments. Here, CLASSIC is used for the two largest catchments, the Thames (39001) and Ouse (27009), whilst the PDM is used for the six smaller catchments (see Table 1). Both models are run at a daily time-step, which is considered appropriate given the area and responsiveness of each modelled catchment.

The PDM is a lumped rainfall–runoff model, requiring inputs of catchment-average rainfall and potential evaporation (PE; the

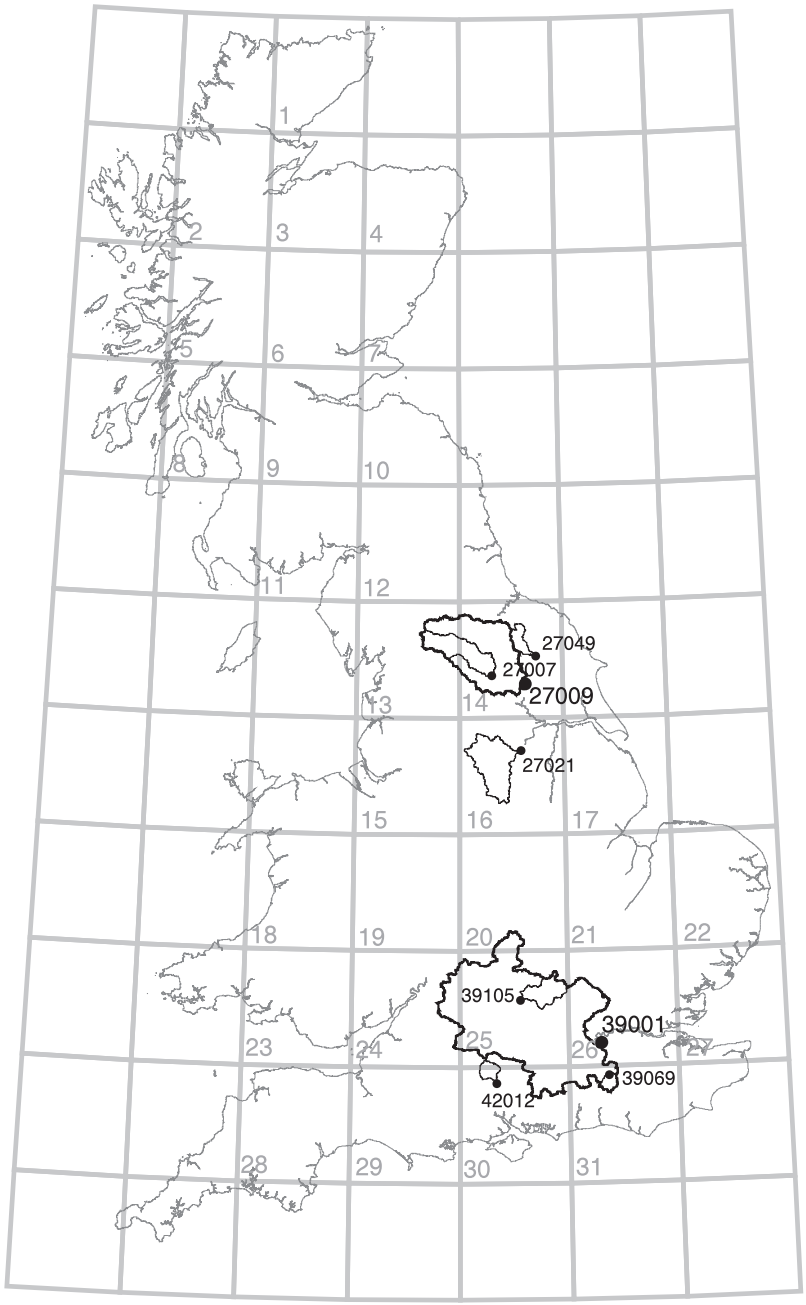

Fig. 1. Locations of the eight catchments in Britain (black) overlaid with the HadAM3-N144 grid (grey). Each catchment boundary includes a dot showing the gauge location, and is labelled with the catchment number. The two larger catchments (27009 and 39001) are shown with thicker boundaries and bigger dots. The 31 grid boxes over Britain classified as land are labelled with the numbers 1–31, in the bottom-left of each box (for reference).

maximum possible amount of evaporation from the land surface, given idealised conditions). It has three conceptual stores; a soil moisture store, and fast and slow flow stores. The soil moisture store accounts for rainfall inputs and evaporative losses, with content of the soil store determining the proportion of PE that actually evaporates. Although the full PDM has a variety of formulations, the version used here has been simplified in order to reduce the problem of equifinality (where multiple parameter sets can result in similar performance) and allow automatic (rather than manual) calibration. The version used has six catchment-specific parameters; two are determined by catchment location, one is set using soils data, and the remaining three require calibration using observed data. Brief information on calibration is given below; see Crooks et al. (2009) for more detail on this version of the model and its calibration, which are refinements of those described by Kay et al. (2007).

The semi-distributed continuous simulation rainfall–runoff model CLASSIC was developed for estimating the impacts of climate and land use change on flows in large catchments, with an emphasis on medium to high flows. The model is applied on a grid square framework and comprises three main modules (soil moisture accounting, drainage and channel routing). The model has a grid to outlet structure in which the simulated runoff from each grid square is routed directly to the catchment outlet rather than through successive grid squares. Calibration is semi-automatic, requiring percentages of six land use types and soil classes for each grid with physiographic information on hill-slope, altitude and the channel network derived from a digital terrain model (DTM).

Table 1
Details of the eight catchments, in two regions of England.

Catchment number	Catchment name (River @ Location)	Catchment area (km^2)	$SAAR_{61-90}$ (mm)	Mean elevation (masl)	Elevation range (masl)	BHP (%)	Start year of flow record
North–East England							
27007	Ure @ Westwick Lock	914.6	1118	264	14–710	15	1961
27021	Don @ Doncaster	1256.2	799	124	4–543	2	1961
27049	Rye @ Ness	238.7	839	227	26–452	0	1974
27009	Ouse @ Skelton	3315.0	914	118	5–714	34	1969
South–East England							
39069	Mole @ Kinnersley Manor	142.0	795	74	48–175	0	1976
39105	Thame @ Wheatley	533.8	644	88	50–268	15	1989
42012	Anton @ Fullerton	185.0	773	103	41–253	100	1975
39001	Thames @ Kingston	9948.0	719	100	5–330	43	1883

$SAAR_{61-90}$: standard average annual rainfall for 1961–1990.
BHP: percentage of catchment underlain by high permeability bedrock.

Climatic inputs of rainfall and PE are required for each grid square, with flow data for the outlet point for final calibration. The grid square size is chosen to be appropriate for the catchment area and the variation of climatic and physiographic conditions within the catchment.

Both models were calibrated using data from the beginning of the flow record (or January 1961 if later; see Table 1) to December 2001, to simulate the full flow range. Different objective functions were used within the calibration procedure, as appropriate to the hydrological role of each parameter; functions used include volume error over monthly and annual time scales, fit between observed and simulated flows weighted to high or low flows and fit of the flood frequency curve. All available flow data were used in the calibration to allow the fit of the flood frequency curve to be reasonably assessed.

A snowmelt module can be used as a pre-processor for the rainfall inputs to both models, essentially delaying the input of water if snowfall occurs. The module, devised by Bell and Moore (1999), uses a simple temperature-related snow store and melt rate and has separate accounting in different elevation zones. The module thus requires a single time-series of temperature data, the altitude to which the temperature relates, and information on catchment area within the different elevation zones (which have a 10 m height). The IHDTM (Morris and Flavin, 1990), which has a 50 m horizontal resolution and a 1 m vertical resolution, was used to derive the latter information. A fixed lapse rate is used to derive the temperature in each elevation zone from the single temperature time-series and its corresponding altitude. See Crooks et al. (2009) for more detail on calibration of the snowmelt module coupled with the hydrological models.

3.2. Attribution system

3.2.1. Description

The Seasonal Attribution Project was conducted to examine the attributable component of the England and Wales Autumn 2000 flood risk. It used the climate*prediction*.net distributed computing facility to produce an unprecedentedly large ensemble of runs from the global atmospheric model HadAM3 run at approximately 100 km resolution (HadAM3-N144; Pope and Stratton, 2002); this version of the model is considered to have an improved representation of storm track behaviour, with the storm track being the dominant large-scale circulation feature in the generation of precipitation over the UK in autumn and winter. This ensemble consisted of five sub-ensembles; one representative of year 2000/2001 conditions including the effect of historical emissions of greenhouse gases (termed "industrial", hereafter "IA2000"), and four representative of what the year 2000/2001 conditions might have been if there had been no 20th Century emissions of greenhouse gases (termed "non-industrial", hereafter "NIA2000").

Table 2 gives the number of usable runs available for each of the five sub-ensembles. All ensembles are driven with observed weekly varying sea surface temperatures (SSTs) from the National Oceanic and Atmospheric Administration's Optimum Interpolation version 2 dataset (Reynolds et al., 2002). The SSTs used in each of the NIA2000 ensembles (a–d) are modified by the subtraction of an estimate of the SST warming pattern (ΔSST) attributable to historical anthropogenic greenhouse gas emissions. These ΔSSTs estimates are derived from different global atmosphere–ocean climate models (a – HadCM3, b – GFDLR30, c – NCARPCM and d – MIROC3.2). Each of the four NIA2000 ensembles is further divided into 10 subsets, which use different estimates of the amplitude of the respective ΔSST spatial pattern (hereafter the ΔSST scaling factor), derived through comparison against the historical record according to an optimal fingerprinting regression method (Nozawa et al., 2005; Stott et al., 2006); these 10 scaling factors evenly sample the distribution of amplitudes of the ΔSST pattern permitted by the observational record constraints. Consistent sea ice coverage is estimated using a simple scaling (see Pall et al., 2011 details). Thus there are effectively 40 NIA2000 subsets, reflecting the 40 different SST (and sea ice) boundary conditions for the atmospheric model simulations. Each run within a subset represents the use of a different set of initial conditions, as does each run within the IA2000 ensemble. Further details of the experimental setup of these atmospheric model simulations are described in Pall et al. (2011).

Every ensemble run simulates the period 1st April 2000–30th March 2001 (although the climate model year in this case consists of 360 days, with 12 months each of 30 days). Daily total precipitation and daily mean temperature data are available for various regions including Britain. The map in Fig. 1 shows the climate model grid over Britain, with the 31 boxes classified as 'land' numbered. It is the precipitation and temperature data from some of the 19 land boxes over England and Wales (numbered 13–31 in Fig. 1) that are used in this paper.

