The Impact of Uncertain Precipitation Data on Insurance Loss Estimates Using a Flood Catastrophe Model

Christopher C. Sampson[1], Timothy J. Fewtrell[2], Fiachra O'Loughlin[3], Florian Pappenberger[4], Paul B. Bates[3], Jim E. Freer[3], Hannah L. Cloke[5]

[1]Corresponding author, School of Geographical Sciences, University of Bristol, BS8 1SS, UK, chris.sampson@bristol.ac.uk

[2]Willis Global Analytics, Willis Group, 51 Lime Street, London, EC3M 7DQ

[3]School of Geographical Sciences, University of Bristol, BS8 1SS, UK

[4]European Centre for Medium-Range Weather Forecasts, Reading, RG2 9AX

[5]Department of Geography and Environmental Science, University of Reading, RG6 6DW

Abstract

Catastrophe risk models used by the insurance industry are likely subject to significant uncertainty, but due to their proprietary nature and strict licensing conditions they are not available for experimentation. In addition, even if such experiments were conducted, these would not be repeatable by other researchers because commercial confidentiality issues prevent the details of proprietary catastrophe model structures from being described in public domain documents. However, such experimentation is urgently required to improve decision making in both insurance and re-insurance markets. In this paper we therefore construct our own catastrophe risk model for flooding in Dublin, Ireland in order to assess the impact of typical precipitation data uncertainty on loss predictions. As we consider only a city region rather than a whole territory and have access to detailed data and computing resources typically unavailable to industry modellers, our model is significantly more detailed than commercial products. The model consists of four components, a stochastic rainfall module, a hydrological and hydraulic flood hazard module, a vulnerability module and a financial loss module. Using these we undertake a series of simulations to test the impact of driving the stochastic event generator with four different rainfall data sets: ground gauge data, gauge corrected rainfall radar, meteorological re-analysis data (ERA-Interim) and a satellite rainfall product (CMORPH). Catastrophe models are unusual because they use the upper three components of the modelling chain to generate a large synthetic database of unobserved and severe loss-driving events for which estimated losses are calculated. We find the loss estimates to be more sensitive to uncertainties propagated from the driving precipitation datasets than to other uncertainties in the hazard and vulnerability modules, suggesting that the range of uncertainty within catastrophe model structures may be greater than commonly believed.

1.0 Introduction and Literature Review

The repeated occurrence of high profile flood events across the British Isles, such as Carlisle in January 2005, Gloucestershire in July 2007 and Dublin in October 2011, has resulted in sustained public, commercial, political and scientific interest in flood risk. Recent catastrophic flood events in other countries, such as the Indus floods in Pakistan (2010), the Australian and Thai floods (2011), and the Central European Floods (2013) have further raised the profile of flood risk through extensive global news coverage. The economic cost associated with flooding is often high. It is estimated that the October and November 2000 floods in the UK caused insured losses of £1.3 billion (Pall et al., 2011), whilst household losses resulting from the summer 2007 floods reached £2.5 billion, with business losses accounting for a further £1 billion (Chatterton et al., 2010; Pitt, 2008). The reinsurance firm Munich Re estimates that total economic losses from the Australian and Thailand events were USD 2.8 billion and USD 40 billion respectively (Munich Re, 2012), whilst the reinsurance firm Swiss Re estimates these figures at USD 6.1 billion and USD 30 billion (Swiss Re, 2012). Much of the total insured loss was from business interruption and contingent business interruption claims, demonstrating the global impact of such events.

Due to the scale of potential losses the insurance and reinsurance industries require accurate flood risk estimates, and the current accepted approach is to use calculation chains comprising linked stochastic and physically-based models. These calculation chains, known as catastrophe or 'CAT' models, are at the core of a methodological framework employed by the insurance industry to produce probabilistic estimates of natural catastrophe risk. First developed in the late 1980s to model earthquake risk, the methodology was widely adopted throughout the 1990s to model a range of hazards such as tropical cyclone windstorms and storm-surge floods (Wood et al., 2005). Today, such models are relied upon by the insurance and risk management industries to guide a wide range of financial decisions (Grossi et al., 2005). Whilst being applicable to a wide range of hazards, commercial 'vendor' CAT models typically share a common structure that can be broken down into four component parts:

i. Stochastic module. The stochastic module is used to generate a database of plausible event driving conditions. In the case of flooding, this could be a database of extreme precipitation events over the catchment(s) that drive fluvial or pluvial risk where the insured assets are located. The stochastic module is typically trained on historically observed data. As observational records of natural hazards are typically short (10^1 years) relative to return periods of interest to the insurance industry (10^2 to 10^4 years), the module must be capable of simulating events whose magnitude exceeds that of the largest observed event.
ii. Hazard module. The hazard module is used to simulate a selection of events from the database generated by the stochastic module. The hazard module needs to produce an estimate of damage-driving characteristics across the area where insured assets are located. In the case of flooding this is likely to take the form of a map of water depths.
iii. Vulnerability module. The vulnerability module calculates the expected damage to assets as a result of the event modelled by the hazard module. These damages are

expressed as a damage ratio that varies between 0 (no damage) and 1 (total loss). Factors influencing the susceptibility of an asset to damage may include terms such as building age, occupancy type, construction materials, or height. These parameters are typically uncertain, and thus vulnerability may be represented by an uncertain measure that maps the expected damage to a particular asset against a continuously variable hazard module output such as water depth and/or velocities. This is often done using a beta distribution with non-zero probabilities for damage ratios of 0 and 1.

iv. Financial module. The financial module transforms the per event damage estimates produced by the vulnerability module into an estimate of insured loss. Estimates of insured losses are generated by aggregating the losses from all assets being considered and applying policy conditions such as limits and deductibles to the total estimate of loss. The financial module resamples the database of simulated events to produce a large number of different time series realisations from which time-aggregated loss curves are produced.

As with any study that involves the modelling of environmental processes, it is important to address the presence of uncertainty within the system. Previous studies that consider flood risk using a model cascade framework have found the ‘driving’ component at the top of the cascade to be the most significant source of uncertainty (Kay et al., 2008; McMillan and Brasington, 2008). Cloke et al. (2012) also highlight the problem of uncertainty propagating from global and regional climate models when attempting to assess flood hazard on the River Severn in the UK. Due to their focus on low frequency, high magnitude events, the stochastic component of a CAT model inevitably has to extrapolate to event scales beyond those in the observational record. As a result, the loss estimates produced by CAT models may be particularly sensitive to the propagation of uncertainty in the data used to drive the stochastic component. If true, this will indicate that CAT model cascades are even more sensitive to driving uncertainties than other previously studied hydrological model cascades. As the stochastic module forms the driving component of a CAT model, this study attempts to assess the uncertainties derived from the choice of data used to calibrate, and therefore govern, the behaviour of the stochastic module. In order to provide context for this analysis, further limited analysis of the effect of parametric uncertainty within the hazard module and uncertainty within the vulnerability model were performed.

When developing a CAT model, it is important to bear in mind that the recent Solvency II legislation in Europe (European Parliament and European Council, 2009) requires that model users are able to understand and communicate how their models function. Many users will not be specialists in the field of environmental sciences and thus such legislation favours simpler model structures. A further reason to favour simpler model structures lies in their ease of application. Simpler models typically require less data than complex models, and therefore should be easier to apply to the wide array of locations that are of interest to insurance markets. It is also important to minimise the computational requirements of the cascade due to the extremely large number of events that may need to be modelled in order to estimate losses at very high return periods. The model structure used for this study was developed with such operational concerns in mind, and as such simple methods capable of delivering adequate performance against historical observations were favoured.

The following section of the literature review briefly explains the choice of model components employed in this study. The methodology that follows explains in more detail how each component functions within a CAT model framework.

1.1 Stochastic Module

Stochastic rainfall models are data-based approaches that use statistical information extracted from observations to parameterise a mechanism used to generate synthetic rainfall records. Such approaches are attractive in this context due to their relative simplicity and low computational costs. Stochastic rainfall models can generally be split into two methodological groups, namely profile-based and pulse-based, although there have been attempts to test alternative approaches including chaotic (Rodriguez-Iturbe et al., 1989; Sivakumar et al., 2001), artificial neural networks (Burian and Durran, 2002), simulated annealing (Bárdossy, 1998) and multiplicative cascade disaggregation (Gaume et al., 2007). Profile-based models typically use statistical distributions to characterise storms in terms of intensity, duration and inter-arrival time, whereas pulse-based models use statistical distributions to define raincells occurring within larger storm units characterised by duration and inter-arrival time distributions. The raincells take the form of pulses with individual durations and intensities, and the total storm intensity at a given time can therefore be calculated through summation of all active cell intensities at that time.

For the purposes of building a flood catastrophe model, it is necessary to select a model formulation that is able to reproduce the extreme events that drive flood risk. Several comparison studies have noted that while pulse-based models are able to simulate storm inter-arrival times and precipitation averages well, their ability to capture extreme statistics is variable and often particularly poor over short timescales (Cameron et al., 2000; Khaliq and Cunnane, 1996; Onof and Wheater, 1993; Verhoest et al., 1997). By comparison, the profile-based models have shown skill at simulating extreme events (Acreman, 1990; Blazkov and Beven, 1997; Cameron et al., 2000), although their ability to perform well for such events is dependent on the length and quality of the historical record used for their calibration. Due to its demonstrated ability to represent a range of different extreme precipitation events, this study employs a model developed from the profile-based Cumulative Distribution Function Generalised Pareto Distribution Model (CDFGPDM) of Cameron et al. (1999).

1.2 Hazard Module

In order to convert the rainfall input from the stochastic module into an estimate of water depths across the spatial domain containing the insured assets, two components are required: a hydrological rainfall-runoff model to produce an estimate of river discharge and a hydraulic model to transform the estimate of river discharge into a map of water depths. Hydrological models vary in complexity from process-rich, spatially distributed models such as the Systeme Hydrologique Europeen (Abbott et al., 1986a, 1986b) and the US Department of Agriculture's Soil and Water Assessment Tool (Muleta and Nicklow, 2005), to simple, spatially lumped conceptual models such as TOPMODEL (Beven and Kirkby, 1979) or HBV (Bergstrom and Forsman, 1973). Increasing model complexity inevitably entails increased dimensionality and data requirements, a situation that is often at odds with the requirements of a CAT model. Furthermore, the fundamental argument as to how much complexity is valuable in a model has not yet been conclusively answered in the literature (Bai et al., 2009; Beven, 1989; Blöschl and Sivapalan, 1995), and a number of studies have found that model performance does not necessarily improve with increased model complexity (e.g. Butts et al., 2004; Reed et al., 2004). As a result, a simple variant of the HBV model (Bergstrom and Forsman, 1973; Bergström and Singh, 1995; Seibert and Vis, 2012) was chosen here thanks to its ease of application, low data and computation cost and demonstrated performance across a large number of studies (Cloke et al., 2012; Deckers et al., 2010; e.g. Seibert, 1999).

