Why extreme floods are more common than you might think

Floods in England and Wales have the potential to cause billions of pounds of damage. You might think such extreme events are rare, but they are likely to occur more frequently than expected. By ***Ross Towe****,* ***Jonathan Tawn*** *and* ***Rob Lamb***

In England, one-sixth of all households – 5.3 million properties – are considered to be at risk of flooding. In the winters of 2013-14 and 2015-16, more than 25 000 homes were flooded, causing an estimated £2.3-3.2 billion in economic damage. Little wonder, then, that flooding is considered the most severe natural hazard in the UK's national risk register (bit.ly/2yaBvUj).

Understanding flood risk is important for a number of parties, including the government, insurance companies and the general public. Flood events like those of previous years require significant response from the emergency services, and in some cases assistance from the army, to protect people and infrastructure against rising waters. However, each major flood event is different, as shown in Figure 1, and this makes preparation and planning a challenge.

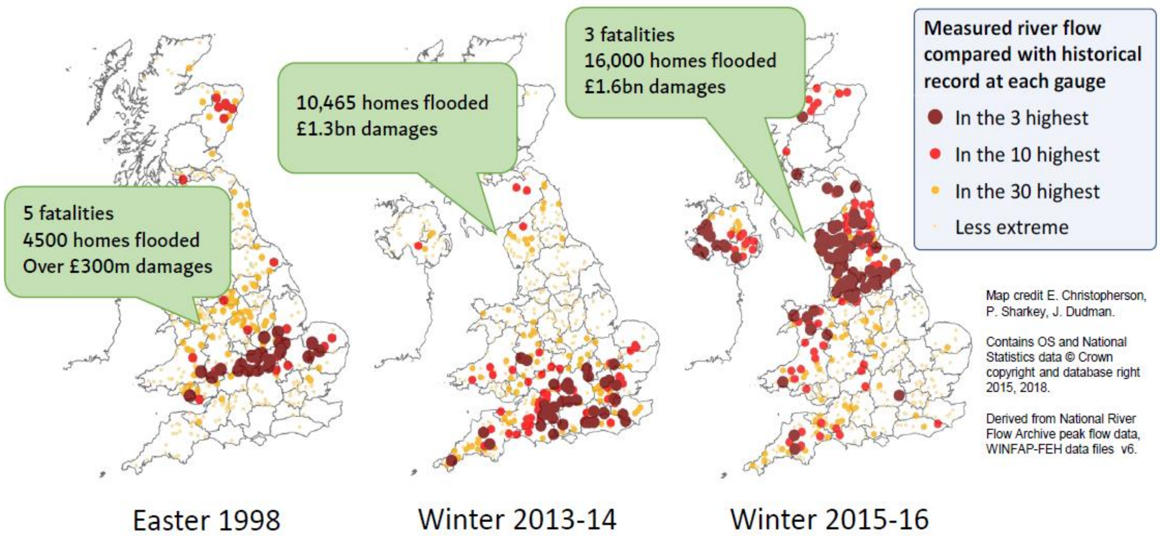

FIGURE 1 Extreme river flows mapped for three significant flood events, where discs represent the rank of the peak flow measured at river gauges. Ranks are relative to the historical records at each individual gauge. Internal borders show Local Resilience Forum boundaries.

In 2016, the UK government set up the National Flood Resilience Review (NFRR) to better understand the drivers of flooding, as well as whether existing methods to deal with the associated risks and damages would be robust over the 10 years following the review (bit.ly/2yrpJsw).

The risk of flooding has traditionally been thought of largely as a location-by-location problem, however the NFRR took a broader view. One of the questions asked was: How likely is extreme river flooding, when viewed at a national scale?

Answering that question required a comprehensive analysis of the historical extremes measured by river flow gauges and the relationship between gauge measurements in different locations. What we found was that extreme flooding events are more common than might be expected.

Creating a model

The severity of a hazard such as a flood is typically communicated through either the annual exceedance probability or the return period. For example, an event with a 100-year return period is, on average, likely to occur once every 100 years (assuming the underlying probability distribution remains unchanged over time) but has a 1% annual exceedance probability of occurring in any given year. An event of this kind might be classed as an "extreme event".

The risk of flooding is usually assessed through an integrated modelling approach composed of three key elements: hazard, vulnerability and consequences. Hazard corresponds to how rivers rise up following extreme rainfall events; vulnerability refers to the state of measures deployed to reduce the risk of flooding, such as flood defences; and consequences refers to the effect that these flood events can have on people and assets.

The aim of our research was to understand the behaviour of the hazard component. The work was conducted by the JBA Trust and statisticians at Lancaster University. Our study focussed on England and Wales, reflecting the scope of the NFRR, as flood risk management is a devolved matter in the UK, with separate arrangements in Scotland and Northern Ireland.