3.2.2. Use of ensemble data to drive hydrological models

The main data input required by the hydrological models (Section 3.1) is rainfall time-series, which are available directly from the ensembles. The second data input required by the hydrological models is PE time-series. However, PE data are not available directly from the ensembles, and neither are all of the atmospheric variables required to estimate the generally preferred Penman–Monteith formulation of PE (Monteith, 1965). Instead, PE has been estimated from the available daily mean temperature (T) time-series using the method of Oudin et al. (2005), which has been shown to work relatively well using climate model temperature data over Britain (Kay and Davies, 2008). The daily mean temperature data

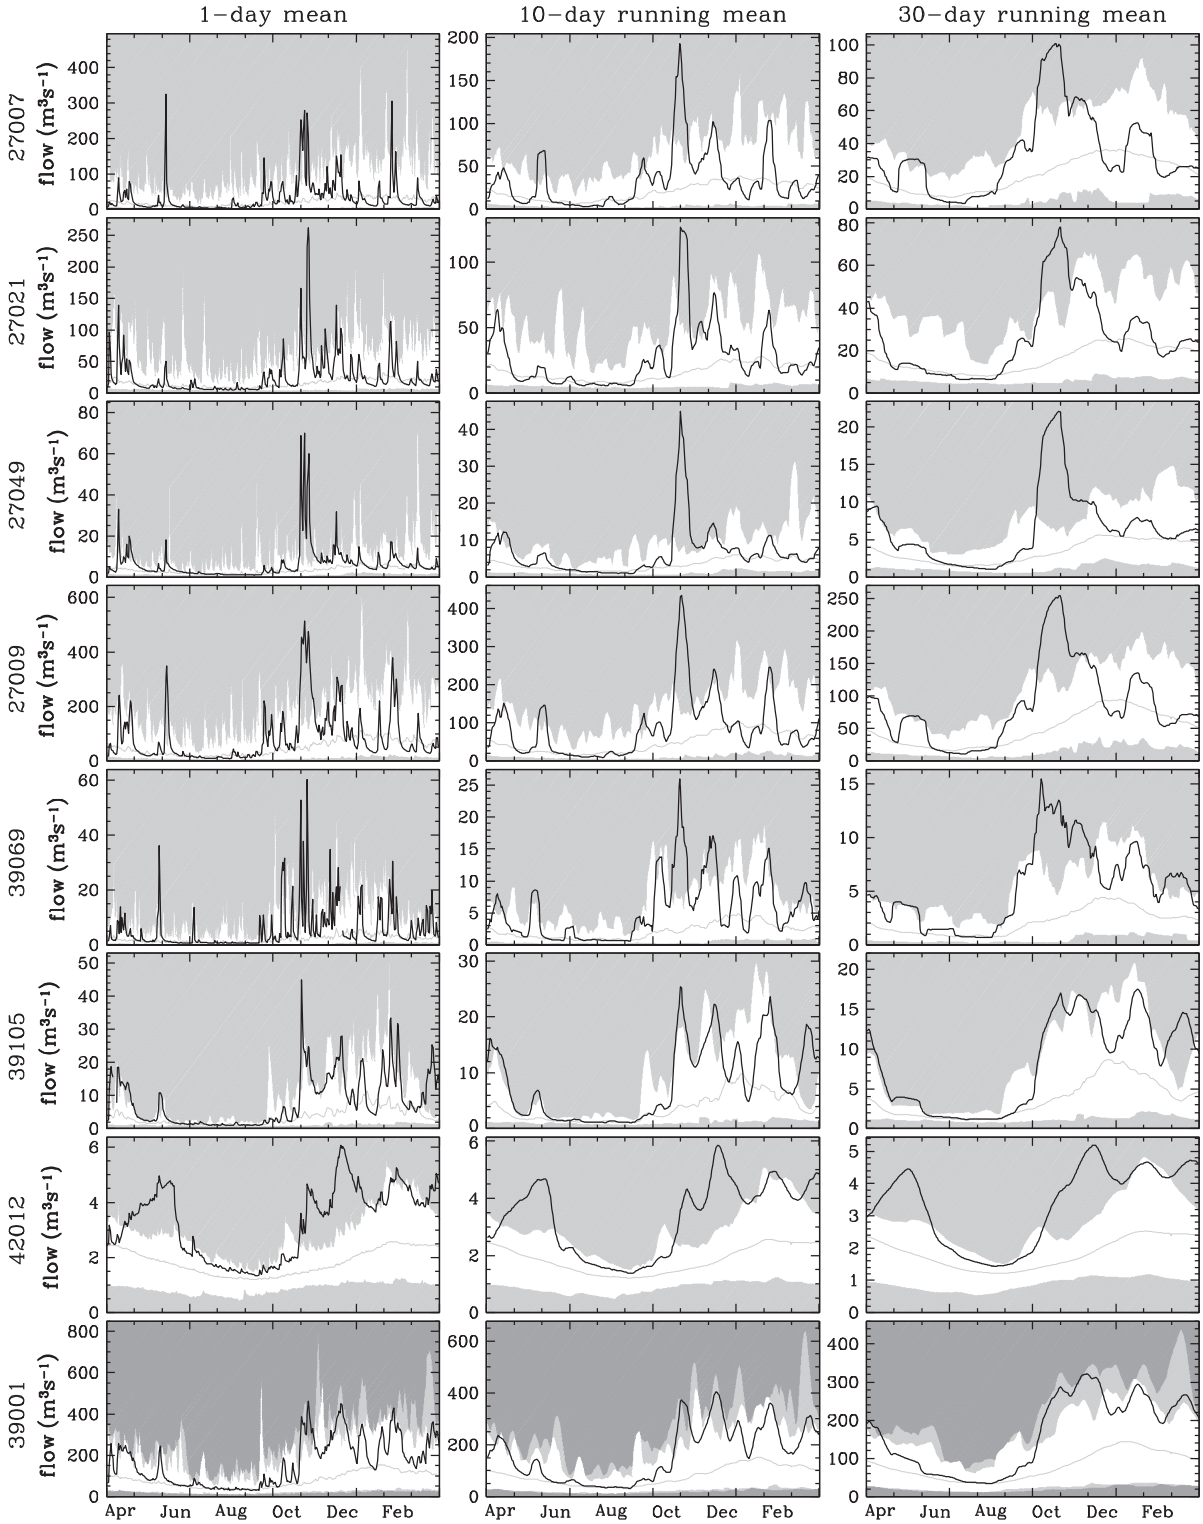

Fig. 2. Observed flows in the eight catchments over the period April 2000–March 2001 (black line), together with mean, maximum and minimum flows (grey line and light grey shading) from any observed data available from 1961 (see Table 1) up to April 2000, at three durations. For catchment 39001, maximum and minimum flows are also shown for observed data from 1883 (darker grey shading). Naturalised flows have been used for catchment 39001; gauged otherwise.

are also required by the snowmelt module (Section 3.1), as is information on the altitude to which the temperature data relate. The latter is available from the orography file of the climate model.

The gridded climate model rainfall and PE data are converted into catchment-average data for input to the PDM, or interpolated onto a finer scale grid for input to CLASSIC, using area-weighting, with additional weighting according to SAAR (standard average annual rainfall) for rainfall (Kay et al., 2006). The SAAR weights add spatial structure to the climate model rainfall data, by allowing for the fact that some areas can receive consistently more rainfall than others (due to topography or exposure for example). Table 3 gives the climate model grid squares and area and SAAR weights

Table 2
The five ensembles of HadAM3-N144 model runs.

Ensemble name	Description	Number of runs available
IA2000	Industrial Autumn 2000	2271
NIA2000a	Non-industrial Autumn 2000, with SSTs adjusted using HadCM3 ΔSSTs	2158[a]
NIA2000b	Non-industrial Autumn 2000, with SSTs adjusted using GFDLR30 ΔSSTs	2159[a]
NIA2000c	Non-industrial Autumn 2000, with SSTs adjusted using NCARPCM1 ΔSSTs	2174[a]
NIA2000d	Non-industrial Autumn 2000, with SSTs adjusted using MIROC3.2 ΔSSTs	2073[a]

[a] These ensembles each to be further divided, into 10 subsets, according to the value of the ΔSST scaling factor used for each run.

Table 3
HadAM3-N144 model grid squares, with their corresponding area and SAAR weights, for each of the eight catchments.

Catchment number	Grid square numbers (see Fig. 1)	Area proportions	SAAR weights
North–East England			
27007	13, 14	0.48, 0.52	0.81, 1.53
27021	16	1.00	1.12
27049	14	1.00	1.20
27009	13, 14	0.25, 0.75	0.66, 1.26
South–East England			
39069	31	1.00	1.02
39105	25	1.00	0.96
42012	25, 30	0.28, 0.72	1.16, 0.99
39001	20, 24, 25, 26, 30, 31	0.02, 0.06, 0.65, 0.15, 0.06, 0.06	1.12, 0.87, 1.05, 1.15, 0.90, 0.91

for each catchment. Note that although these numbers are given for all eight catchments, they are not used directly for the two larger catchments (27009 and 39001), since the CLASSIC model requires the information for each hydrological model grid square rather than for the catchment as a whole. For each catchment, the climate model temperature data are simply taken from the climate model grid box that contains the largest proportion of the catchment's area (see Table 3), and used along with that grid box's altitude (taken from the orography file of the climate model).

3.2.3. Comparison of ensemble data over England and Wales

This section presents a comparison of the distributions of annual (and seasonal) rainfall, daily mean temperature and PE between the IA2000 and NIA2000 ensembles for each of the 19 'land' climate model boxes over England and Wales (numbered 13–31 in Fig. 1). It also compares the distributions from the IA2000 ensemble with gridded observed data (available on a 5 km × 5 km grid over the UK; see Annex 1 of Jenkins et al., 2007), interpolated onto the climate model grid. The latter comparison can give some indication of the characteristics of the climate model data compared to observations. However, it would not be appropriate to use the results of this comparison for the purposes of bias-correction. This is because the IA2000 ensemble is designed to represent possible weather patterns specific to the period April 2000–March 2001, and is thus run with SSTs and other boundary conditions specific to that period. This specific set of boundary conditions may be conducive to certain anomalous climatic conditions relative to the period of the available observed data (April 1961–March 2001) (Blackburn and Hoskins, 2001; Massacand, 2003). The comparison of the distributions from the IA2000 ensemble and those from observed data for April 1961–March 2001 must therefore be interpreted carefully.