In order to translate estimates of river discharge into maps of water depth across a domain, an additional hydraulic modelling component is required. The flow of water in urban areas is inherently multi-dimensional and requires a model of commensurate dimensionality able to run at the fine spatial resolutions needed to represent urban environments where vulnerability to losses will be most critical. The computational expense of such simulations has resulted in a research drive to develop efficient methods of modelling high resolution two-dimensional shallow water flows. Hunter et al. (2008) benchmarked a suite of commercial and research 2D codes on a small urban test scenario and found all to give plausible results, with predicted water depths typically differing by less than the vertical error in the topographical error despite the model governing equations varying from full 2D shallow-water equations to x-y decoupled analytical approximations to the 2D diffusion wave. These results are supported by further recent studies that have found highly efficient simplifications of the 2D shallow water equations to be appropriate for a number of urban inundation modelling (Neal et al., 2011; Néelz and Pender, 2010). As a result, this study employs the latest inertial formulation of the highly efficient 2D storage cell inundation model LISFLOOD-FP (Bates et al., 2010). This approach offers a more sophisticated representation of flow dynamics than the methods adopted by most vendor CAT models; vendor models typically represent the channel and floodplain using a 1D model, with a limited number of models also offering 2D modelling of 'off-floodplain' processes (AIR Worldwide, 2013; RMS, 2006).

1.3 Vulnerability Module

Flood damage models typically use water depths to predict damage based on a depth-damage function derived from empirical data (Black et al., 2006; Merz and Thieken, 2009; Merz et al., 2004), synthetic data (Penning-Rowsell et al., 2005), or a combination of both (ICPR, 2001). Studies have demonstrated significant variation in the curves produced by each methodology (Merz and Thieken, 2009; Merz et al., 2010), with the greater accuracy of empirical data compared to synthetic data (Gissing and Blong, 2004) being countered by the limited transferability of empirical data between sites (Smith, 1994). Depth damage functions are inherently uncertain due to the large number of factors that may influence the level of damage that results from a water depth. These include, but are not limited to, building type, building construction method, building age, building condition and precautionary measures). Although there is ongoing research into the possibility of accounting for these factors explicitly within multivariate depth-damage functions (Kreibich et al., 2010; Merz et al., 2013), such methods have not been widely adopted within the insurance market as a lack of observed damage data in most regions prevents calibration of such complex functions. Many commercial models instead attempt to represent much of the total CAT model uncertainty within the vulnerability module by sampling around the depth-damage curve. This is typically done using beta distributions to represent the probabilities of experiencing a range of damage ratios of between 0 and 1 for a given water depth. As the focus of this study is on the uncertainty due to driving precipitation data, we employ fixed depth-damage curves for most of our experiments. However, as recent studies (Jongman et al., 2012; Moel and Aerts, 2010) have suggested that the vulnerability module may be the dominant source of uncertainty, we also undertake a limited analysis using uncertain vulnerability curves in section 3.4 in order to provide an indication of relative contributions to modelled uncertainty. The curves and distribution parameters were supplied by Willis Global Analytics and were derived from a combination of synthetic and empirical data, claims data, and industry expertise.

1.4 Financial Module

Due to their proprietary nature, public domain literature describing the financial component of CAT models is very limited. Generally the role of financial modules is to transform damage estimates from the vulnerability module into estimates of insured ground up loss (i.e. loss before application of deductibles and/or reinsurance) before aggregating the location-specific losses to produce portfolio-wide loss estimates for a given event. These can then be transformed into estimates of gross insured loss by applying policy conditions such as deductibles, coverage limits, triggers, reinsurance terms, etc. (Grossi et al., 2005). Where the hazard module is computationally expensive, the financial module is often used to fit curves to the loss distributions generated by calculation chain, allowing much larger synthetic databases of event losses to be generated by subsequent resampling of the distributions. The primary output of a financial model takes the form of a curve that describes the probability of exceeding a certain level of loss within a fixed time period (typically annual). The two most common exceedence probability (EP) curves are the annual occurrence exceedence probability (OEP), representing the probability of a single event loss exceeding a certain level in a given year, and the aggregate exceedence probability (AEP), representing the probability of aggregate losses exceeding a certain level in a given year. Details of the financial module employed in this study are shown in section 2.2.4.

2.0 Study Site, Data and Methodology

Dublin, Ireland, was selected as the test site for this study due to its flood prone nature and the availability of suitable data sources. Historically, Dublin has been prone to fluvial, pluvial and tidal flooding, with fluvial risk being largely concentrated along two rivers, namely the River Dodder and the River Tolka. The River Dodder has its source in the Wicklow Mountains to the South of the city and drains an area of approximately 113 km^2. High rainfall intensities over the peaks of the Wicklow Mountains (annual totals can reach

2000 mm) coupled with steep gradients results in the River Dodder exhibiting flashy responses to storm events, with a typical time to peak of less than 24 hours. The River Tolka has its source in gently sloping farmland to the North West of the city and drains an area of approximately 150 km^2; it exhibits a slightly less flashy response than the Dodder with a time to peak of approximately 24 hours. As a result of the short catchment response times, sub-daily (ideally hourly) rainfall data are required to drive hydrological models of the rivers. Both catchments contain a mixture of urban and rural land use. Figure 1 is a map showing the location of these rivers and their respective catchment boundaries upstream of their gauging stations, as well as the boundary of the hydraulic model, the location of river gauging stations and the location of rain gauges. The calculation chain uses hydrological models of the Dodder and Tolka catchments to drive a hydraulic model of the rivers as they flow through the city and out into Dublin Bay. A third major river, the River Liffey, is also shown. The Liffey is not modelled in this study as its flow is controlled by three reservoirs that supply a hydroelectric generator upstream; serious flooding downstream of these features has not been observed since their construction was completed in 1949. River flow records are available from 1986 to present on the River Dodder and 1999 to present on the River Tolka.

[FIGURE 1 AROUND HERE]

In section 2.1, the four types of precipitation data (ground rain gauge, radar, meteorological reanalysis and satellite) used to drive the model are introduced along with the methods used to derive a catchment average precipitation series from each type of data. This step was required as using the stochastic module to generate extremely long (>500,000 years) spatial rainfall fields on an hourly time step would not have been computationally feasible, nor was it necessary given the input requirements of the simple hydrological model used here. The four types of precipitation data were chosen to represent the range of rainfall products available, from the high resolution localised gauge and radar data to the coarser (but globally available) reanalysis and satellite products. The record lengths of the different data sources were variable, but all four were available for the period January 2002 – May 2009; for experiments comparing the different data sources this was the period used.

In section 2.2, the components and data used to build and calibrate the stochastic, hazard, vulnerability and financial modules are presented.

2.1.1 Rain Gauge Record

The catchments surrounding Dublin are relatively well served by a network of rain gauges operated by Dublin City Council and the Irish weather service, Met Éireann. The gauges are primarily daily, with hourly weather stations sited at Dublin airport and Casement aerodrome. However, the network is subject to the usual limitations of gauge data which include missing data and inconsistent recording periods across the network. While some of the daily rain gauges have been operating for over 100 years, others were recently installed or retired. The gauges shown in figure 1 are the ones selected for use in this study following a significant pre-processing effort to check the availability of uninterrupted records from each gauge for periods coinciding with the available river flow records.

The daily catchment average time series were constructed by generating a gridded precipitation record at 50 m resolution for each of the catchments; the relatively fine grid was chosen due to the negligible computational cost of this process. The contribution of each daily gauge within a catchment to a given grid cell was calculated using an inverse distance weighting function. The difference in altitude between a given gauge and grid cell was also accounted for by correction using a precipitation-altitude gradient derived from the gauge record. Once the precipitation in all cells within a catchment was calculated, the catchment average precipitation was obtaining by averaging the value across all cells. The daily record was then distributed according to the nearest hourly station (Casement Aerodrome in the Dodder; Dublin Airport in the Tolka) to produce an hourly catchment average record.

2.1.2 Radar Record

The radar rainfall data were provided by the Met Éireann from a C-band radar located at Dublin Airport. A number of different products are produced for this radar, and the 1 km pre-gridded 15 minute Precipitation Accumulation (PAC) product is used in this study. The PAC product estimates the rainfall intensity at 1 km above the topographical surface, and the data were supplied for the period 2002 – 2009. Pre-processing was required to remove an echo signal present over mountainous parts of the Dodder catchment that was expressed in the data as anomalous near-continuous low intensity rainfall. An hourly timestep catchment average series was generated by averaging the cells that fell within the boundaries of a catchment. Whilst radar data are able to provide an estimate of the spatial distribution of precipitation, correction using ground-based observations is required in order for reasonable estimates of rainfall intensities (Borga, 2002; Germann et al., 2006; O'Loughlin et al., 2013; Steiner et al., 1999). Adjustment factors were therefore used to match the radar-derived catchment rainfall volume to the gauge-derived catchment rainfall volume on a three-monthly basis. The adjustment factor values were assumed to be time invariant for the duration of each three month period (Gjertsen et al., 2004).

2.1.3 ECMWF ERA-Interim Reanalysis

ERA-Interim is a global atmospheric reanalysis produced by the European Centre for Medium-Range Weather Forecasts (ECMWF) (Dee et al., 2011). The reanalysis covers the period 1979-present and produces gridded surface parameters. The ERAI configuration has a spectral T255 horizontal resolution, which corresponds to approximately 79 km spacing on a reduced Gaussian grid. The vertical resolution is using 60 model levels with the top of the atmosphere located at 0.1 hPa. ERA Interim data have been used in a wide range of applications such as mapping of drought, fire, flood and health risk (Pappenberger et al., 2013). Precipitation data are available in the form of 3-hour rainfall accumulation totals. Three–hourly timestep catchment average precipitation time series were produced using a weighted average of the ERA-Interim cells that covered the catchment, where weights were assigned based on the fraction of the catchment covered by each cell.

2.1.4 CMORPH Satellite Precipitation

The Climate Prediction Center morphing method (CMORPH) precipitation record is produced by using motion vectors derived from half-hourly interval geostationary satellite infrared imagery to propagate passive microwave precipitation estimates (Joyce et al., 2004). Data are available from 1998 – present day at a 3 hourly timestep on a 0.25 degree spatial grid. Three-hourly timestep catchment average precipitation time series were produced in the same way as with the ERA-Interim reanalysis data.

2.2.0 Catastrophe Model Framework

The CAT model framework employed in this study replicates the logic used by proprietary commercial models but uses detailed and transparent components that allow us to experiment in a controlled and repeatable fashion. The stochastic event generator creates a long time series of rainfall events that are used to drive the hazard module. When a flood event occurs, the predicted water depths are input into the vulnerability module to produce an estimate of loss. The event ID and loss ratio (event loss expressed as a percentage of the total sum insured across the portfolio) are recorded in an event loss table. The number of events occurring in each year is also recorded. Finally, the financial module resamples the event loss table in order to produce an aggregate annual loss exceedence probability (AEP) curve. Table 1 summarises the implications of a number of key uncertainties and assumptions present in the four modules.

[TABLE 1 AROUND HERE]

As we demonstrate in section 3.0, the sampling uncertainty associated with extreme events can be large. This is because different realisations of events with a common return period produce different losses, and multiple stochastic model runs of a given length may generate very different sets of extreme events. Whilst it is possible to handle this uncertainty by producing an extremely large stochastic event set, using the hazard module to simulate every small scale event that occurs in such a large event set is not computationally feasible. This computational restraint requires that a simple event similarity criterion based on hydrograph peak and hydrograph volume is used to test for similar previously simulated events. Events are only simulated with the hydraulic model if the hydrograph peak or hydrograph volume on either river differs from a previously simulated event by more than a preset threshold of 10%. If this requirement is not met then it is assumed that a similar event has already been simulated, and the calculated loss from this earlier simulation is selected and added again to the event loss table.