We considered 916 river flow gauges across England and Wales, each with at least 20 years of observational data.[1] Within hydrology, for a long time the standard approach was to analyse each gauge site in turn, fitting an extreme value model to gauged data and using it to estimate the appropriate return period for the design of flood defences. However, these traditional approaches fail to account for any relationships between gauge measurements in different locations or at different times. Statisticians refer to these relationships as spatial or temporal dependences.

Understanding spatial dependence is of utmost importance when large areas, such as England and Wales, are considered. Flooding is unlikely to be extreme simultaneously at all gauges but can occur in more than one place at a time. Furthermore, historical flooding shows us that as events become more extreme, they also become more localised. This behaviour can be seen in Figure 2. In this illustration, two gauges (both marked with triangles) have experienced a flood event; the remaining gauges, represented as discs, are coloured according to the probability that they too will record a flood event at some point over the course of a week. Note how the dark-coloured discs are spread widely when the T-year return level is lower, but become fewer and closer together as the T-year return level increases (see box for more detail).

Spatial dependence

To help us visualise the spatial dependence of river flow gauges, let X_i be the flow at gauge i in an event and $x_{i,T}$ be the T-year return level at that gauge. Conditional on being above the T-year level at gauge i, we want to determine the probability that the flow at gauge j, X_j, is also above its respective T-year return level, $x_{j,T}$. This is equivalent to being interested in the following probability

$$\chi_{j|i}(T)=P(X_j>x_{j,T}|X_i>x_{i,T}),$$

which provides us with an estimate of gauge j experiencing a flood event given that gauge i has experienced a flood event. To account for the time lag between events at different gauges, a window of a week is considered; this window is consistent with analysis of the lags over which dependence in extreme flows is apparent in large river basins in England and Wales.

The empirical conditional probability estimates in Figure 2 show that the estimates decay as we consider higher levels, i.e. larger T, and that this pattern of decay changes when different conditioning gauges are considered.

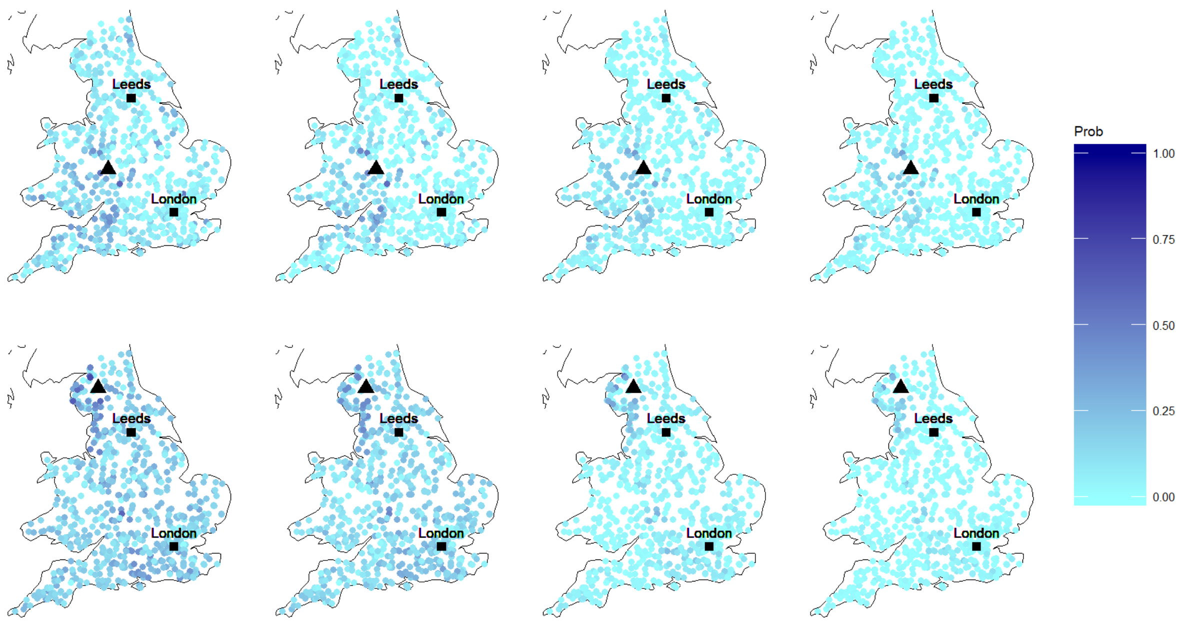

FIGURE 2 Comparisons of $\chi_{j|i}$ (T) for two conditioning gauges (represented by the triangle symbol), situated in the River Severn catchment (top row) and in Cumbria (bottom row). The first two columns are empirical estimates with T=5 months and 1 year respectively. The last two columns are estimates obtained from the statistical model for T=10 and 100 years respectively. The estimates show that conditioning on different locations changes the spatial extent of the estimates. When conditioning on the gauge in Cumbria, it is clear that distance is not the only factor; for example, there is stronger dependence with gauges on a

southern transect than with those gauges in North Wales, despite the latter being closer geographically to the conditioning gauge.