Fig. 3a compares the cumulative distribution (cdf) of annual total rainfall from each ensemble within each of the 19 'land' climate model boxes over England and Wales. This shows that there is generally little difference between the IA2000 ensemble and any of the NIA2000 ensembles, except for NIA2000c which is slightly shifted rightwards of the other cdfs, indicating higher rainfall, particularly for more westerly grid boxes. For spring (March–May) and autumn (September–November) rainfall (not shown) there is very little difference between the cdfs of any of the ensembles. For summer (June–August) rainfall, the cdfs for each of the NIA2000 ensembles are slightly shifted rightwards compared to the cdf of the IA2000 ensemble in most of the 19 climate model grid boxes (not shown), indicating slightly higher summer rainfall in a non-industrial climate. For winter (December–February) rainfall, there is more difference between the NIA2000 cdfs, with that for sub-ensemble NIA2000c generally being shifted rightwards from the cdf for the IA2000 ensemble, which is in turn slightly rightwards of the cdfs for the other NIA2000 ensembles in most of the 19 climate model grid boxes (not shown). This indicates slightly higher winter rainfall in the NIA2000c climate and perhaps slightly lower winter rainfall in the other NIA2000 climates compared to the IA2000 climate. Note that in Fig. 3 (and later in Fig. 4) the cdfs shown are for all 10 ΔSST scaling factors lumped together for each NIA2000 sub-ensemble, so each NIA2000 cdf represents the most likely climate for that ΔSST spatial pattern and does not show the range of possible cdfs within that climate due to the ΔSST scaling factor.

Also shown for each climate model grid box in Fig. 3a is the cdf of annual total rainfall from observed data for April 1961–March 2001 (available on a 5 km × 5 km grid, but interpolated onto the climate model grid using area and SAAR weighting, as described for the catchment-average derivation in Section 3.2.2). The comparison is generally quite good, particularly for grid boxes not on the south or west coasts. The same is true for rainfall in each season (not shown). Note that the comparison for coastal grid boxes is complicated by the lack of observed data for the sea parts of the box (for example for box number 23), even though the SAAR weighting should partially compensate for this in the case of the rainfall comparison. Fig. 3a also shows how extreme the rainfall in April 2000–March 2001 was, as the vertical line showing this value is positioned far into the upper percentiles of the distribution of observed rainfall. The same is true for Autumn 2000 rainfall (not shown).

Fig. 4 compares the daily rainfall extremes (above the 99th percentile) in autumn for the IA2000 and NIA2000 ensembles. It shows lower extreme autumn daily rainfall totals in each of the NIA2000 sub-ensembles compared to the IA2000 ensemble. The difference for the 99.5 percentile, averaged across the 19 boxes, ranges between −4% (for NIA2000c) and −8% (for NIA2000a and

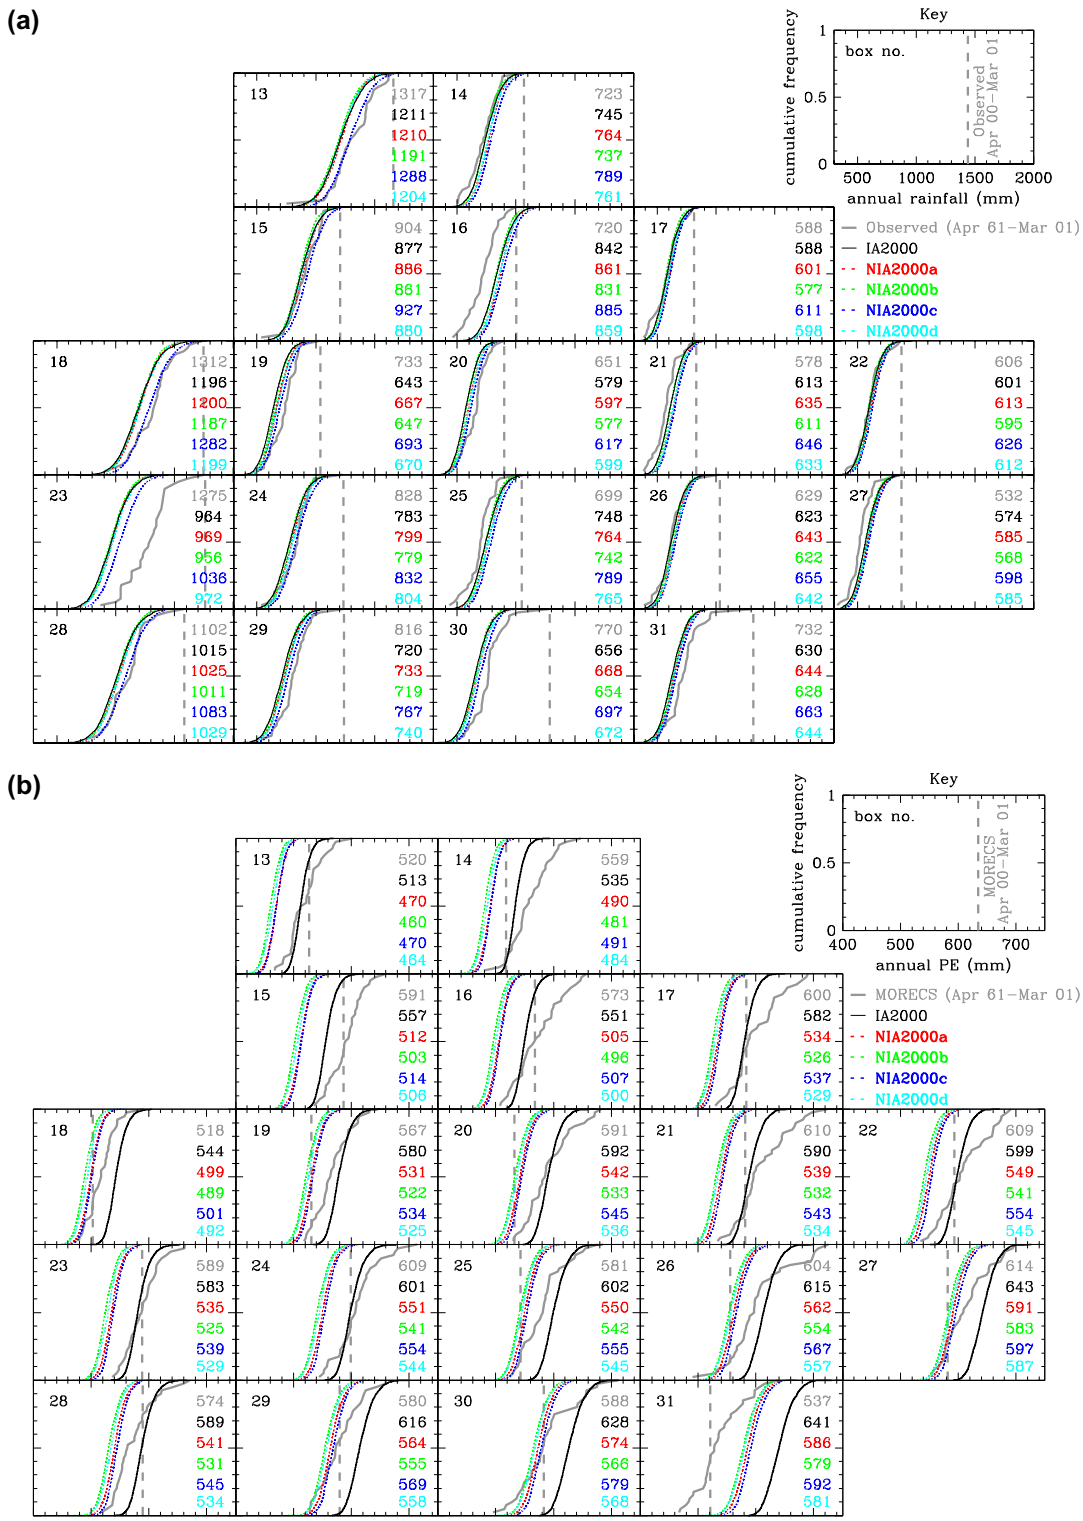

Fig. 3. Cumulative distributions of annual total: (a) rainfall and (b) PE from the IA2000 ensemble (solid black), each of the four NIA2000 ensembles (dotted) and from (a) observed rainfall and (b) MORECS PE for April 1961–March 2001 (solid grey), for each of the 19 'land' climate model grid boxes over England and Wales (numbered in the top-left corner of each box). The numbers down the right of each graph give the mean value from each distribution, in the order of the key (top-right). The dashed vertical grey lines show the (a) observed rainfall and (b) MORECS PE for April 2000–March 2001.

d). In winter, the daily rainfall extremes in NIA2000 sub-ensembles a, b and d are similarly lower than IA2000 (by an average of between −7% and −10%), but NIA2000c has extremes similar to or higher than IA2000 (by an average of about +2%; not shown).

A comparison of the cdfs of annual-average daily mean T from each ensemble within each of the 19 'land' climate model boxes over England and Wales shows that there is a large difference between the IA2000 ensemble and each of the NIA2000 ensembles (not shown), with each of the latter shifted leftwards of the former, indicating lower temperatures in a non-industrial climate. The same is true for daily mean T in each season (not shown).

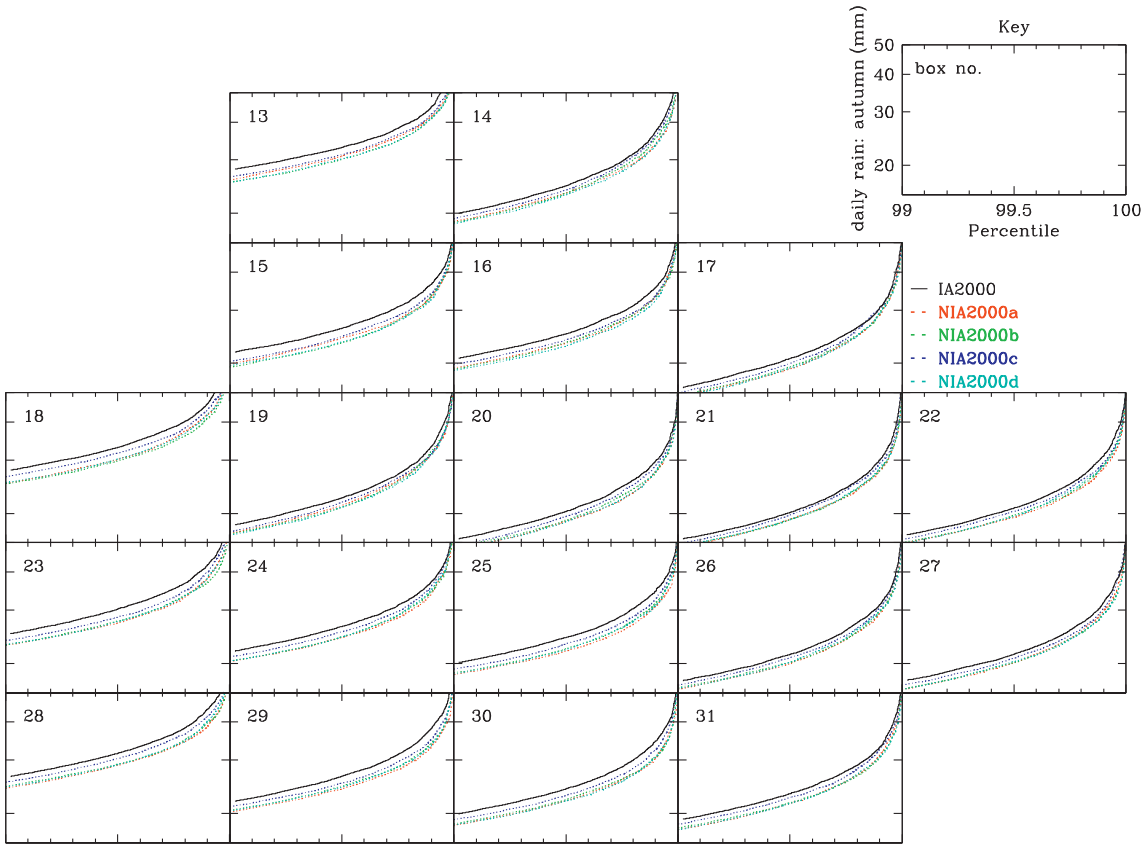

Fig. 4. Distributions of autumn daily rainfall extremes from the IA2000 ensemble (solid black) and each of the four NIA2000 ensembles (dotted) for each of the 19 'land' climate model grid boxes over England and Wales (numbered in the top-left corner of each box).