2.2.1 Stochastic Rainfall Module

The Cumulative Distribution Function Generalised Pareto Distribution Model CDFGPDM employed here uses statistical distributions to define storms in terms of mean durations, intensities and inter-arrival times. The CDFGPDM is a profile-based stochastic rainfall model that generates a series of independent rainstorms and ‘inter-arrival’ periods (dry-spells) via a Monte Carlo sampling procedure. The model retains the Eagleson (1972) approach of characterising a storm in terms of inter-arrival time, duration and mean intensity whilst incorporating a profiling component to distribute the total precipitation throughout the duration of the storm. Storms in the observational record are classed by duration and their intensities are recorded using empirical cumulative distribution functions (CDFs). In order to enable the simulation of storms of greater duration or intensity than in the observational record, the tails of the CDFs are modelled using maximum likelihood Generalised Pareto Distributions (GPD). The threshold above which the GPD was fitted depended on the number of observations in each class and ranged from the 75^{th} to 95^{th} quantile. The empirical CDFs are then combined with their modelled GPD tails to generate hybrid distributions from which storm characteristics can be sampled. Previous studies have argued that rainfall runoff models can be realistically driven by such a model structure as the shape parameter within the GPD allows a wide range of upper tail shapes to be adequately captured (Cameron et al., 2000, 1999). Following Cameron et al. (1999) we here define a rainstorm as any event with an intensity of $\geq$ 0.1 mm/hour, a duration of $\geq$ 1 hour and an inter-arrival time of $\geq$ 1 hour, where no zero-rainfall periods are permitted within a storm. It should be noted that for the ERA-Interim and CMORPH driven models, the minimum duration and inter-arrival times were 3 hours due to the 3 hour timestep of these products. This definition encapsulates all recorded precipitation in the 1 hour interval historical records available for Dublin, making it appropriate for characterisation and subsequent generation of continuous rainfall records. The rainstorm generation procedure is identical to the method detailed in Cameron et al. (1999). In order to evaluate the model's ability to recreate the extremes seen in the observed series, a total of 50 synthetic series of 40 years length were simulated using the rain gauge derived series for the Dodder catchment. The annual maximum rainfall totals (ANNMAX) for each duration class were extracted from the synthetic series and plotted against their counterparts from the observed catchment average series (figure 2). The reduced variate plots show that the observed ANNMAX values are well bracketed by those from the 50 synthetic series, indicating the ability of the model to recreate a reasonable distribution of extreme events suited to a study of flood risk.

[FIGURE 2 AROUND HERE]

Due to the need to limit model complexity and computational expense, it was necessary to assume a spatially uniform rainfall across the modelled catchments. Such an assumption may be justified for Dublin as the modelled catchments are relatively small (<130 km^2) and floods in this region are driven by large weather systems such as frontal depressions and decaying hurricanes rather than by small scale convective cells. The gauge-based catchment average records produced for the Dodder and Tolka catchments were tested for correlation, yielding a Pearson's linear correlation coefficient of 0.89 and a Kendall tau of 0.69. These values indicate that rainfall in the two catchments is indeed strongly correlated; however the lack of perfect correlations implies that the approach will result in a slight overestimation of domain-total rainfall for a given event. The assumption allows a spatially uniform, time varying rainfall series to be generated for all catchments by training the CDFGPDM on a single, centrally located, observation site. However, due to significant variation in altitude across the domain, it was necessary to correct the rainfall intensities of the generated series for each catchment as the observed precipitation intensity distributions varied between the catchment mean records and the central training site. To achieve this, a quantile-quantile bias correction method (Boé et al., 2007) was used on each observed record type in turn, where adjustment factors for each quantile bin were obtained by comparing the observed time series at the training site to the observed catchment average rainfall series. Therefore, for each of the modelled catchments, a different set of adjustment factor values were generated for the ground gauge, radar, ERA-Interim and CMORPH data, allowing precipitation time series to be generated in which the correct precipitation intensity distributions of each individual catchment are persevered despite all catchments sharing a common temporal rainfall pattern.

2.2.2 Hazard Module

The hazard module consists of a hydrological model and a hydraulic model. The hydrological model employed here is the widely used conceptual rainfall runoff model HBV (Bergstrom and Forsman, 1973; Bergström and Singh, 1995). While there are many variants of the HBV model, the one used for this study is most closely related to HBV Light (Seibert and Vis, 2012). The model uses precipitation, temperature and potential evaporation as inputs, the latter of which is calculated from extraterrestrial radiation and temperature using the McGuiness model (McGuinness and Bordne, 1972), to produce an estimate of river discharge at the gauge station locations shown in figure one with an hourly timestep. Model calibration was undertaken to generate behavioural parameter sets for each precipitation data source in each catchment. Initially, the 15-parameter space was explored using Monte Carlo simulation and parameter ranges were set by visually identifying upper and lower limits from the resultant simulations. Where the model did not exhibit detectable parameter range limits, ranges from previous studies were employed (Abebe et al., 2010; Cloke et al., 2012; Shrestha et al., 2009). Once defined, the parameter ranges were sampled using Latin hypercube Monte Carlo sampling to produce 100,000 parameter sets, a number of samples which proved computationally feasible whilst providing adequate exploration of the parameter space. The parameter sets were then used to simulate discharge during a period for which observations were available, and those that failed to produce behavioural simulations, defined by a Nash-Sutcliffe (NS) score exceeding a threshold of 0.7 (Nash and Sutcliffe, 1970), were discarded. The choice of performance measure and threshold used to define what constitutes a behavioural simulation is necessarily subjective (Beven and Freer, 2001); NS was chosen as it is particularly influenced by high flow performance, and the threshold of 0.7 was selected following visual inspection of hydrographs generated from a preliminary sample of parameter sets. In order to assign weights, the behavioural parameter sets were then ranked and weighted by their ability to minimise error in the top 0.1% of the flow duration curve. Due to computational constraints imposed by the subsequent hydraulic model, the number of behavioural parameter sets was limited to the 100 highest ranked sets. Weighting was performed by calculating the inverse sum of absolute errors between the simulated and observed series in the top 0.1% of the flow duration curve for each of the behavioural parameter sets. These values were then normalised to give the best performing parameter set a weight of 1 and the worst a weight of 0. This approach favours behavioural parameter sets that best simulate high-flow periods and is therefore appropriate for a study concerned with flood risk.

Initially, attempts were made to calibrate HBV using each precipitation data type. However, only those simulations driven using the gauge-derived precipitation data were able to satisfy the behavioural NS threshold in all catchments. Models driven using ECMWF and CMORPH data were especially poor; this may be explained by their reduced spatial and temporal resolution compared to the gauge and radar data. As the model was only able to adequately represent observed catchment flow characteristics using the behavioural parameter sets identified using gauge data, it was therefore decided to employ these parameter sets for all simulations. The very large number of event simulations required to produce an EP curve precluded HBV parametric uncertainty from being incorporated directly into the CAT model; such an approach would have further increased the required computational resource to an unfeasible level. Due to this limitation, the highest ranked parameter set produced using gauge data was used to generate the EP curves. The impact of parametric uncertainty is addressed separately on an event basis in section 3.3, where the weighted behavioural parameter sets are used to produce uncertain loss estimates with 5-95% confidence intervals for four synthetic flood events.

The hydraulic model LISFLOOD-FP (Bates and De Roo, 2000) is used to generate flood inundation maps from the event hydrographs produced by HBV. The configuration employed here uses a subgrid representation of the channel (Neal et al., 2012b) coupled to a 2D flood plain model that uses a simplified 'inertial formulation' of the shallow water equations (Bates et al., 2010) solved using the numerical method of de Almeida et al. (2012). The channel models include weirs and were constructed using surveyed river cross sections supplied by Dublin City Council, and the digital elevation model (DEM) for the 144 km^2 2D hydraulic model was constructed from 2 m resolution bare-earth LiDAR data that was coarsened to 10 m and 50 m resolution (1,440,000 and 57600 cells respectively) using bilinear resampling (Fewtrell et al., 2008). Where >50% of the surface area of a cell was occupied by building(s), identified through Ordinance Survey Ireland data, the cell elevation was increased by 10 m to become a 'building cell'. Model calibration of channel floodplain friction was undertaken by driving the hydraulic model with observed discharges and comparing the observed and simulated flood inundation extents for the August 1986 Hurricane Charlie and the November 2002 flood events. These are the largest events for which observed discharge and inundation data are available, with the 2002 event generating \$47.2 million in unindexed losses (AXCO, 2013), and have been attributed with ~700 and ~100 year return periods respectively (RPS Consulting Engineers, 2008; RPS MCOS, 2003). The extent of the larger 1986 event was digitised from hand drawn post-event flood outline maps, which included indications of dominant flow directions, although the completeness of these maps is uncertain. The November 2002 flood outlines were supplied by Dublin City Council. Both of these datasets will be subject to considerable uncertainty as they were constructed from eye witness accounts and post-event ground based observations; they should therefore be considered as approximations of the true maximum extents. Observed and simulated flood outlines for the calibration events are shown in figure 3. The quantitative

F-squared performance measure (Werner et al., 2005) was calculated for each calibration run, with the optimised model yielding values of 0.62 and 0.44 for the 10 m and 50 m resolution models respectively. Some of the variation between the observed and simulated extents may be explained by errors in the observed data; some may also be explained by land development and engineering works that occurred between the events and the date on which the modern DEM terrain data were collected; this latter factor may have an especially strong influence for the 1986 event results. Nevertheless, the F-squared values still compare favourably with a previous study of urban inundation modelling (Fewtrell et al., 2008), in which it is noted that performance of models in urban areas is strongly affected by the ability of the DEM to represent urban structures; subsequent studies have also highlighted the influence of detailed terrain features on urban inundation processes (Fewtrell et al., 2011; Sampson et al., 2012). These findings are further evidenced here, as the reduced representation of buildings on the 50 m DEM removes flow restrictions and results in an overestimation of flood extents with a corresponding reduction in water depths near the channel. Despite this, qualitative assessment of the modelled dynamics with reference to the observations suggests that, at both resolutions, the model is capturing the dominant process well, with water entering the floodplain in the correct areas. Unfortunately, the computational expense of the 10 m resolution model was several orders of magnitude greater than the 50 m model, resulting in simulation times of several hours compared to ~ 20 seconds for a 48 hour event. Due to this cost, the 50 m model was adopted for use within the CAT model. Whilst this will result in some lost predictive skill relative to the 10 m model, the representation of 2D flow both on and off the floodplain ensures the model remains more sophisticated than the 1D or quasi-2D approaches typically employed by vendor CAT models. The implication to loss estimates of this decision is briefly discussed in section 3.3.