The first two columns of Figure 2 are empirical estimates based on historical data, but we cannot rely on empirical estimates when considering events more extreme than those recorded in the past. Instead we need a statistical model that has some underlying justification, and at this point we call on extreme value theory – which helps us understand the probability of events that are far removed from the average of all those we have seen before.

The statistical model needs to take account of both the probability of flooding at each location as well as the dependence between extreme events at different locations, and also missing data. Critically, we need the model to fit the extreme events best, and so focus the fit on these values.

The joint distribution is separated into a series of conditional distributions, where each gauge is taken in turn, with a statistical model estimated conditional on river flow at that gauge being large. We use an asymptotically motivated dependence model,[2] which is flexible enough to handle the property of weakening of dependence with level of extremity of the event that was identified empirically in Figure 2. The last two columns of Figure 2 show the model-based estimates. This figure also shows that our model captures this decay of dependence at higher levels in a way that is consistent with the empirical decay in $\chi_{j|i}(T)$.

The easiest method to estimate probabilities of interest is through a method known as Monte Carlo numerical integration. We simulated 10 000 years' worth of events, including events that are larger than those observed in the data for at least one site, but with the dependence structure of these events being consistent with the features of those observed in the historical record. This simulated set of events allows us to estimate summary statistics for a range of flood severities to help characterise the behaviour of flood events with minimal Monte Carlo error. (Uncertainties can also be obtained – although not shown here – through a computationally intensive process known as a non-parametric bootstrap.)

Questions and answers

One of the main questions we sought to answer was, "What is the probability, P, that a 1-in-100-year event occurs somewhere in England and Wales this year?" This question allows us to assess how likely flooding is when viewed at a national scale, and the simulated event set gives us the ability to answer this question for a range of return periods, not just a 100-year event.

If we had not attempted to model the dependence between gauges in a way that reflected the observed relationships over space and time, we would be left with two overly simplified ways to approach this question: either we assume complete dependence over all gauges, or assume that each gauge is independent of all others. These two approaches are compared in Figure 3. If complete dependence is assumed (red curve), then all of the sites behave the same; they all get their T-year event together. Therefore, under this scenario, the probability for a 1-in-T-year event is $P_D=1/T$ for any given T. In the case of independence, where there is no relationship between any of the river flow gauges (blue curve), the probability is $P_I=1-(1-1/T)^{916}$, where 916 is the number of river flow gauges used in our study.

Under the assumption of complete dependence, the probability of observing at least one 100-year event in any given year is 1%, but if we assume complete independence, the probability is approximately 100%. The huge difference between these estimates shows the importance of modelling the relationship between river flow gauges based on observed dependencies.

For the modelled approach (the black curve in Figure 3), the probability of observing a 1-in-100-year event at any of the 916 gauges in any given year is 78%; for a 1-in-1000-year event, the probability is 20%. Therefore, it is likely that a 100-year extreme flow event occurs somewhere on the rivers of England and Wales in any given year, and more extreme events still have a non-negligible probability.