The cdf from the IA2000 ensemble is also compared to the distribution of annual-average daily mean *T* from observed data for April 1961–March 2001 (available on a 5 km × 5 km grid, but interpolated onto the climate model grid using area weighting). As for rainfall, this comparison is generally quite good, particularly for non-coastal grid boxes over England and Wales (not shown). The same is generally true for autumn and winter daily mean *T* (not shown). For summer daily mean *T*, the cdf from the IA2000 ensemble is generally shifted rightwards from the observed cdf (not shown), suggesting a possible warm bias from the climate model in the summer months (although see the discussion at the start of this section regarding the limitations of this sort of comparison). There is also a possible warm bias in spring, although less pronounced than in summer (not shown).

As for the comparison for daily mean *T*, a comparison of the cdfs of annual total PE from each ensemble within each of the 19 'land' climate model boxes over England and Wales shows that there is a large difference between the IA2000 ensemble and each of the NIA2000 ensembles (Fig. 3b), with each of the latter shifted leftwards of the former, indicating lower PE in a non-industrial climate. The same is true for PE in each season (not shown).

Also shown for each climate model grid box in Fig. 3b is the cdf of annual total PE data for April 1961–March 2001 from MORECS (available on a 40 km × 40 km grid but interpolated onto the climate model grid using area weighting). Comparison is made with MORECS data as they are derived from observed meteorological data (using a Penman–Monteith formulation) and provide the PE data used for hydrological model calibration. The comparison of the cdf of annual total PE from the IA2000 ensemble with that from MORECS is not as good as the comparison for rainfall or *T* data. The IA2000 ensemble data generally have a narrower range of values (steeper slope to their cdf) than the MORECS data in each climate model grid box, and whilst the mean values are similar for some grid boxes they are quite different for other grid boxes, with the shift not having a consistent direction for all grid boxes. The same is true for autumn PE (not shown). For PE in the remaining seasons, whilst the range of values from the IA2000 ensemble is still narrower than that from the MORECS data, the cdf is generally shifted leftwards for winter and spring PE and rightwards for summer PE, indicating a possible low-bias in winter/spring PE and a high-bias in summer PE from the climate model. The latter is consistent with the possible summer warm bias found from the comparison of daily mean *T* data.

However, the fact that the IA2000 PE data are derived, out of necessity, using a *T*-based formulation (see Section 3.2.2) rather than a Penman–Monteith formulation, like that used within MORECS, makes direct comparison with MORECS more difficult. Kay and Davies (2008) compared MORECS PE with *T*-based PE derived from MORECS' *T* data (for 1961–1990), for regions over the North and South of Britain. They showed that the *T*-based PE was lower than MORECS in winter and spring but higher than MORECS in summer (and similar in autumn), but that annual means were similar (with *T*-based PE just slightly less than MORECS PE). It is thus possible that some of the 'bias' in seasonal PE from the IA2000 ensemble compared to MORECS is actually due to the use of the *T*-based formulation of PE. However, this does not help to explain the differences seen even for annual PE.

Despite the differences seen in the PE data, the ensembles of climate model data are used as-is to drive the hydrological models, since it is rainfall that is the key driver of flooding and the comparison of the rainfall data from the IA2000 ensemble with observed data is relatively good.

3.2.4. Ensemble data for North–East and South–East England

Differences outlined above between the IA2000 and NIA2000 ensembles are evident in the mean annual and seasonal rainfall, temperature and PE values for the nine grid boxes used in the hydrological modelling. Differences in these means are given in Table 4, where values are averages for three grid squares for the North–East region and six grid squares for the South–East region (see Table 3).

The rainfall values in Table 4 show that NIA2000c is the only sub-ensemble which is wetter in all seasons than IA2000, contributing to the noticeable difference between this sub-ensemble and the other three in terms of annual total of rainfall (Fig. 3a). The contrast is greater for the North–East region than the South–East. All NIA2000 sub-ensembles are wetter than IA2000 in the summer and drier (apart from NIA2000c) in the winter. NIA2000b is drier than IA2000 in all seasons apart from summer, and is the only ensemble which is drier than IA2000 annually. In spring and autumn, differences between the NIA2000 sub-ensembles and the IA2000 ensemble are small.

All NIA2000 sub-ensembles are colder than the IA2000 ensemble, in all seasons, though autumn is generally the season with the greatest difference, especially in the South–East. The lower temperatures also mean lower PE in the NIA2000 sub-ensembles compared to the IA2000 ensemble, in all seasons, although with the greatest differences in summer and smallest in winter. NIA2000b has the greatest difference in PE annually, and NIA2000c the least difference.

Although the annual differences between the NIA2000 and IA2000 ensembles are not great, the differences in seasonality of rainfall, with drier summers and wetter winters (apart from NIA2000c) in the industrial climate, combined with the higher PE, can impact on flood potential. Differences between the NIA2000 ensembles may also cause differences in high flow response.

3.3. Modelling methodology

For each catchment, each year of climate model-derived data is set in the right place within the observed data series (that is, replacing the year from April 2000), and used to drive the hydrological model (including the snowmelt module). The PDM runs, for the six smaller catchments, use the observed data from the start year given in Table 1, whilst the CLASSIC runs, for the two largest catchments, use the observed data from January 1999. From each modelled flow series, the peak flow occurring within the 3 months (90 climate-model days) October–December 2000 is extracted. This is done for running-mean flows over 10- and 30-day durations, as well as for the daily flows output by the hydrological models, as the autumn 2000 floods were generally more unusual over longer durations. For the *n*-day running-mean flows, the peak extraction looks at any *n*-day period starting within October–December 2000. The period October–December was chosen as no significant flooding occurred in England in Autumn 2000 before October, whilst in some catchments the flooding did not reach its peak until December of that year (including modelled catchments 39001 and 42012; see Fig. 2).

The need to keep the results separated by ΔSST scaling factor as well as sub-ensemble results in 41 sets of peaks at each duration (4 NIA2000 sub-ensembles × 10 ΔSST scaling factors + 1 IA2000 ensemble). The peaks corresponding to each set can then be combined into flood frequency plots (showing peak exceedance against return period), using the Gringorten plotting position (Gringorten, 1963). That is, the *i*th of *N* peaks (when ordered from smallest to largest) is plotted at position $p = (i - 0.44)/(N + 0.12)$, with the return period given by $1/(1 - p)$. The results can be presented as a set of 12 plots for each catchment (4 NIA2000 sub-ensembles × 3 durations), with 11 sets of data on each plot (1 from the IA2000 sub-ensemble, plus 10 from the different ΔSST scaling factors for the appropriate NIA2000 sub-ensemble).

An example set of flood frequency plots is shown in Fig. 5, for catchment 27007. An inspection of these plots shows that higher peaks are generated with the NIA2000c sub-ensemble than the IA2000 ensemble, particularly for longer durations, which is compatible with the higher rainfall and lower PE in all seasons for this sub-ensemble. The other NIA2000 sub-ensembles generally show the opposite direction of change, with mostly higher peaks in the IA2000. The extremity of the flows in autumn 2000, particularly for the 10-day duration, is evident from the position of the dotted horizontal lines relative to the dashed lines in Fig. 5.

For each catchment, additional sets of runs are performed to investigate the possibly reduced potential of large snowmelt-induced floods in an industrial as opposed to non-industrial climate, due to higher winter temperatures. In these runs, the ensemble data are placed in the correct position within the observed data (as before) and used to drive the hydrological model first with and then without the snowmelt module. From each modelled flow series, the peak flow occurring within the 6 months (180 climate-model days) October 2000–March 2001 is extracted. This longer period is chosen so that it covers the coldest, winter period, and continues up to the end of the 1-year climate model simulations, during the spring warming, (see Section 3.2.1). As before, the extraction is done for running-mean flows over 10- and 30-day durations, as well as for the daily flows output by the hydrological models.

3.4. Fraction of attributable risk

The sets of simulated peaks can then be analysed in terms of their probability of exceeding a given threshold, with a comparison between the exceedance probabilities for the IA2000 and NIA2000 sub-ensembles giving an indication of change in risk due to anthropogenic emissions (assuming a fixed relationship between hazard and damage). This change can be denoted as the fraction of attributable risk (*FAR*);

$$FAR = 1 - NIE/NE \quad (1)$$

Table 4
Differences in mean seasonal rainfall, temperature and PE in the IA2000 ensemble compared to each NIA2000 ensemble (a–d), averaged over grid squares for two regions of England.