[FIGURE 3 AROUND HERE]

2.2.3 Vulnerability Module

A synthetic portfolio of insured properties, modelled on real data, was provided by Willis Global Analytics for use in this study. This was necessary to preserve the anonymity of real policy holders, and the portfolio was built by resampling a distribution of asset values for the region. As is common for insurance portfolios, the data were aggregated to postcode level. The portfolio took the form of an insured sum for three lines of business (residential, commercial and industrial) for each postcode area. It is common practice in industry to disaggregate such datasets using proxy data (Scott, 2009), and the approach adopted here to use the National Oceanic and Atmospheric Administration (NOAA) Impervious Surface Area (ISA) dataset as a proxy for built area (Elvidge et al., 2007). This method assumes a linear relationship between the percentage of a grid cell that is impervious and its insured value, and allows the sum insured within each postcode to be distributed around the postcode area based on ISA pixel values. From these data we built a simple industry exposure database (IED) that contained the values of insured assets for each line of business within each grid cell.

When a cell is flooded, the damage sustained within the cell is calculated using depth-damage functions supplied by Willis Global Analytics that were derived from historical data of floods in European cities. In this paper we employ both a simplified deterministic depth damage curve approach and a more sophisticated uncertain vulnerability function. The simplified approach involves separate curves for the residential, commercial and industrial lines of business that relate the water depth within a cell to the percentage of the cell's insured value that is lost. These simple curves therefore represent a mean damage ratio and were used for all experiments other than the vulnerability uncertainty analysis in order to reduce computational cost and better isolate the subject of each experiment. The more sophisticated functions used in the vulnerability uncertainty analysis sample around the fixed curves using modified beta distributions. Here, the depth in a cell determines the mean damage ratio as well as the probabilities of zero damage (P0) and total loss (P1). A stratified antithetic sample of values between 0 and 1 is performed, with all values below P0 being assigned a damage ratio of 0 and all values above P1 being assigned a damage ratio of 1. The values between P0 and P1 are rescaled to between 0 and 1 and used to sample from a beta distribution whose parameters are calculated based on the mean damage ratio, P0, P1 and an assumed variance. The result is a sample of damage ratios, with a mass of values at zero, a mass of values at one, and an intermediary range drawn from a beta distribution. As the water depth in a cell increases, the mass of zero damages becomes smaller, the mass of total losses becomes larger, and the mean of the intermediary sampled beta distribution moves towards one (total loss). This method is currently used by Willis on an operational basis and therefore represents industry practice at the date of publication.

2.2.4 Financial Module

The financial module employed here is used to aggregate simulated losses from the hazard module across a specified aerial unit (here the entire domain) before generating and resampling occurrence and loss distributions from the results. The occurrence distribution represents the distribution of event counts for a given time period (here defined as one year) using an empirical CDF. The main body of the loss distribution is modelled using an empirical CDF, with a GPD fitted to the tail to produce a smooth curve where data are sparse. A synthetic series can then be rapidly generated by adopting a Monte Carlo resampling method. This procedure samples first from the occurrence distribution to find the number (n) of events occurring in a given year. The loss distribution is then sampled n times to assign a loss to each event. Finally, the annual aggregate loss is found by summing the losses for that year. By repeating this process a large number of times, multiple synthetic series can be generated. From these series, an annual AEP curve can be generated that includes confidence intervals derived from the spread of values at any given return period. The annual AEP curve is a standard insurance tool that is used to express the expected probability of exceeding a given level of loss over a one year period, i.e. the expected '1 in 100 year loss' is equivalent to a loss with an annual exceedence probability (AEP) of 0.01.

3.0 Results - Event Sampling Uncertainty

A known source of uncertainty within a CAT model originates in the event generation procedure used to build an event set. This is referred to as 'primary uncertainty' by the insurance industry (Guin, 2010). A key difficulty in calculating the expected loss at a given AEP is that the predicted insured loss will vary from one model run to another due to the random component of the stochastic module. One method of reducing this 'sampling uncertainty' is to simulate a series that is considerably longer than the desired recurrence interval (Neal et al., 2012a). Alternatively a large number of realisations can be simulated, and the expected loss can then be defined by the mean loss across the realisations. The second method also allows the sampling uncertainty to be investigated by looking at the spread of values across the realisations. The number of realisations that it is feasible to simulate is determined by the required series length and the available computational resource. Here the stochastic module is trained using the rain gauge record and used to generate 500 realisations of a 1000 year rainfall series in order to investigate the effect of sampling uncertainty on the 1-in-1000 year loss.

The object of this experiment is to determine the number of realisations required to adequately capture the range of possible losses at a given event scale. One way to examine such 'sampling uncertainty is to assemble batches of realisations and observe how key descriptors (such as the mean loss or standard deviation of losses) vary between batches. By altering the number of realisations in each batch, it is possible to observe how the variation of descriptors between batches changes as the batch size changes. It is then possible to predict the expected average variation, in terms of the descriptors, between the simulated batch of *n* realisations and any other batch of *n* realisations.

To do this, the maximum losses recorded in each of the 500 realisations were randomly sampled to produce batches containing 5, 10, 25, 50, 100 or 250 loss ratios ('batch A'). The process was repeated to produce a second batch ('batch B') of identical size to batch A. The mean and standard deviation of loss ratios in batch A ($\bar{L}_A$ and s_A) were then calculated and compared to their equivalent values in Batch B ($\bar{L}_B$ and s_B), yielding two simple measures:

$$M = |\bar{L}_A - \bar{L}_B| \qquad (1)$$

$$S = |s_A - s_B| \qquad (2)$$

[FIGURE 4 AROUND HERE]

By repeating this process a large number of times (10,000 for each batch size), the expected uncertainty due to sampling variability can be assessed. The results of this experiment are shown in figure 4a, where *M* is expressed as a percentage of the mean 1-in-1000 year loss across all 500 realisations and *S* is equivalently expressed as a percentage of the standard deviation across all 500 realisations. The plots show that differences between batches A and B decrease as the number of samples within a batch increases, with the median value of *M* decreasing from 23.0% to 3.8% as the batch size increases from 5 to 250. This finding can be explained by the underlying distribution of loss ratios being increasingly well represented as the sample size is increased; this is observed in the diminishing value of *S* as sample size increases. By transforming the median values of *M* with reciprocal $1/M^2$ and fitting a linear regression model, the expected value of M for the 500 realisations was calculated as 2.7%. This indicates that the mean loss ratio of any 500 simulated realisations will typically differ from any other batch of 500 realisations by ~ 3% of the mean loss ratio itself; the same process yields a value of 2.3% for the standard deviations (figure 4b). Primary uncertainty is an accepted facet of catastrophe modelling and, relative to inherent aleatory uncertainty, uncertainty of this order due to sampling variability is reasonable (Guin, 2010). Whilst the uncertainty caused by sampling variability could be reduced by significantly increasing the number of realisations simulated, the additional computational cost of such an increase would be large and the benefit questionable in the presence of other uncertainties within the calculation chain. For the purpose of this study we identify 50 realisations as the minimum required; at this level the mean and median values of M and S are <10% of the mean and standard deviation of all 500 realisations respectively. The practical implication of this analysis is that it is necessary for the hazard module to simulate >50 time series of length equal to the return period of interest.

3.1 Variability across data sources

The availability and quality of observed precipitation records varies greatly between sites. In order to investigate how the use of different types of precipitation data might affect predicted losses, each of the data types described in section 2.1 was used to train the stochastic module. The training record length was defined by the longest period for which a continuous record was available from all data sources; this ran from the 1st January 2002 to the 1st May 2009. This period is clearly shorter than ideal and it is likely that the true variability within each data source is underrepresented as a result; however it was necessary to ensure that the records were of equal length over the same period in order to fairly compare between data types. All parameters in the hazard, vulnerability and financial modules were identical across the simulations. Taking a maximum return period of interest to be the 1-in-10,000 year event, 500,000 years' worth of simulations was performed for each data type (giving the required 50 realisations of the 1-in-10,000 year event). The annual aggregate EP curves resulting from these model runs are shown in figure 5, with uncertainty bounds that represent the 5 – 95% confidence intervals generated by the financial module. Also plotted are the modelled losses of two observed historical floods (August 1986 and November 2002), produced by driving the hydraulic and vulnerability components with observed river discharges.

[FIGURE 5 AROUND HERE]

It is immediately apparent from figure 5 that the different precipitation data sets produce very different EP curves despite the fact that each record covered the same spatial area over a common period of time. At certain points the difference can be as great as an order of magnitude – for example, the ERA-Interim driven model predicts a 1-in-100 year (AEP = 10^{-2}) loss ratio of 0.02% whereas the CMORPH driven model predicts a loss ratio of 0.17%. The pronounced differences between the curves can be explained in terms of the ability of each of the data sources to represent the local rainfall patterns. The gauge and radar driven models produced EP curves of similar shape, with losses from the radar driven model being slightly lower than from the gauge record. Their relative similarity compared to the ERA-Interim and CMORPH driven models was expected as both are detailed local data sources rather than global products. Furthermore the adjustment factors for radar rainfall intensity were derived from the gauge record so that the two records had equal 3-monthly rainfall volumes. As a result, storms were usually captured in both records and attributed with similar rainfall totals, yielding similar stochastic model calibrations and therefore similar loss projections.

The curves produced by the ERA-Interim and CMORPH driven models differ greatly from those produced by the local gauge and radar datasets. The ERA-Interim curve shows only gradual growth in losses as the return period increases to the maximum modelled value of the 1-in-10,000 year event, and at all return periods the ERA-Interim model under predicts compared to the other data sources. By contrast, the losses predicted by the CMORPH driven model are consistently higher than the others, especially at lower return periods. Figure 6a shows cumulative daily precipitation for all four data types. As previously found by Kidd et al. (2012) in a study of rainfall products over Northwest Europe, CMORPH is found to consistently underestimate rainfall totals compared to the local data whereas ERA-Interim consistently overestimates rainfall totals. Given the pattern of cumulative rainfall totals, the opposite pattern found in the loss projections is initially surprising. However, once hourly rainfall intensities are considered (figure 6b) the findings can be explained. CMORPH is found to underestimate rainfall totals in this region because of the limited sensitivity of satellite products to very low intensity rainfall ('drizzle') (Kidd et al., 2012). However, it exhibits higher rainfall intensities in the upper (>95th) quantiles of rainfall intensity than the other records. Severe storms in the CMORPH record typically had slightly higher rainfall volumes than the same storms in other records, the result of which is an increased expected loss at all return periods. ERA-Interim has the opposite problem whereby the frequency of low intensity precipitation is over predicted and high intensity precipitation is severely underestimated.

[FIGURE 6 AROUND HERE]

3.2 Uncertainty due to record length

A similar approach to the above comparison between data sources was adopted to examine the sensitivity of projected losses to the length of record used to train the stochastic module. For this test the gauge precipitation data were cropped to produce training records of 5, 10, 20 and 40 years in length. The training records share a common end date (September 2011) and therefore the longer records extend further into the past. As with the data sources test, all other parameters were held constant across the other components, and the resulting EP curves are plotted in figure 7. The EP curves demonstrate that altering the training record length has a significant impact on the projected losses for a given return period. At AEP = 10^{-2}, the median expected loss ratio ranges from 0.05 to 0.28; at AEP=10^{-3}, representing the 1-in-1000 year event, the expected loss ratios vary from 0.12 to 0.60. The relative overestimation of loss ratios by the 5 year training data set demonstrates how the presence of a large event in a short training set is able to skew the results. There are two storms that generate exceptionally high precipitation volumes in the 40 year observed record, and the second of these falls within the final five years that form the 5 year training record. When trained with this short record, the stochastic module inevitably over predicts the rate of occurrence of such storms, leading to an overestimation of expected flood losses. Modelled uncertainty increases as the return period increases; in the case of the 10 year training period, the range of modelled losses at the 10^{-4} AEP level is greater than the median estimate of 0.36%.