	North–East England				South–East England			
	a	b	c	d	a	b	c	d
Rainfall (mm)								
Annual	−13	12	−55	−9	−17	2	−39	−19
Spring	−1	5	−6	−3	0	4	−2	−3
Summer	−28	−10	−22	−22	−29	−15	−22	−25
Autumn	6	6	0	5	1	0	−3	0
Winter	9	11	−28	11	11	13	−13	10
Temperature (°C)								
Annual	1.4	1.6	1.4	1.6	1.5	1.7	1.4	1.7
Spring	1.3	1.4	1.3	1.3	1.2	1.4	1.2	1.2
Summer	1.3	1.7	1.1	1.5	1.6	1.9	1.4	1.8
Autumn	1.5	1.8	1.6	1.8	1.7	1.9	1.6	1.9
Winter	1.6	1.6	1.8	1.8	1.6	1.5	1.3	1.7
PE (mm)								
Annual	45	54	44	51	52	60	48	58
Spring	14	15	14	13	14	15	13	13
Summer	17	24	16	21	22	27	19	26
Autumn	9	11	9	11	11	13	11	13
Winter	5	5	5	5	5	5	4	6

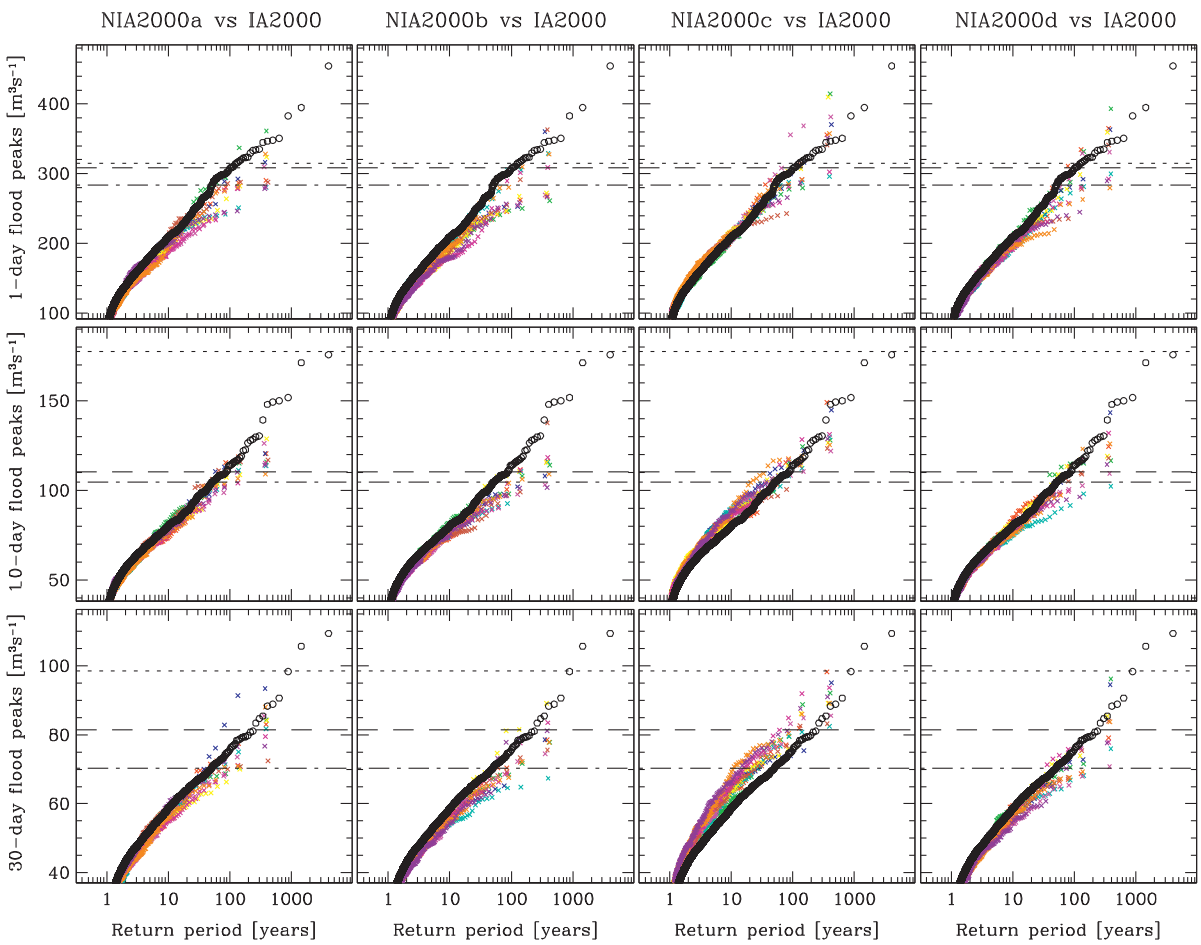

Fig. 5. Flood frequency plots for catchment 27007, comparing the results from the IA2000 ensemble (black circles) to those from NIA2000 ensembles a–d (columns) for each ΔSST scaling factor (crosses), at three durations (rows). Also shown, as horizontal lines on each graph, are peak flows (i) simulated in October–December 2000 using observed input data (dotted line), (ii) simulated prior to April 2000 using observed input data (dashed line) and (iii) simulated in October–December 2000 using the IA2000 ensemble and corresponding to a 50-year return period (dot-dashed line).

where *IE* is the fraction of the IA2000 runs exceeding the given threshold, and *NIE* is the fraction of the NIA2000 runs exceeding the given threshold (Allen, 2003). A *FAR* > 0 indicates an increased chance of such flooding in an industrial as opposed to non-industrial climate, with a value closer to 1 being a very much increased chance, while a *FAR* < 0 indicates a decreased chance of such flooding in an industrial as opposed to non-industrial climate. The *NIE* and *FAR* are calculated separately for each NIA2000 sub-ensemble/ΔSST scaling factor subset. In addition, a bootstrapping technique is applied (whereby 100 new sets of IA2000 and NIA2000 peaks are produced by random resampling, with replacement, from within each of the original sets) to estimate a distribution of *FAR* in each case. The distributions for the 10 ΔSST scaling factor subsets within a sub-ensemble can then be amalgamated, to produce a *FAR* distribution for each sub-ensemble. The latter can then be further amalgamated, to produce an overall *FAR* distribution.

One issue with this method of estimating *FAR* is the choice of threshold to apply. If the chosen threshold is too high, then it is possible for one (or both) of *IE* and *NIE* to be zero, resulting in a *FAR* of, respectively, $-\infty$ or 1. A seemingly natural choice of threshold, given the aim of studying the Autumn 2000 floods, would be the observed Autumn 2000 flow peak. A better alternative would be the Autumn 2000 flow peak simulated by the hydrological model driven with observed input data, as this reduces possible problems with inaccurately measured or missing observed flow data and with potential model bias. However, this threshold is too high for some of the catchments modelled here (see example in Fig. 5), resulting in difficulties in estimating *FAR*. Using the highest flow simulated with observed input data prior to Autumn 2000 does not always improve things significantly (Fig. 5). The difficulty with using a threshold derived from flows simulated with observed data could stem from two sources: the relatively small ensemble sizes in comparison to the rarity of the Autumn 2000 floods (especially for the NIA2000 sub-ensembles when they have to be further sub-divided by ΔSST scaling factor, see Table 2), and possible bias in the climate model data resulting in simulated flow peaks being too low compared to observed flow peaks. It is difficult to say which source is most important, although it is assumed that any bias in the IA2000 climate applies equally to each of the NIA2000 climates.

Instead then, the threshold is taken from the flow peaks simulated under the IA2000 ensemble, by selecting the simulated peak which is closest to having a 50-year return period. This removes the possibility of a *FAR* of $-\infty$ as, by definition of the Gringorten plotting position (see Section 3.3), 45 of the 2271 IA2000 runs will have a peak above this threshold (i.e. the value of *IE* is fixed at 45/2271 = 0.019815, so cannot be zero). However, a *FAR* of 1 is still possible since some of the sets of NIA2000 peaks could lie completely below this threshold (i.e. *NIE* = 0). In the pseudo-ensembles, generated by resampling, the value of *IE* is not fixed but is unlikely to be zero.

Choosing the threshold in this way provides a threshold that is more consistent, both between different catchments and between

different durations. However, it is possible that alternative choices of the threshold would result in a different estimation of *FAR*, especially where the industrial and non-industrial flood frequency curves have a different shape. The threshold corresponding to the IA2000 50-year return period was chosen so as to still be quite extreme, but so that most sub-ensembles had a reasonable number of runs exceeding the threshold. The use of alternative thresholds is discussed at the end of the paper.

3.5. Antecedent conditions

In order to assess the potential influence of antecedent conditions in the period prior to April 2000 on flows in Autumn 2000, further sets of IA2000 runs are performed (including the snowmelt module) for the four catchments in South–East England (as these have the greatest range of permeability; Table 1). In these runs, the climate model-derived data from the IA2000 ensemble are placed in different positions within the observed data series. That is, they are used to replace the observed data for the year from April 1962, then for the year from April 1963, and so on up to the year from April 2001. In each case, flood peaks are extracted as before (from within the 3 months October–December of the replaced year), and a count is made of the number of runs exceeding a fixed threshold. The chosen threshold is that derived as described in Section 3.4, by taking the simulated peak closest to having a 50-year return period from the set of simulations where the IA2000 ensemble replaces the 'correct' year (from April 2000).

By definition, the number of exceedances of the threshold when the IA2000 ensemble replaces the observed data for the year from April 2000 is 45. The question is how much the exceedance count varies when the IA2000 ensemble replaces alternative years of the observed data, and whether the value of 45 is about average or not. The overall level of variation in the exceedance count for a catchment indicates the degree to which the rainfall, temperature and evaporation in the months prior to April affects the flows in the following October–December for that catchment. An average exceedance count close to 45 indicates that the conditions prior to April 2000 were not that influential, whereas an average further from 45 indicates a greater influence of the period prior to April 2000.

4. Results

4.1. Comparison between FAR distributions for October–December flow peaks

The *FAR* distribution boxplots in Fig. 6 show that there is a wide range of uncertainty in the *FAR* values for each of the eight study catchments. The *FAR* distribution has a relatively wide range for a given NIA2000 sub-ensemble (due to the 10 ΔSST scaling factors and to the resampling methodology applied; Section 3.4), but there is even greater difference between NIA2000 sub-ensembles. In general, for each catchment and duration, sub-ensemble NIA2000b gives the highest median value of *FAR* (especially at shorter durations). Sub-ensemble NIA2000c always gives the lowest median value of *FAR*, in many cases by quite a wide margin over the other three NIA2000 sub-ensembles. This is especially the case at longer durations, where the NIA2000c median *FAR* is often below the 25th percentile *FAR* for each of the other three sub-ensembles. This large difference is due to the NIA2000c sub-ensemble having higher rainfall in all seasons than the IA2000 ensemble and the other three NIA2000 sub-ensembles (Section 3.2.4).

The large differences between the *FAR* distributions from the four NIA2000 sub-ensembles mean that the overall *FAR* distributions (which combine the distributions from the corresponding four NIA2000 sub-ensembles, assuming equal likelihood) are quite wide. Their 5th–95th percentile range extends to either side of *FAR* = 0 in all cases, and even their 25th–75th percentile range extends to either side of *FAR* = 0 in the majority of cases (except for the more northerly catchments at the 1-day duration). But the median of the overall *FAR* distribution is positive, at all durations, for all catchments except 42012 [recall that *FAR* > 0 indicates an increased chance of such flooding in an industrial as opposed to non-industrial climate, whereas *FAR* < 0 indicates a decreased chance].

There are relatively small differences between the *FAR* distributions for any of the catchments apart from 42012, which has negative median *FAR* values for all NIA2000 ensembles (except for NIA2000b at the 1- and 10-day durations). The differing response of 42012 is due to the fact that flow in this catchment is by far the most groundwater-dominated (highest BHP; Table 1) of the eight example catchments. For permeable catchments the higher rainfall and lower PE during the summer in the non-industrial climates (Table 4) mean that groundwater stores are not depleted as much as in the industrial climate, and thus these catchments are potentially more responsive to autumn and winter rainfall in a non-industrial climate than in an industrial climate. The results demonstrate the long hydrological 'memory' of such catchments, so that differences and changes over months and seasons have more effect than those from single extreme events.

In contrast, higher summer rainfall has less impact on the rainfall–runoff response in autumn and winter for non-permeable catchments, so values of *FAR* for these catchments more directly reflect differences in the same month/season. At the 1-day duration, *FAR* values are generally higher for the North–East catchments than for those in South–East England, which probably reflects a difference between the two regions for changes in autumn rainfall (Table 4) combined with faster rainfall–runoff response times. For the North–East region all NIA2000 sub-ensembles (apart from c) have lower autumn rainfall than IA2000 (see Table 4), whereas for the South–East region only NIA2000a is lower for all grid boxes with NIA2000b for some boxes (that is, there is variation in catchment change depending on location within the region). Differences in 1-day flows are (except for permeable catchments) also dependent on differences in daily rainfall intensity. As shown in Fig. 4, intensities of daily rainfall in all NIA2000 sub-ensembles are lower than those in IA2000; lower extreme daily rainfall in NIA2000 ensembles would contribute to higher values of *FAR* for 1-day duration.

The pattern of change in *FAR* with duration differs between catchments; for some catchments the overall median *FAR* increases with duration (e.g. 39069), for some catchments the overall median *FAR* decreases as duration increases (e.g. 27007, 27021, 27049), and for other catchments the median *FAR* is largest/smallest at the middle (10-day) duration (e.g. 39105, 27009). Similar behaviour is seen when the results for individual NIA2000 sub-ensembles are considered. The variable pattern of *FAR* with duration is the combination of catchment response time, seasonal changes in rainfall and balance with evaporation and temporal distribution of rainfall. For instance, it would generally be expected that smaller, faster responding catchments (with low BHP, high SAAR and steep slopes) would show more sensitivity to changes in rainfall at shorter durations, as with catchments 27007, 27021 and 27049. Catchment area is also a factor in the relationship between duration and *FAR*, as larger catchments, particularly those with moderate to high BHP, have the potential for mean flows to be affected over longer durations.

4.2. Comparison between FAR distributions for October–March flow peaks, with and without the snowmelt module

The *FAR* distribution boxplots in Fig. 7 show that *FAR* with the snowmelt module is generally lower than *FAR* without the

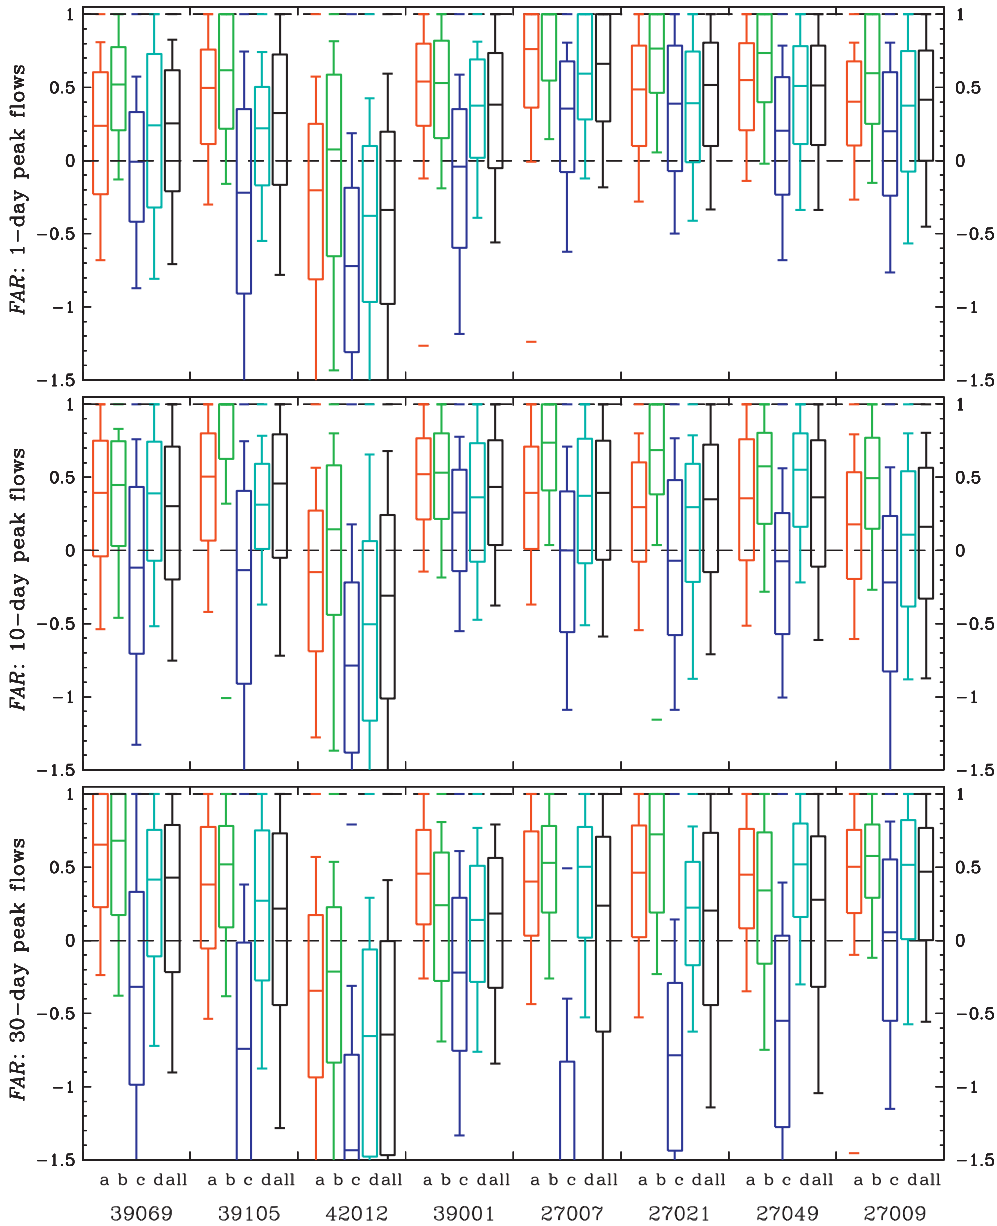

Fig. 6. Boxplots showing the distributions of *FAR* for each catchment, for peak flows in October–December 2000 at each of three durations (1, 10 and 30 days). The distribution is shown for the four NIA2000 ensembles (a–d) separately and aggregated (all). The box denotes the 25th and 75th percentiles, with the median (50th percentile). The whiskers denote the 10th and 90th percentiles, with additional markers outside the whiskers indicating the minimum and maximum values (if in the plotted range −1.5–1). In each case the hydrological modelling includes the snowmelt module.

snowmelt module (for flood peaks in the 6-month period October–March). This means that the reduced likelihood of snowmelt-induced floods in the warmer temperatures of the industrial climate moderates the increased flood chance due to other sources of flooding. That is, when the snowmelt module is not used the NIA2000 sub-ensembles cannot generate the snowmelt-induced flood peaks to go with their colder climates, and so the NIA2000 results generally sit below the IA2000 results, meaning greater positive *FAR* values. When the snowmelt module is used, the lower temperatures of the NIA2000 climates lead to an increased number of snowmelt-induced flood peaks than in the higher temperatures of the IA2000 climate, thus bringing the NIA2000 results up towards (or even above) the IA2000 results, meaning lower (or negative) *FAR* values. This effect is more pronounced for the more northerly catchments, and at shorter durations for the more southerly catchments. It is still less pronounced for catchment 42012.

4.3. Overall comparison between FAR distributions

A comparison between the coloured boxes in Fig. 7 (for flood peaks in the 6-month period October–March) with those in Fig. 6 (for flood peaks in the 3-month period October–December) shows similar relative behaviour of the NIA2000 sub-ensembles and catchments (described in Section 4.1) for the two time periods. For instance, the NIA2000b sub-ensemble generally gives the highest *FAR* values whilst the NIA2000c sub-ensemble always gives the lowest *FAR* values. However, the *FAR* ranges using the October–March period are sometimes wider than their October–December equivalents, and often positioned with a lower median value. Thus in almost all cases the 25th–75th percentile range includes *FAR* = 0 (the exceptions being catchments 27049 and 27007 at the 1-day duration), and in over half of cases the median of the overall *FAR* is now negative rather than positive. This is likely to be due to