[FIGURE 7 AROUND HERE]

3.3 Hazard module uncertainty

In order to provide some context for the uncertainty associated with the choice of driving data, the uncertainty resulting from the choice of parameter set used with HBV was also investigated. Due to computational limitations it was not feasible to produce EP curves for a large number of parameter sets, so instead we focussed on individual events. The largest event was extracted from each of four 500 year runs of the stochastic module. Each event was then simulated using the 100 best performing HBV parameter sets, all of which had previously been selected and assigned weights as described in section 2.2.2. The resulting hydrographs were then used to drive the hydraulic model, and the event loss from each simulation was calculated and weighted according to their respective parameter set weights. Figure 8 shows each event hyetograph, the range of hydrographs produced by the different parameter sets on both the Dodder and Tolka rivers, and the resulting weighted CDF of loss ratios. The weighted 95% confidence interval values for peak discharge, hydrograph volume and loss ratio are shown in table 2.

[FIGURE 8 AROUND HERE]

[TABLE 2 AROUND HERE]

The results of this exercise demonstrate the impact of parametric uncertainty within the hydrological model on expected losses. For the smallest of the events (event 3), the ratio of the 95^{th} to 5^{th} quantile peak discharges for the Dodder and Tolka was ~1.1. Despite these relatively modest increases, the ratio of 95^{th} to 5^{th} quantile losses across the whole domain was ~1.7. For a larger event (event 4), the equivalent 95^{th} to 5^{th} quantile peak discharge ratio increased to ~1.2 and yielded a ratio of losses of ~3.25.

The high sensitivity of expected losses to relatively smaller percentage changes in hydrograph peak or volume is due to the fact that losses are only affected by the part of the hydrograph that drives flood inundation – namely the portion of flow that is out-of-bank. This region of the hydrograph is clearly sensitive to parametric uncertainty, leading to the high degree of uncertainty in modelled losses exhibited here. It should also be noted that these results are sensitive to the subjective choice of behavioural threshold and performance measures employed. Had a higher threshold been chosen, the available parameter space from which behavioural sets could be selected would be smaller, leading to a reduction in the modelled loss ratio uncertainty. However, despite parametric uncertainty clearly being important, in the context of this study the choice of driving precipitation data source remains the greater source of uncertainty in modelled losses.

As noted in the hazard module description (section 2.2.2), the high computational cost of hydraulic simulations on a 10m grid prevented the finer resolution model from being adopted. The earlier qualitative assessment of the hydraulic model at 50 m relative to 10 m indicated that both exhibited similar first order dynamics, with the coarser model producing a greater simulation extent with reduced water depths as a result of the reduced building blockages and terrain smoothing. In order to provide a general indication as to how this might affect loss estimates, the losses from the 10 m and 50 m calibration simulations were calculated. These calculations yielded loss ratios of 0.101 and 0.146 respectively, indicating that areas of deep localised flooding present in the 10 m simulations were generating high losses not adequately captured by the 50 m model. However, although a more detailed study is required before firm conclusions can be drawn regarding the importance of hydraulic model resolution in this context, this result does suggest that the contribution of the hydraulic model to the total hazard model uncertainty may be small relative to the hydrological model.

3.4 Vulnerability module uncertainty

Contemporary CAT models typically account for uncertainty within the vulnerability module by using historical claims data to develop a distribution of damage ratios for any given water depth as described in sections 1.3 and 2.2.3. In order to investigate the uncertainty imparted onto the EP curves by the vulnerability module, the 500,000 years' worth of hazard module simulations performed for section 3.1 were coupled to the uncertain vulnerability module. This process generated EP curves for each data source in which the 5-95% confidence intervals are defined by uncertainty within the vulnerability module (figure 9).

[FIGURE 9 AROUND HERE]

Figure 9 demonstrates that the uncertainty imparted by the vulnerability module is large relative to uncertainty generated by the financial model (figure 5) for small to moderate event scales (1 in 10 to 1 in ~250 year). However, for the more extreme events the two contribute uncertainty of a broadly similar magnitude. This is due to the nature of uncertainty within the vulnerability module. At small event scales the vulnerability module is able to generate a wide range of loss ratios even when water depths are relatively low. This produces significant uncertainty within the EP curve relative to a model that uses fixed depth-damage curves, as loss ratios from the fixed curves will typically be low when water depths are shallow. However, during more extreme events where high loss ratios dominate the curve due to increased water depths, the relative uncertainty of the vulnerability model is seen to decrease as both the uncertain and fixed vulnerability methods cannot generate losses exceeding 1 (total loss). This exhibition of asymptotic behaviour highlights the fact that uncertainties vary both in absolute terms and relatively to each other as event scale changes.

4.0 Discussion

The results presented above examine how the loss estimates produced by a flood catastrophe model are affected by the choice of data used to drive the model's stochastic component. Parametric uncertainty from the hydrological model has also been examined on an event basis to contextualise the scale of uncertainty induced by the stochastic component and uncertainty from the vulnerability module has also been modelled. The findings highlight the difficulty in producing robust EP curves using a cascade methodology, as the uncertainty associated with each component is large and increases as event scale increases. Furthermore, not all sources of uncertainty have been considered – for example flood defence failure rates. Despite this, the model presented here is very detailed compared to standard industry practice, and contains detailed local information (such as river channel geometry and features) that would often be unavailable under the time and financial constraints of most commercial catastrophe modelling activities. The required computational resource would also exceed what is practicably available if models of this detail were extended to cover entire national territories. As a result, the uncertainty estimates made in this study are likely to be conservative. The CMORPH and ERA-Interim precipitation records have global coverage and are typical of the kind of product that could be used to drive a commercial CAT model. However, the hydrological model was unable to generate behavioural results when driven by these data sources, indicating their inability to produce realistic storm precipitation and thus runoff in the modelled catchments. It is therefore unsurprising that they generated EP curves that were both very different to each other and to the curves produced using more detailed local records. Examination of the observed precipitation records reveals that the precipitation intensity distributions vary significantly between the data sources. The observed records are relatively short; a common record across all four data sources was only available for a little over seven years due to the short length of radar records and gaps in the ground gauge data. The divergence in estimates of precipitation totals for heavy storms between the observational records is reflected in the synthetic series produced by the stochastic module, and this divergence inevitably continues as the simulated event scale increases. This results in the pronounced differences in higher return period loss estimates produced by the model when trained with each of the data sources in turn. Whilst access to longer overlapping records might have reduced the severity of this divergence, the consistently different storm rainfall intensities recorded by the four data types means that the stochastic module would still be expected to generate very different estimates of high return period rainfall events depending on which data it was driven with. It is also worth noting at this point that we did not consider the parametric uncertainty associated with fitting GPDs to the precipitation intensity and duration tails; this source of epistemic uncertainty is likely to be large given the relatively short rainfall records to which the GPDs are fitted and therefore the true uncertainty is most likely greater than reported here. Unfortunately, investigating the impact of this on modelled losses would have required a number of runs of the entire model cascade that was computationally prohibitive.

The EP curves were also found to be sensitive to the length of record used to train the stochastic module. Unfortunately, satellite and model reanalysis precipitation records are typically short (CMORPH runs from the mid-1990's; ERA-Interim from 1979) and the results presented here demonstrated significant differences between the EP curves produced by records of 5, 10, 20 and 40 years in length. Lack of available data prevented longer records from being tested, but our results do indicate that extra care is required when using short (<10 years) records due to the ability of a single extreme observation to skew results. Furthermore, the fact that there is an appreciable difference between the 20 and 40 year curves suggests that records of at least 40 years in length should be used where possible. Future reanalysis products hoping to extend records further back in time may help to alleviate this issue; the European Reanalysis of Global Climate Observations (ERA-CLIM) project aims to provide a 100 year record dating back to the early 20th century. The impact of parametric uncertainty within HBV should also be of concern to practitioners. The model in this study was calibrated with detailed precipitation and discharge records and might therefore be considered tightly constrained compared to commercial models that will have to operate at national scales. Despite this, the variation in predicted loss ratios over a range of behavioural parameter sets for individual events was very large. Due to computational constraints we were unable to also consider uncertainty in the hydraulic model component of the hazard module, although it is believed that the hydraulic model is a relatively minor source of uncertainty in this context (Apel et al., 2008a). Former studies have indicated that topography is the dominant driver of uncertainty within hydraulic models if we consider the inflow boundary condition uncertainty to be associated with the hydrological model (Fewtrell et al., 2011; Gallegos et al., 2009; Schubert et al., 2008; Yu and Lane, 2006), and given the differences seen between the calibration runs at 10 m and 50 m resolution (figure 3) it is very likely that the uncertainty reported in this study is an underestimate of the total uncertainty present within the hazard module.

The final uncertainty source considered was the vulnerability module. This module was found to contribute significantly to the uncertainty at smaller event scales but, due to the inherently asymptotic nature of a damage function, its relative contribution was shown to decrease as event scale increased. Of particular interest is the fact that, in contrast to some previous studies (e.g. Moel and Aerts, 2010), the vulnerability module uncertainty is smaller than the uncertainty resulting from choice of data used to drive the hazard module. This is likely due to such studies using relatively constrained event scenarios in which under which hazard uncertainty is more limited than in a stochastic model. Studies which considered a wider range of events (Apel et al., 2008b; Merz and Thieken, 2009) have found uncertainty in the features controlling the occurrence and magnitude of events (e.g. stage discharge relationships, flood frequency analysis) to be similar to or greater than the vulnerability uncertainty, especially at larger event scales.

Spatial scales are an important consideration in the context of this study. The catchments modelled in this study are relatively small, and it is reasonable to suggest that the relatively coarse reanalysis and satellite products might perform better for major rivers where fluvial floods are driven by rainfall accumulations over longer time periods and large spatial areas. Some of their inherent traits, such as tendency for the reanalysis product to persistently 'drizzle' while underestimating storm rainfall accumulations, will negatively impact their applicability to flood modelling across most catchment scales although the severity of the effect may reduce as catchment sizes increase. However, it is wrong to assume that the dominant driver of flood risk is always large events on major rivers. A significant proportion of insurance losses resulting from the 2007 UK floods and 2013 Central European Floods can be classified ‘off-floodplain’ – that is to say they occurred either as a result of surface water (pluvial) flooding or as a result of fluvial flooding in small catchments (Willis, personal communication). This suggests that even when considering large events, the ability to produce realistic hazard footprints in small catchments remains critical and thus for practitioners concerned about such events, the findings of this paper remain relevant.

When considered together, the above findings make it difficult to commend a stochastic flood model driven by precipitation data as a robust tool for producing EP curves for use in portfolio analysis. The sensitivity of the stochastic component to the driving data is of fundamental concern due to the high degree of uncertainty in observed precipitation extremes, suggesting that alternative driving mechanisms such as flood frequency analysis should be evaluated in this context. Furthermore, the results demonstrate sensitivity to model parametric uncertainty that will be difficult to overcome. However, these shortcomings do not mean that such a model has no value. Although it may be difficult to use such a system to project accurately how often events of a certain magnitude will occur, and thus estimate probable losses over a given time window, the model could still be used to assess the relative risk of assets within a portfolio. We argue that understanding and quantifying the uncertainties generated by the stochastic and hazard modules for a given portfolio may be important to managing assets effectively. Although the computational demand of the hazard module in particular will likely render this unfeasible on an operational basis, studies such as this may be used to inform judgments regarding the total uncertainty within such model structures. A valuable exercise for users of commercial models may be to compare such findings to the uncertainty generated by their own models, many of which may attempt to account for hazard uncertainty via sampling widened distributions within the vulnerability module.

5.0 Conclusions

In this study, stochastic, hazard, vulnerability and loss modules have been assembled into a cascade framework that follows the same principles as an insurance catastrophe model. The model operates by generating a large synthetic series of events in the stochastic component which is then simulated by the hazard component. The vulnerability component assesses the damage and loss caused by each event, building up a database of occurrence intervals and event losses. Finally, the loss component resamples from the modelled occurrence and loss distributions, producing exceedence probability curves that estimate the expected annual aggregate loss for a range of return periods. The model simulates fluvial flood risk in Dublin, Ireland, and the components were calibrated using local historical observations where appropriate data were available.

A number of different precipitation datasets were tested with the model, including high resolution local gauge and radar records, model reanalysis records (ERA-Interim) and satellite records (CMORPH). The exceedence probability curves produced by the model were found to be very sensitive to the choice of driving precipitation data, with different driving datasets producing loss estimates that varied by more than an order of magnitude in some instances. Examination of the observational records reveals that the precipitation intensity distributions over a common period vary markedly between the different data types. These differences are inevitably reflected in the output produced by the stochastic module and result in large differences in the modelled magnitude of high return period events. The calculation chain was also found to be sensitive to the length of observational record available, with the presence of a large event in a short training set resulting in severe overestimation of losses relative to models driven by a longer record. The sensitivity of the model to parameterisation of the hydrological model was tested on an event basis. Modelled loss ratios were found to be highly sensitive to the choice of parameter set. Despite all being classified as behavioural, the loss ratios for one event varied by up to six times dependent on the parameter set selected. Finally uncertainty in the vulnerability module was considered. Due to the asymptotic nature of damage functions it was found to be a larger relative contributor at small event scales than large, although even at large scales its contribution remained high. However, the impact of both hydrological parameter uncertainty and vulnerability uncertainty were both smaller than the impact of uncertainty within the driving precipitation data.

Considered together, the results of this study illustrate the difficulty in producing robust estimates of extreme events. The uncertainty in the observed record, along with the short length of records relative to return periods of interest, is of particular concern as observed differences diverge when the event scale is extrapolated far beyond what has historically been observed. A lack of suitable observational data for model calibration makes it challenging to envisage how similar methods to those employed in this study could be used to produce the national scale models required by industry without uncertainty bounds becoming unmanageably high. Further issues that will compound these problems are the scarcity of data relating to the condition and location of flood defences, another important source of uncertainty (Gouldby et al., 2008), and the requirement to build models in data-poor developing regions where insurance market growth is greatest. The results of this study have emphasised the dramatic impact of data uncertainties on loss estimates, and it is important that the users and developers of catastrophe models bare such results in mind when assessing the validity of the uncertainty mechanisms within their models. At present, the combination of short record lengths and highly uncertain precipitation intensities during storm events make it difficult to recommend the use of rainfall-driven model cascades to estimate fluvial flood risk, especially where estimates of return period are necessary. Looking forward, increased resolution regional reanalysis products with improved rainfall process representation may help to reduce these uncertainties as may the assimilation of local data into global observational datasets to produce improved regional calibrations for rainfall products (Dinku et al., 2013). Further effort should also be concentrated on developing alternative means of characterising the loss driving properties of river basins. One such alternative may be to revisit methods based on geomorphology and flood frequency analysis (Leopold and Maddock, 1953; Meigh et al., 1997) in conjunction with modern observational databases (such as the Global Runoff Data Centre) and remotely sensed data. As supercomputing power continues to grow exponentially, large ensemble stochastic frameworks that combine such approaches will likely become tenable projects over the coming decade.

Acknowledgements

The work conducted for this project by Christopher Sampson formed part of his PhD research at the University of Bristol funded by a Natural Environment Research Council (NERC) Combined Award in Science and Engineering (NE/H017836/1) in partnership with the Willis Research Network (WRN). The authors would like to thank Tony Maguire and Michael Curtis at Dublin City Council for their assistance in supplying valuable data and local historical knowledge that informed this study, Ciaran Commins at Met Éireann for providing the radar data, and the ECMWF for supplying the ERA-Interim reanalysis data.

References

Abbott, M. B., Bathurst, J. C., Cunge, J. A., O'Connell, P. E. and Rasmussen, J.: An introduction to the European Hydrological System — Systeme Hydrologique Europeen, "SHE", 1: History and philosophy of a physically-based, distributed modelling system, Journal of Hydrology, 87(1–2), 45–59, doi:10.1016/0022-1694(86)90114-9, 1986a.

Abbott, M. B., Bathurst, J. C., Cunge, J. A., O'Connell, P. E. and Rasmussen, J.: An introduction to the European Hydrological System — Systeme Hydrologique Europeen, "SHE", 2: Structure of a physically-based, distributed modelling system, Journal of Hydrology, 87(1–2), 61–77, doi:10.1016/0022-1694(86)90115-0, 1986b.

Abebe, N. A., Ogden, F. L. and Pradhan, N. R.: Sensitivity and uncertainty analysis of the conceptual HBV rainfall–runoff model: Implications for parameter estimation, Journal of Hydrology, 389(3-4), 301–310, doi:10.1016/j.jhydrol.2010.06.007, 2010.

Acreman, M. C.: A simple stochastic model of hourly rainfall for Farnborough, England / Un modèle stochastique simple des pluies horaires de Farnborough, Angleterre, Hydrological Sciences Journal, 35(2), 119–148, doi:10.1080/02626669009492414, 1990.

AIR Worldwide: The AIR Inland Flood Model for Great Britain, [online] Available from: http://www.air-worldwide.com/Publications/Brochures/documents/AIR-Inland-Flood-Model-for-Great-Britain/ (Accessed 10 August 2013), 2013.

De Almeida, G. A. M., Bates, P., Freer, J. E. and Souvignet, M.: Improving the stability of a simple formulation of the shallow water equations for 2-D flood modeling, Water Resources Research, 48(5), doi:10.1029/2011WR011570, 2012.

Apel, H., Aronica, G. T., Kreibich, H. and Thieken, A. H.: Flood risk analyses—how detailed do we need to be?, Nat Hazards, 49(1), 79–98, doi:10.1007/s11069-008-9277-8, 2008a.

Apel, H., Merz, B. and Thieken, A. H.: Quantification of uncertainties in flood risk assessments, International Journal of River Basin Management, 6(2), 149–162, doi:10.1080/15715124.2008.9635344, 2008b.

AXCO: Insurance Market Reports: Ireland, s.l., 2013.

Bai, Y., Wagener, T. and Reed, P.: A top-down framework for watershed model evaluation and selection under uncertainty, Environmental Modelling & Software, 24(8), 901–916, doi:10.1016/j.envsoft.2008.12.012, 2009.

Bárdossy, A.: Generating precipitation time series using simulated annealing, Water Resour. Res., 34(7), PP. 1737–1744, doi:10.1029/98WR00981, 1998.

Bates, P. D., Horritt, M. S. and Fewtrell, T. J.: A simple inertial formulation of the shallow water equations for efficient two-dimensional flood inundation modelling, Journal of Hydrology, 387(1-2), 33–45, doi:10.1016/j.jhydrol.2010.03.027, 2010.

Bates, P. D. and De Roo, A. P. J.: A simple raster-based model for flood inundation simulation, Journal of Hydrology, 236(1-2), 54–77, doi:10.1016/S0022-1694(00)00278-X, 2000.

Bergstrom, S. and Forsman, A.: Development of a conceptual deterministic rainfall-runoff model, NORDIC HYDROLOGY, VOL 4, NO 3, P 147-170, 1973. 19 FIG, 3 TAB, 5 REF. [online] Available from: http://www.csa.com/partners/viewrecord.php?requester=gs&collection=ENV&recid=7401128 (Accessed 1 February 2013), 1973.

Bergström, S. and Singh, V. P.: The HBV model., in Computer models of watershed hydrology., pp. 443–476. [online] Available from: http://www.cabdirect.org/abstracts/19961904773.html (Accessed 2 February 2013), 1995.

Beven, K.: Changing ideas in hydrology — The case of physically-based models, Journal of Hydrology, 105(1–2), 157–172, doi:10.1016/0022-1694(89)90101-7, 1989.

Beven, K. and Freer, J.: A dynamic TOPMODEL, Hydrological Processes, 15(10), 1993–2011, doi:10.1002/hyp.252, 2001.

Beven, K. J. and Kirkby, M. J.: A physically based, variable contributing area model of basin hydrology / Un modèle à base physique de zone d'appel variable de l'hydrologie du bassin versant, Hydrological Sciences Bulletin, 24(1), 43–69, doi:10.1080/02626667909491834, 1979.

Black, A., Werritty, A. and Paine, J.: Financial costs of property damages due to flooding: the Halifax Dundee Flood Loss Tables 2005, University of Dundee, University of Dundee., 2006.

Blazkov, S. and Beven, K.: Flood frequency prediction for data limited catchments in the Czech Republic using a stochastic rainfall model and TOPMODEL, Journal of Hydrology, 195(1-4), 256–278, 1997.

Blöschl, G. and Sivapalan, M.: Scale issues in hydrological modelling: A review, Hydrological Processes, 9(3-4), 251–290, doi:10.1002/hyp.3360090305, 1995.

Boé, J., Terray, L., Habets, F. and Martin, E.: Statistical and dynamical downscaling of the Seine basin climate for hydro-meteorological studies, International Journal of Climatology, 27(12), 1643–1655, doi:10.1002/joc.1602, 2007.

Borga, M.: Accuracy of radar rainfall estimates for streamflow simulation, Journal of Hydrology, 267(1–2), 26–39, doi:10.1016/S0022-1694(02)00137-3, 2002.

Burian, S. J. and Durran, S. R.: Evaluation of an artificial neural network rainfall disaggregation model, Water Sci. Technol, 45(2), 99–104, 2002.

Butts, M. B., Payne, J. T., Kristensen, M. and Madsen, H.: An evaluation of the impact of model structure on hydrological modelling uncertainty for streamflow simulation, Journal of Hydrology, 298(1–4), 242–266, doi:10.1016/j.jhydrol.2004.03.042, 2004.

Cameron, D., Beven, K. and Tawn, J.: An evaluation of three stochastic rainfall models, Journal of Hydrology, 228(1-2), 130–149, doi:10.1016/S0022-1694(00)00143-8, 2000.