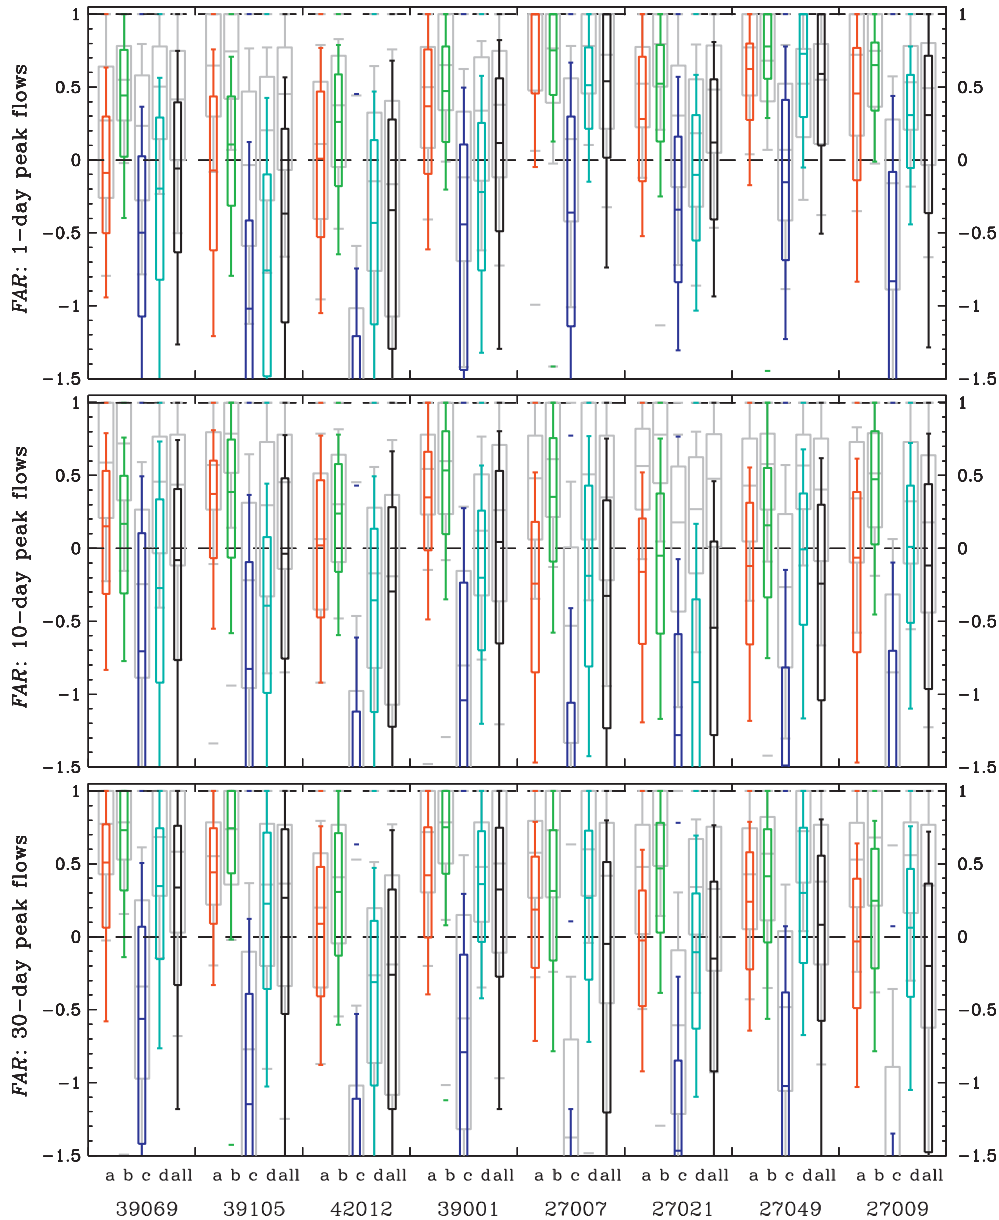

Fig. 7. As Fig. 6 but for flood peaks in October 2000–March 2001, where the hydrological model has been run both with (black/coloured boxes/lines) and without (wider light grey boxes/lines) the snowmelt module. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

the enhanced effect of snowmelt peaks (and their reduction in a warmer climate) over the extended period (October–March cf. October–December), which covers the winter and start of spring rather than ending in early/mid winter. However, at the 30-day duration catchments 39001 and 42012 have higher median *FAR* values (overall, and for all NIA2000 ensembles apart from c) when using the 6-month rather than 3-month analysis period, reflecting the effect of their relatively long response time.

Overall median values of *FAR*, for the four NIA2000 ensembles combined, are given in Table 5 for each catchment and for the 3- and 6-month analysis periods. Negative values of *FAR* are shown in italics, with thresholds of +0.33 and −0.5 used to distinguish notably high or low values respectively (in bold). Particular features of Table 5 are:

- positive *FAR* for both regions (excepting catchment 42012) for the 3-month period;
- negative *FAR* for the 6-month period apart from North–East England at the 1-day duration and South–East England at the 30-day duration;
- high *FAR* for North–East England at the 1-day duration, for both the 3- and 6-month periods;
- all negative *FAR* for catchment 42012.

In general, the industrial and non-industrial climates seem to be less differentiated when considered in terms of flows within the whole autumn/winter period than they are for autumn flows alone, shown by the mean values over all eight catchments.

4.4. The effect of antecedent conditions

Fig. 8 shows the bearing that catchment permeability has on the influence of antecedent conditions on simulated flows during the climate-model derived year. For the four catchments in

Table 5
Overall median values of FAR for each catchment (for NIA2000 sub-ensembles a–d combined) and the mean FARs over the four catchments in each region and over all eight catchments, at each duration for each analysis period. Negative FAR are shown in italics, with FAR > 0.33 and FAR < −0.5 in bold (corresponding respectively to a greater than 50% increase in risk from non-industrial to industrial climates or from industrial to non-industrial climates).

Catchment	3-month period (October–December)			6-month period (October–March)		
	1-day	10-day	30-day	1-day	10-day	30-day
North–East England						
27007	**0.66**	**0.39**	0.24	**0.54**	*−0.32*	*−0.05*
27021	**0.52**	**0.35**	0.20	0.12	***−0.54***	*−0.15*
27049	**0.51**	**0.36**	0.28	**0.59**	*−0.24*	0.08
27009	**0.41**	0.16	**0.47**	0.31	*−0.12*	*−0.20*
Mean	**0.53**	0.32	0.30	**0.39**	*−0.31*	*−0.08*
South–East England						
39069	0.25	0.30	**0.43**	*−0.06*	*−0.08*	**0.34**
39105	0.32	**0.46**	0.22	*−0.37*	*−0.04*	0.27
42012	*−0.34*	*−0.31*	***−0.64***	*−0.35*	*−0.29*	*−0.26*
39001	**0.38**	**0.43**	0.18	0.12	0.04	0.32
Mean	0.16	0.22	0.05	*−0.16*	*−0.09*	0.17
Mean (all)	**0.34**	0.27	0.17	0.11	*−0.20*	0.04

South–East England, the count of the number of IA2000 runs with (October–December) peaks exceeding a fixed threshold is shown, where the IA2000 runs have been used to replace each year of observed data in turn from 1962 to 2001. The threshold is the 50-year return period flow from the IA2000 ensemble placed in the 'correct' position; replacing the observed data for the year from April 2000. The greatest variability in exceedance count occurs for catchment 42012, followed by catchment 39001, with much lower variability for the other two catchments. The pattern of variability follows the order of bedrock high permeability (Table 1) where catchment 42012 has the highest value of the four catchments (100), followed by 39001 (43), 39105 (15) and 39069 (0). These results illustrate the impact of catchment properties and antecedent conditions on flood response. Fig. 8 also highlights the 'wetness' of the autumn/winter of 2000/01 as, at each duration, the highest exceedance count for catchments 42012 and 39001 is generally when the IA2000 runs have been used to replace the year from April 2001.

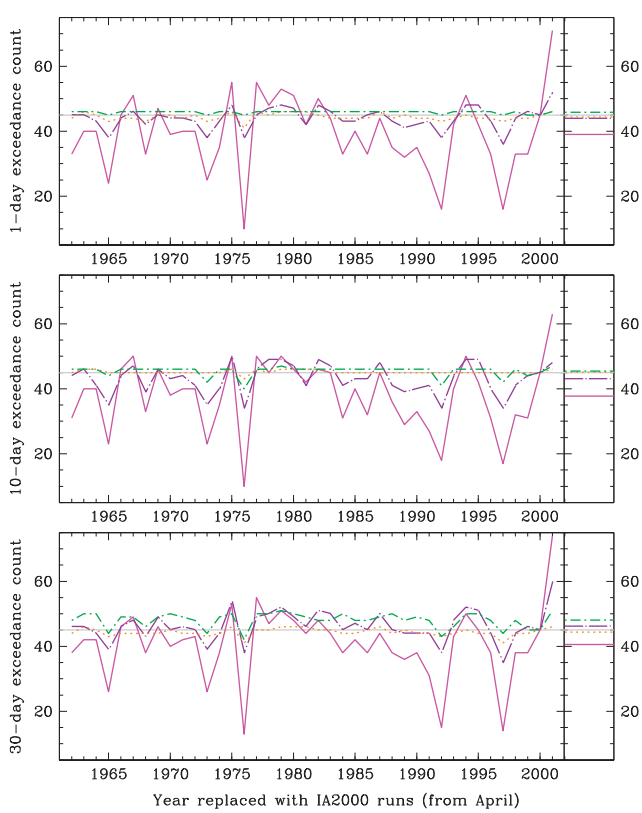

Fig. 8. The number of IA2000 runs exceeding a fixed threshold, where the runs are used to replace each year of observed data in turn from 1962 to 2001, at three durations and for the four catchments in South–East England: 42012 – solid (magenta), 39001 – dotted/long-dashed (purple), 39069 – dotted/short-dashed (dark green), 39105 – dotted (orange). The solid grey horizontal line shows the standard exceedance count (45), where the IA2000 runs replace the correct year of observed data (from April 2000). The horizontal lines on the right-hand side of each plot show the mean exceedance count over 1962–2001 for each catchment. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.).

By definition, the threshold exceedance count when the IA2000 ensemble replaces the observed data for the year from April 2000 is 45 (see Section 3.4). Catchment 42012 is the only catchment with a mean count relatively far from, and lower than, this standard count, indicating a large influence of the catchment's weather in the period prior to April 2000 on its flows from Autumn 2000. The question is then whether the chance of the antecedent conditions that contributed to the Autumn/Winter 2000 flooding in catchment 42012 (and in catchments like it) has been changed due to anthropogenic emissions. Any such changes will not be imposed via the use of the 1-year IA2000 and NIA2000 ensemble runs, and so the FAR distribution calculated here may not be a true representation of the effect of anthropogenic emissions on flooding in permeable catchments.

5. Summary and conclusions

This paper has used hydrological modelling, together with ensembles of climate model data representing the year April 2000–March 2001 in industrial (IA2000) and non-industrial (NIA2000) climates, to assess the impact of greenhouse gas emissions on the chance of occurrence of floods like those in England in Autumn 2000. Four sub-ensembles of the non-industrial climate were used (a–d, representing uncertainty in the spatial pattern of non-industrial sea-surface temperatures, SSTs) each of which was further sub-divided into 10 subsets (according to the ΔSST scaling factor applied, representing uncertainty in the amplitude of the ΔSST spatial pattern). Eight catchments were studied, in two regions of England (South–East and North–East). The climate model data were used to drive the hydrological model for each catchment, and the peak flow occurring within a given analysis period was extracted from each modelled flow series, for flows at three durations (1, 10 and 30 days). The analysis periods were the 3 months October–December 2000, to span the main period of flooding that occurred in Autumn 2000, and the 6 months October 2000–March 2001, to

cover the following winter/early spring. The changed possibility of flood occurrence due to industrial emissions was then assessed through calculation of the fraction of attributable risk (*FAR*), which compares the proportion of runs in the industrial and non-industrial ensembles having peak flows exceeding a given threshold. The threshold was set to the 50-year return period IA2000 flow for each catchment, rather than being dependent on the actual Autumn 2000 flows, in order to avoid problems with estimation of *FAR*. Bootstrapping was used to help estimate the range of uncertainty in *FAR* in each case.

Combining results for all four NIA2000 sub-ensembles, positive median values of *FAR* for the 3-month analysis period indicate that, for all catchments (apart from 42012) at all durations, there is an increased chance of high flows in an industrial climate. For the 6-month analysis period, it is mostly only the catchments in the North–East region at the 1-day duration and the Thames (39001) and two of its sub-catchments at the 30-day duration which indicate this increased chance. In general, the industrial and non-industrial climates seem to be less differentiated when considered in terms of flows within the whole autumn/winter period than they are for autumn flows alone. For permeable catchment 42012, generally negative median values of *FAR* indicate a decreased chance of high flows in an industrial compared to non-industrial climate. But an analysis of the effect of antecedent conditions showed that, for such catchments, a longer period of climate data than 1-year is probably required to obtain representative values of *FAR*.