Cameron, D. S., Beven, K. J., Tawn, J., Blazkova, S. and Naden, P.: Flood frequency estimation by continuous simulation for a gauged upland catchment (with uncertainty), Journal of Hydrology, 219(3-4), 169–187, 1999.

Chatterton, J., Viavattene, C., Morris, J., Penning-Rowsell, E. and Tapsell, S.: The costs of the summer 2007 floods in England, Environment Agency, Bristol., 2010.

Cloke, H. L., Wetterhall, F., He, Y., Freer, J. E. and Pappenberger, F.: Modelling climate impact on floods with ensemble climate projections, Quarterly Journal of the Royal Meteorological Society, doi:10.1002/qj.1998, 2012.

Deckers, D. L. E. H., Booij, M. J., Rientjes, T. H. M. and Krol, M. S.: Catchment Variability and Parameter Estimation in Multi-Objective Regionalisation of a Rainfall–Runoff Model, Water Resour Manage, 24(14), 3961–3985, doi:10.1007/s11269-010-9642-8, 2010.

Dinku, T., Hailemariam, K., Maidment, R., Tarnavsky, E. and Connor, S.: Combined use of satellite estimates and rain gauge observations to generate high-quality historical rainfall time series over Ethiopia, International Journal of Climatology, doi:10.1002/joc.3855, 2013.

Eagleson, P. S.: Dynamics of flood frequency, Water Resour. Res., 8(4), PP. 878–898, doi:10.1029/WR008i004p00878, 1972.

Elvidge, C. D., Tuttle, B. T., Sutton, P. C., Baugh, K. E., Howard, A. T., Milesi, C., Bhaduri, B. and Nemani, R.: Global Distribution and Density of Constructed Impervious Surfaces, Sensors, 7(9), 1962–1979, doi:10.3390/s7091962, 2007.

European Parliament and European Council: Directive 2009/138/EC of the European Parliament and of the Council of 25 November 2009 on the taking-up and pursuit of the business of Insurance and Reinsurance (Solvency II), Official Journal of the European Union, L 335/1, 1–155, 2009.

Fewtrell, T. J., Bates, P. D., Horritt, M. and Hunter, N. M.: Evaluating the effect of scale in flood inundation modelling in urban environments, Hydrol. Process., 22(26), 5107–5118, doi:10.1002/hyp.7148, 2008.

Fewtrell, T. J., Duncan, A., Sampson, C. C., Neal, J. C. and Bates, P. D.: Benchmarking urban flood models of varying complexity and scale using high resolution terrestrial LiDAR data, Physics and Chemistry of the Earth, Parts A/B/C, 36(7-8), 281–291, doi:10.1016/j.pce.2010.12.011, 2011.

Gallegos, H. A., Schubert, J. E. and Sanders, B. F.: Two-dimensional, high-resolution modeling of urban dam-break flooding: A case study of Baldwin Hills, California, Advances in Water Resources, 32(8), 1323–1335, doi:10.1016/j.advwatres.2009.05.008, 2009.

Gaume, E., Mouhous, N. and Andrieu, H.: Rainfall stochastic disaggregation models: Calibration and validation of a multiplicative cascade model, Advances in water resources, 30(5), 1301–1319, 2007.

Germann, U., Galli, G., Boscacci, M. and Bolliger, M.: Radar precipitation measurement in a mountainous region, Quarterly Journal of the Royal Meteorological Society, 132(618), 1669–1692, doi:10.1256/qj.05.190, 2006.

Gissing, A. and Blong, R.: Accounting for variability in commercial flood damage estimation, Australian Geographer, 35(2), 209–222, doi:10.1080/0004918042000249511, 2004.

Gjertsen, U., Salek, M. and Michelson, D. B.: Gauge adjustment of radar-based precipitation estimates in Europe, in Proceedings of ERAD, vol. 7. [online] Available from: http://copernicus.org/erad/2004/online/ERAD04_P_7.pdf (Accessed 5 February 2013), 2004.

Gouldby, B., Sayers, P., Mulet-Marti, J., Hassan, M. A. A. M. and Benwell, D.: A methodology for regional-scale flood risk assessment, Proceedings of the ICE - Water Management, 161(3), 169–182, doi:10.1680/wama.2008.161.3.169, 2008.

Grossi, P., Kunreuther, H. and Patel, C. C.: Catastrophe modeling a new approach to managing risk, Springer Science+Business Media, New York, NY. [online] Available from: http://www.ebrary.com/ (Accessed 25 January 2013), 2005.

Guin, J.: Understanding Uncertainty, AIR Currents [online] Available from: http://www.air-worldwide.com/Publications/AIR-Currents/attachments/AIR-Currents--Uncertainty/ (Accessed 15 February 2013), 2010.

Hunter, N. M., Bates, P. D., Neelz, S., Pender, G., Villanueva, I., Wright, N. G., Liang, D., Falconer, R. A., Lin, B., Waller, S., Crossley, A. J. and Mason, D. C.: Benchmarking 2D hydraulic models for urban flooding, Proceedings of the ICE - Water Management, 161(1), 13–30, doi:10.1680/wama.2008.161.1.13, 2008.

ICPR: Atlas of flood danger and potential damage due to extreme floods of the Rhine, IKSR, Koblenz. [online] Available from: http://www.iksr.org/index.php?id=212&L=3 (Accessed 2 April 2012), 2001.

Jongman, B., Kreibich, H., Apel, H., Barredo, J. I., Bates, P. D., Feyen, L., Gericke, A., Neal, J., Aerts, J. C. J. H. and Ward, P. J.: Comparative flood damage model assessment: towards a European approach, Nat. Hazards Earth Syst. Sci., 12(12), 3733–3752, doi:10.5194/nhess-12-3733-2012, 2012.

Joyce, R. J., Janowiak, J. E., Arkin, P. A. and Xie, P.: CMORPH: A Method that Produces Global Precipitation Estimates from Passive Microwave and Infrared Data at High Spatial and Temporal Resolution, Journal of Hydrometeorology, 5(3), 487–503, doi:10.1175/1525-7541(2004)005<0487:CAMTPG>2.0.CO;2, 2004.

Kay, A. L., Davies, H. N., Bell, V. A. and Jones, R. G.: Comparison of uncertainty sources for climate change impacts: flood frequency in England, Climatic Change, 92(1-2), 41–63, doi:10.1007/s10584-008-9471-4, 2008.

Khaliq, M. N. and Cunnane, C.: Modelling point rainfall occurrences with the modified Bartlett-Lewis rectangular pulses model, Journal of hydrology, 180(1-4), 109–138, 1996.

Kidd, C., Bauer, P., Turk, J., Huffman, G. J., Joyce, R., Hsu, K.-L. and Braithwaite, D.: Intercomparison of High-Resolution Precipitation Products over Northwest Europe, Journal of Hydrometeorology, 13(1), 67–83, doi:10.1175/JHM-D-11-042.1, 2012.

Kreibich, H., Seifert, I., Merz, B. and Thieken, A. H.: Development of FLEMOcs – a new model for the estimation of flood losses in the commercial sector, Hydrological Sciences Journal, 55(8), 1302–1314, doi:10.1080/02626667.2010.529815, 2010.

Leopold, L. and Maddock, T.: The hydraulic geometry of stream channels and some physiographic implications, United States Geological Survey Professional Paper, Washington, (252), 1953.

McGuinness, J. L. and Bordne, E. F.: A comparison of lysimeter-derived potential evapotranspiration with computed values, US Dept. of Agriculture. [online] Available from: http://books.google.co.uk/books?hl=en&lr=&id=oqYoAAAAYAAJ&oi=fnd&pg=PA1&dq=A+comparison+of+lysimeter-derived+potential+evapotranspiration+with+computed+values&ots=xuB5thWSdc&sig=JTCrgCiHZ613JYpCAx_e9eohMLI (Accessed 5 February 2013), 1972.

McMillan, H. K. and Brasington, J.: End-to-end flood risk assessment: A coupled model cascade with uncertainty estimation, Water Resour. Res., 44, 14 PP., doi:10.1029/2007WR005995, 2008.

Meigh, J. R., Farquharson, F. A. K. and Sutcliffe, J. V.: A worldwide comparison of regional flood estimation methods and climate, Hydrological Sciences Journal, 42(2), 225–244, doi:10.1080/02626669709492022, 1997.

Merz, B., Kreibich, H. and Lall, U.: Multi-variate flood damage assessment: a tree-based data-mining approach, Nat. Hazards Earth Syst. Sci., 13(1), 53–64, doi:10.5194/nhess-13-53-2013, 2013.

Merz, B., Kreibich, H., Schwarze, R. and Thieken, A.: Review article "Assessment of economic flood damage," Nat. Hazards Earth Syst. Sci., 10(8), 1697–1724, doi:10.5194/nhess-10-1697-2010, 2010.

Merz, B., Kreibich, H., Thieken, A. and Schmidtke, R.: Estimation uncertainty of direct monetary flood damage to buildings, Nat. Hazards Earth Syst. Sci., 4, 153–163, 2004.

Merz, B. and Thieken, A. H.: Flood risk curves and uncertainty bounds, Nat Hazards, 51(3), 437–458, doi:10.1007/s11069-009-9452-6, 2009.

Moel, H. and Aerts, J. C. J. H.: Effect of uncertainty in land use, damage models and inundation depth on flood damage estimates, Natural Hazards, 58(1), 407–425, doi:10.1007/s11069-010-9675-6, 2010.

Muleta, M. K. and Nicklow, J. W.: Sensitivity and uncertainty analysis coupled with automatic calibration for a distributed watershed model, Journal of Hydrology, 306(1–4), 127–145, doi:10.1016/j.jhydrol.2004.09.005, 2005.

Munich Re: Significant natural catastrophes worldwide 2011: 10 costliest natural catastrophes ordered by overall losses, Munich Re, Munich. [online] Available from: http://www.munichre.com/app_pages/www/@res/pdf/NatCatService/annual_statistics/2011/2011_mrnatcatservice_natural_disasters2011_ordered_by_eco_en.pdf (Accessed 19 December 2012), 2012.

Nash, J. E. and Sutcliffe, J. V.: River flow forecasting through conceptual models part I — A discussion of principles, Journal of Hydrology, 10(3), 282–290, doi:10.1016/0022-1694(70)90255-6, 1970.

Neal, J., Keef, C., Bates, P., Beven, K. and Leedal, D.: Probabilistic flood risk mapping including spatial dependence, Hydrological Processes, 27(9), 1349–1363, doi:10.1002/hyp.9572, 2012a.

Neal, J., Schumann, G. and Bates, P.: A subgrid channel model for simulating river hydraulics and floodplain inundation over large and data sparse areas, Water Resources Research, 48(11), doi:10.1029/2012WR012514, 2012b.

Neal, J., Villanueva, I., Wright, N., Willis, T., Fewtrell, T. and Bates, P.: How much physical complexity is needed to model flood inundation?, Hydrological Processes, 26(15), 2264–2282, doi:10.1002/hyp.8339, 2011.

Néelz, S. and Pender, G.: Benchmarking of 2D Hydraulic Modelling Packages, Environment Agency, Bristol., 2010.

O'Loughlin, F., Bruen, M. and Wagener, T.: Parameter sensitivity of a watershed-scale flood forecasting model as a function of modelling time-step, Hydrology Research, 44(2), 334, doi:10.2166/nh.2012.157, 2013.

Onof, C. and Wheater, H. S.: Modelling of British rainfall using a random parameter Bartlett-Lewis rectangular pulse model, Journal of Hydrology, 149(1-4), 67–95, 1993.

Pall, P., Aina, T., Stone, D. A., Stott, P. A., Nozawa, T., Hilberts, A. G. J., Lohmann, D. and Allen, M. R.: Anthropogenic greenhouse gas contribution to flood risk in England and Wales in autumn 2000, Nature, 470(7334), 382–385, doi:10.1038/nature09762, 2011.

Pappenberger, F., Wetterhall, F., Dutra, E., Di Giuseppe, F., Bogner, K., Alfieri, L. and Cloke, H.: Seamless forecasting of extreme events on a global scale, in Climate and Land Surface Changes in Hydrology, IAHS Publ. 359, Gothenburg, Sweden., 2013.

Penning-Rowsell, E., Johnson, C., Tunstall, S., Tapsell, S., Morris, J., Chatterton, J., Green, C., Wilson, T., Koussela, K. and Fernandez-Bilbao, A.: The benefits of flood and coastal risk management: a manual of assessment techniques. [online] Available from: http://www.vliz.be/imis/imis.php?module=ref&refid=117138 (Accessed 4 February 2013), 2005.

Pitt, M.: Learning lessons from the 2007 floods: an independent review by Sir Michael Pitt, Cabinet Office, London., 2008.

Reed, S., Koren, V., Smith, M., Zhang, Z., Moreda, F., Seo, D.-J. and DMIP Participants, and: Overall distributed model intercomparison project results, Journal of Hydrology, 298(1–4), 27–60, doi:10.1016/j.jhydrol.2004.03.031, 2004.

RMS: Germany River Flood, [online] Available from: https://support.rms.com/publications/Germany_Flood.pdf (Accessed 10 August 2013), 2006.

Rodriguez-Iturbe, I., Power, B. F. D., Sharifi, M. B. and Georgakakos, K. P.: Chaos in rainfall, Water Resour. Res., 25(7), PP. 1667–1675, doi:10.1029/WR025i007p01667, 1989.

RPS Consulting Engineers: River Dodder Catchment Flood Risk Management Plan: Hydrological Analysis Report, Dublin City Council, Dublin., 2008.

RPS MCOS: River Tolka Flooding Study: Final Report, Dublin City Council, Dublin., 2003.

Sampson, C. C., Fewtrell, T. J., Duncan, A., Kashif, S., Horritt, M. S. and Bates, P. D.: Use of terrestrial laser scanning data to drive decimetric resolution urban inundation models, AWR, 41, 1–17, doi:10.1016/j.advwatres.2012.02.010, 2012.

Schubert, J., Sanders, B., Smith, M. and Wright, N.: Unstructured mesh generation and landcover-based resistance for hydrodynamic modeling of urban flooding, Advances in Water Resources, 31(12), 1603–1621, doi:10.1016/j.advwatres.2008.07.012, 2008.

Scott, I.: Distributing Aggregated Census Data to higher geographic resolutions using GIS datasets, City University, London., 2009.

Seibert, J.: Regionalisation of parameters for a conceptual rainfall-runoff model, Agricultural and Forest Meteorology, 98–99, 279–293, doi:10.1016/S0168-1923(99)00105-7, 1999.

Seibert, J. and Vis, M. J. P.: Teaching hydrological modeling with a user-friendly catchment-runoff-model software package, Hydrology and Earth System Sciences Discussions, 9(5), 5905–5930, doi:10.5194/hessd-9-5905-2012, 2012.

Shrestha, D. L., Kayastha, N. and Solomatine, D. P.: A novel approach to parameter uncertainty analysis of hydrological models using neural networks, Hydrol. Earth Syst. Sci., 13(7), 1235–1248, 2009.

Sivakumar, B., Sorooshian, S., Gupta, H. V. and Gao, X.: A chaotic approach to rainfall disaggregation, Water Resour. Res., 37(1), PP. 61–72, doi:10.1029/2000WR900196, 2001.

Smith, D. I.: Flood damage estimation- A review of urban stage-damage curves and loss functions, Water S. A., 20(3), 231–238, 1994.

Steiner, M., Smith, J. A., Burges, S. J., Alonso, C. V. and Darden, R. W.: Effect of bias adjustment and rain gauge data quality control on radar rainfall estimation, Water Resources Research, 35(8), 2487–2503, doi:10.1029/1999WR900142, 1999.

Swiss Re: SIGMA 2/2012, Swiss Re, Zurich. [online] Available from: http://media.swissre.com/documents/Press_Release_sigma_2_2012_en.pdf (Accessed 19 December 2012), 2012.

Verhoest, N., Troch, P. A. and De Troch, F. P.: On the applicability of Bartlett-Lewis rectangular pulses models in the modeling of design storms at a point, Journal of Hydrology, 202(1-4), 108–120, 1997.

Werner, M., Hunter, N. and Bates, P.: Identifiability of distributed floodplain roughness values in flood extent estimation, Journal of Hydrology, 314(1-4), 139–157, doi:10.1016/j.jhydrol.2005.03.012, 2005.

Wood, R. M., Drayton, M., Berger, A., Burgess, P. and Wright, T.: Catastrophe Loss Modelling of Storm-Surge Flood Risk in Eastern England, Philosophical Transactions: Mathematical, Physical and Engineering Sciences, 363(1831), 1407–1422, 2005.

Yu, D. and Lane, S. N.: Urban fluvial flood modelling using a two-dimensional diffusion-wave treatment, part 1: mesh resolution effects, Hydrol. Process., 20(7), 1541–1565, doi:10.1002/hyp.5935, 2006.

Tables

	Key Data	Key Uncertainties	Key Implications	Additional Assumptions	Key Implications
Stochastic Module	- Rainfall - DEM	Short observational record lengths	Limited data to constrain GPD fits to tails of rainfall intensity/duration distributions	Uniform rainfall pattern	Likely over-estimation of modelled losses
		Precipitation intensities vary between data sources	Modelled losses highly sensitive to chosen data source		
Hazard Module: Hydrological Model	- Rainfall - Temperature - PET - Discharge	Parametric uncertainty	Modelled losses sensitive to parameterisation and calibration	River Dodder artificial reservoirs not modelled	Overestimation of losses, especially when antecedent conditions are dry and reservoir level would be low
		Small number of flood events in discharge records			
		Observed flood discharges uncertain			
		Small number of flood events in discharge records choice of behavioural performance measure	Uncertainty range dependent on performance measure		
Hazard Module: Hydraulic Model	- Discharge - Flood extents - River channel geometry - DEM	Errors in observed extents	Unknown sensitivity of modelled losses to hydraulic model structure	No significant flood defence additions since observed events	Model may simulate losses in newly defended areas
		Observed flood discharges uncertain			
		Roughness coefficients		Choice of which events to simulate based on hydrograph peak and volume	Relationship between hydrograph properties and loss may be oversimplified
		Unrepresented channel features			
		DEM resolution	Flood extents and depths influenced by DEM; losses not grid independent	Depth in building cell assumed to be mean of surround cell depths	Likely over-estimation of modelled losses
Vulnerability Module	- Water depths - Postcode areas - Depth damage curves - ISA data	ISA data is low resolution	Likely errors in loss calculation as true location of assets is unknown	Fixed damage fixed for a given flood depth	Loss for a given depth would really vary depending on building type
		Depth damage curves highly uncertain	Loss estimates vary depending on choice of depth damage curve	Period of inundation not considered	Possible Over/under estimation of losses for short/long duration events respectively
Financial Module	- Event loss table	-	-	Policy terms such as deductibles and limits not included	Overestimation of losses compared to 'real' portfolios

Table 1. Table showing the required data sources for each module, along with key uncertainties, assumptions and their respective implications for modelled losses.

	Measure	Event 1	Event 2	Event 3	Event 4
Dodder	Peak Discharge (m^3)	212 - 256	185 - 226	185 - 203	250 – 291
	Volume (x 10^7 m^3)	1.69 – 1.89	1.66 – 1.84	1.76 – 1.97	1.74 – 1.92
Tolka	Peak Discharge (m^3)	125 - 150	130 - 147	113 - 124	118 – 139
	Volume (x 10^7 m^3)	1.50 – 1.64	1.54 – 1.64	1.49 – 1.60	1.35 – 1.47
Entire Domain	Loss Ratio (%)	0.03 – 0.14	0.04 – 0.07	0.03 – 0.05	0.04 – 0.13

Table 2. Weighted 5^{th} – 95^{th} quantile values for event based HBV uncertainty simulations.

Figure Captions

Figure 1. Map of Dublin region. Modelled rivers are shown by thick blue lines. Hydrological model boundaries are shown in red. Hydraulic model boundary is shown in yellow. Rain gauge locations shown by black crosses. River flow gauges are shown by yellow triangles.

Figure 2. Modelled annual maximum rainfall totals for each duration class compared to observational record for the Dodder catchment. The annual maxima for each class in the 40 year catchment average observed record are ranked and plotted using Gringorten plotting positions (black circles). The process was repeated for 50 x 40 year simulated series (grey crosses).

Figure 3. Hydraulic model calibration results. Red shaded area shows observed flood extent. Blue outline shows flood outline from 10 m resolution model. Yellow outline shows flood outline from 50 m resolution model. Underlying DEM is 10 m resolution.

Figure 4. Box plots that show the variation between two batches of simulations reducing as the number of simulations in each batch increases. The top plot (**4a)** shows the difference between the means of the two batches, expressed as a percentage of the mean loss across all 500 simulations. The bottom plot (**4b**) shows the difference between the standard deviations of the two batches, expressed as a percentage of the standard deviations across all 500 simulations.

Figure 5. Exceedence probability plots produced by the catastrophe model when trained using the four different precipitation datasets. The grey shaded area denotes the 5-95% confidence intervals generated by the financial model. The losses simulated when the hydraulic and vulnerability modules are driven with observed flows for two historical events are shown for reference.

Figure 6. Top plot (**6a**) showing cumulative precipitation for each source. Bottom plot (**6b**) shows anomalies in >90th quantile precipitation intensities between gauge and other sources.

Figure 7. Exceedence probability plots produced by the catastrophe model when trained using the gauge record cropped to four different lengths. The grey shaded area denotes the 5-95% confidence intervals generated by the financial model. The losses simulated when the hydraulic and vulnerability models are driven with observed flows for two historical events are shown for reference.

Figure 8. Plots showing event hyetographs and hydrographs for the River Dodder (rows 1 and 2) and River Tolka (rows 3 and 4), and cumulative distribution function plots of modelled losses across the entire domain (row 5). The number of parameter sets simulating discharge at or above a given level at time *t* is represented by the hydrograph colour, ranging from all 100 (dark blue) to 1 (dark red). The weighted 5th - 95th quantile values from these plots are shown in table 2.

Figure 9. Exceedence probability plots produced by the model when trained using the four different precipitation datasets. The grey shaded area denotes the 5-95% confidence intervals generated by uncertainty within the vulnerability model.