Apart from catchment 42012, there are generally smaller differences between results for different catchments than there are between the results for the four NIA2000 sub-ensembles for a catchment. The highest median values of *FAR* for each catchment and duration are generally obtained with the NIA2000b sub-ensemble, followed by NIA2000a and NIA2000d, with easily the lowest values of *FAR* from the NIA2000c sub-ensemble. This is the same ordering as was found by Pall et al. (2011), from their empirical modelling of total daily runoff over England and Wales. Each NIA2000 ensemble has higher summer rainfall and lower PE than the IA2000 ensemble, and so has the *potential* to generate higher autumn flows than IA2000. This is indeed the case for NIA2000c, which thus results in lower values of *FAR*. However, autumn rainfall in the other NIA2000 ensembles is lower so this potential is not realised, resulting in autumn flows that are higher in IA2000 and so higher values of *FAR* when using NIA2000 sub-ensembles a, b and d. If one consequence of industrialisation is increasing seasonality of rainfall, with a shift in rainfall pattern to higher winter rainfall and lower summer rainfall, then climatic conditions during the autumn become more critical in determining whether rainfall generates extreme flows at this time of year. It is difficult to say whether any one NIA2000 sub-ensemble is more or less likely than the others, thus each has been given equal weight in the overall FAR distributions presented here.

However, definitive conclusions are difficult as there are wide bands of uncertainty in *FAR*, with the 5th–95th percentile range spanning $FAR = 0$ in the vast majority of cases, for both analysis periods, and even the 25th–75th percentile range spanning $FAR = 0$ in many cases, especially for the 6-month analysis period. This uncertainty comes from the selection of multiple patterns and amplitudes of the non-industrial SSTs as well as from the bootstrapping (which should pick up the presence of outliers).

All of the reported results were based on the threshold from the 50-year return period IA2000 flow. In general, use of a lower threshold than this results in median values of *FAR* closer to zero, with a narrower uncertainty band. This indicates less difference between the IA2000 and NIA2000 ensembles for less extreme flood events. It is difficult to extend this to higher thresholds though, due to the relatively small ensemble sizes, particularly for the NIA2000 sub-ensembles as they have to be further sub-divided according to the ΔSST scaling factor applied.

Running the rainfall–runoff model with and without the snowmelt module demonstrated another aspect of the effect of differing rainfall and temperature on the generation of peak flows. The differences impact on flows in opposite directions, illustrating the need to ensure that the full picture has been taken into account when drawing conclusions about impacts of change. Not all catchments will experience the same impacts, due to the complex interaction of factors operating at a variety of temporal and spatial scales.

The catchment-specific results presented here (at the 1-day duration for the 3-month analysis period) generally have a lower median *FAR*, and wider uncertainty range, than the results presented by Pall et al. (2011) for autumn England and Wales total daily runoff. However, there are a number of methodological differences which could have led to this apparent discrepancy:

- Continuous simulation hydrological models vs. empirical linear transfer function models.
- Catchment modelling (with catchment-specific input data) vs. regional modelling and area-averaging (with England and Wales average rainfall inputs).
- Use of effective rainfall derived from rainfall, soil-moisture accounting and PE (estimated from temperature data) vs. use of a fixed fraction of rainfall as effective rainfall.
- Inclusion of a snowmelt module vs. purely rainfall–runoff modelling.
- Analysis period October–December vs. September–November.
- Threshold defined by 50-year return period IA2000 flow vs. threshold defined by highest Autumn 2000 flow simulated using re-analysis rainfall data (corresponding to approximately the 9-year return period IA2000 flow).

The latter point could well account for much of the difference in the width of the uncertainty bands, since uncertainty is reduced for lower thresholds. A number of the other points could account for the difference in median *FAR*, particularly the catchment vs. regional modelling and the formulation used for catchment losses (PE).

Another factor is the relatively low resolution of the climate model relative to the size of most of the catchments being modelled. Ideally a regional climate model would be nested within the global climate model, to bring the climate model resolution closer to the scales on which precipitation and hydrology operate. However, the use of rainfall–runoff models calibrated for specific catchments with differing physical catchment properties shows the hydrological consequences of seasonal differences in precipitation and temperature on generation of extreme flows. The different analyses show that a number of factors, from a hydrological perspective, require careful consideration in interpretation of the results, in particular, catchment response time, the impact of preceding conditions on subsequent flows, the role of snowmelt in generation of high flows and the threshold discharge used to calculate the difference in risk between a non-industrial and industrial climate. Changes in annual rainfall between the climatic regimes may be small but the breakdown into seasonal changes, combined with changes in daily rainfall intensity, has greater consequence for impacts on flow. Changes in the water balance (balance between rainfall and evaporation) during the summer and autumn are of particular importance in looking at influences on flows during the autumn and winter when floods are most likely to occur.

Acknowledgements

We thank all of the volunteers who ran the climate model simulations using the BOINC open source computing platform, and the

climateprediction.net team for their technical support. P.P. was supported by a NERC CASE studentship with Risk Management Solutions Ltd., by WWF International, and by NCCR climate. D.A.S. received partial support from the UK Department for Environment, Food, and Rural Affairs and additional support from the US Department of Energy's Office of Science, Office of Biological and Environmental Research, and the US National Oceanic and Atmospheric Administration's Climate Program Office, all via the International ad hoc Detection and Attribution Group.

References

Allen, M., 2003. Liability for climate change. Nature 421, 891–892.

Bell, V.A., Moore, R.J., 1999. An elevation-dependent snowmelt model for upland Britain. Hydrological Processes 13, 1887–1903.

Blackburn, M., Hoskins, B.J. 2001. Atmospheric Variability and Extreme Autumn Rainfall in the UK. <http://www.met.rdg.ac.uk/~mike/autumn2000.html>.

CEH Met Office, 2001. To what degree can the October/November 2000 Flood Events be Attributed to Climate Change? Defra FD2304 Technical Report, Centre for Ecology and Hydrology, Wallingford.

Crooks, S.M., Kay, A.L., Reynard, N.S., 2009. Regionalised Impacts of Climate Change on Flood Flows: Hydrological Models, Catchments and Calibration. Report to Department for Environment, Food and Rural Affairs, FD2020 project milestone, CEH Wallingford, November 2009, 59pp.

Crooks, S.M., Naden, P.S., 2007. CLASSIC: a semi-distributed modelling system. Hydrology and Earth System Sciences 11 (1), 516–531.

Gringorten, I.I., 1963. A plotting rule for extreme probability paper. Journal of Geophysical Research 68, 813–814.

Hall, J.W., Twyman, C., Kay, A.L., 2005. Influence diagrams for representing uncertainty in climate-related propositions. Climatic Change 69 (2–3), 343–365.

IPCC, 2007. Climate Change 2007: Synthesis Report. Contribution of Working Groups I, II and III to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change. IPCC, Geneva, Switzerland, 104pp (Core Writing Team, Pachauri, R.K., Reisinger, A. (Eds.)).

Jenkins, G.J., Perry, M.C., Prior, M.J.O., 2007. The Climate of the United Kingdom and Recent Trends. Met Office Hadley Centre, Exeter, UK.

Kay, A.L., Davies, H.N., 2008. Calculating potential evaporation from climate model data: a source of uncertainty for hydrological climate change impacts. Journal of Hydrology 358, 221–239. doi:10.1016/j.jhydrol.2008.06.005.

Kay, A.L., Reynard, N.S., Jones, R.G., 2006. RCM rainfall for UK flood frequency estimation. I. Method and validation. Journal of Hydrology 318, 151–162.

Kay, A.L., Jones, D.A., Crooks, S.M., Kjeldsen, T.R., Fung, C.F., 2007. An investigation of site-similarity approaches to generalisation of a rainfall–runoff model. Hydrology and Earth System Sciences 11, 500–515.

Marsh, T.J., Dale, M., 2002. The UK floods of 2000–2001: a hydrometeorological appraisal. Journal of the Chartered Institution of Water and Environmental Management 16, 180–188.

Massacand, A.C., 2003. Forecasting of Extreme Seasonal Precipitation: Insight into the ECMWF Potential. ECMWF Technical Memorandum No. 415, 16pp.

Monteith, J.L., 1965. Evaporation and environment. Symposia of the Society for Experimental Biology 19, 205–234.

Morris, D.G., Flavin, R.W., 1990. A digital terrain model for hydrology. In: Proc. 4th International Symposium on Spatial Data Handling, Zurich, vol. 1, pp. 250–262.

Nozawa, T., Nagashima, T., Shiogama, H., Crooks, S.A., 2005. Detecting natural influence on surface air temperature change in the early twentieth century. Geophysical Research Letters 32, L20719.

Oudin, L., Hervieu, F., Michel, C., Perrin, C., Andreassian, V., Anctil, F., Loumagne, C., 2005. Which potential evapotranspiration input for a lumped rainfall–runoff model? part 2 — towards a simple and efficient potential evapotranspiration model for rainfall–runoff modelling. Journal of Hydrology 303, 290–306.

Pall, P., 2006. Constraints on, and Attribution of, Changes in Extreme Precipitation Under Climate Change. PhD Thesis, University of Oxford.

Pall, P., Aina, T., Stone, D.A., Stott, P.A., Nozawa, T., Hilberts, A.G.J., Lohmann, D., Allen, M.R., 2011. Anthropogenic greenhouse gas contribution to flood risk in England and Wales in autumn 2000. Nature 470, 382–386.

Pope, V.D., Stratton, R.A., 2002. The processes governing horizontal resolution sensitivity in a climate model. Climate Dynamics 19, 211–236.

Reynolds, R.W., Rayner, N.A., Smith, T.M., Stokes, D.C., Wang, W.Q., 2002. An improved in situ and satellite SST analysis for climate. Journal of Climate 15, 1609–1625.

Stone, D.A., Allen, M.R., 2005. The end-to-end attribution problem: from emissions to impacts. Climatic Change 71, 303–318.

Stott, P.A., Stone, D.A., Allen, M.R., 2004. Human contribution to the European heatwave of 2003. Nature 432, 610–614.

Stott, P.A., Mitchell, J.F.B., Allen, M.R., Delworth, T.L., Gregory, J.M., Meehl, G.A., Santer, B.D., 2006. Observational constraints on past attributable warming and predictions of future global warming. Journal of Climate 19, 3055–3069. doi:10.1175/JCLI3802.1.

Thompson, N., Barrie, I.A., Ayles, M., 1982. The Meteorological Office Rainfall and Evaporation Calculation System: MORECS (July 1981). Hydrological Memorandum No. 45, Met Office, Bracknell.