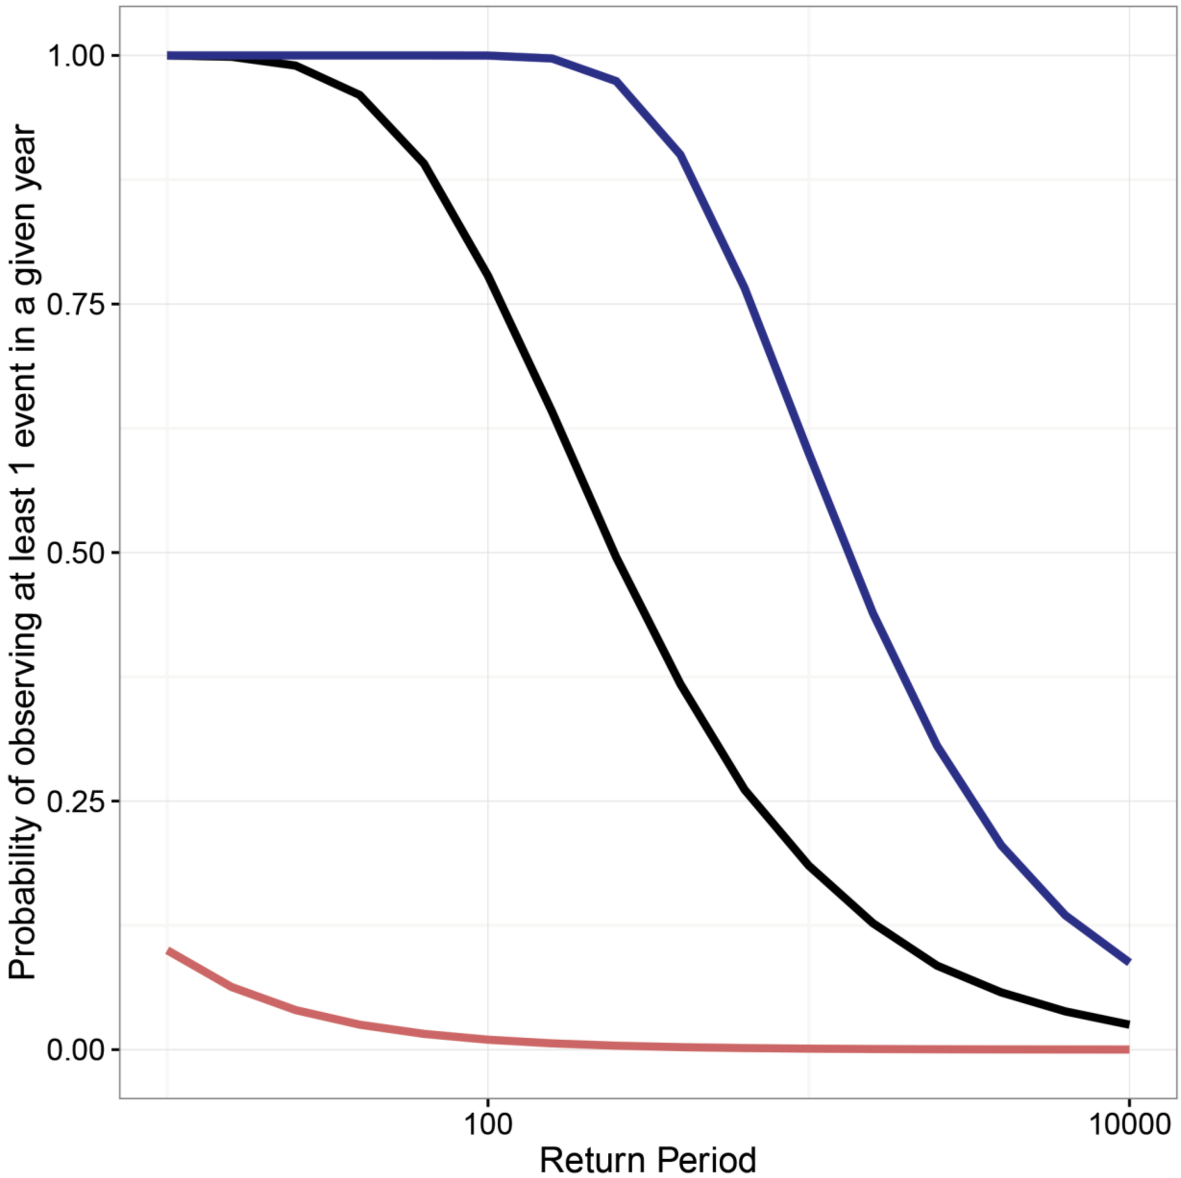

FIGURE 3 Comparison of the three dependence models used to estimate the probability of observing extremely high rates of water, above a T-year return period, flowing through at least one river gauge in a given year. The black curve gives estimates of our model for the dependence between river gauges, inferred from observations. The red curve uses a false assumption of complete dependence, while the blue curve uses a false assumption of independence.

The model applied here can be used to assess the probability of flooding for any reasonable local standard of protection. The values T = 100 and T = 1000 are of practical importance because they are used in defining flood risk zones in official flood maps.

Another question we might ask is, "On average, in a given flood event, how many gauges exceed a 10-year return period if at least one of them observes a 10-year return period?" Figure 4 answers this question, providing an understanding of the spatial extent of flood events and how it changes for a range of return-period events. For a typical event with at least one gauge station above a 10-year return period, we might expect 11 gauges in total to experience a return period higher than 10 years. For a 1-in-10 000-year event, the average number of gauges experiencing a return period higher than 10 000 years is 3.

Interest also lies in understanding how many Local Resilience Forums (LRFs) are affected by a single extreme event. LRFs are administrative regions, generally coinciding with police areas, in which multi-agency responses to natural hazards are organised; within England and Wales there are 42 LRFs, and their boundaries are shown in the maps in Figure 1. The estimates for at least r regions being above a T-year return period in any given year is shown in Figure 4 (right), with the r=1 case consistent with Figure 3. In any given year there is a 35% probability of a 1-in-100-year event occurring in at least four LRFs.

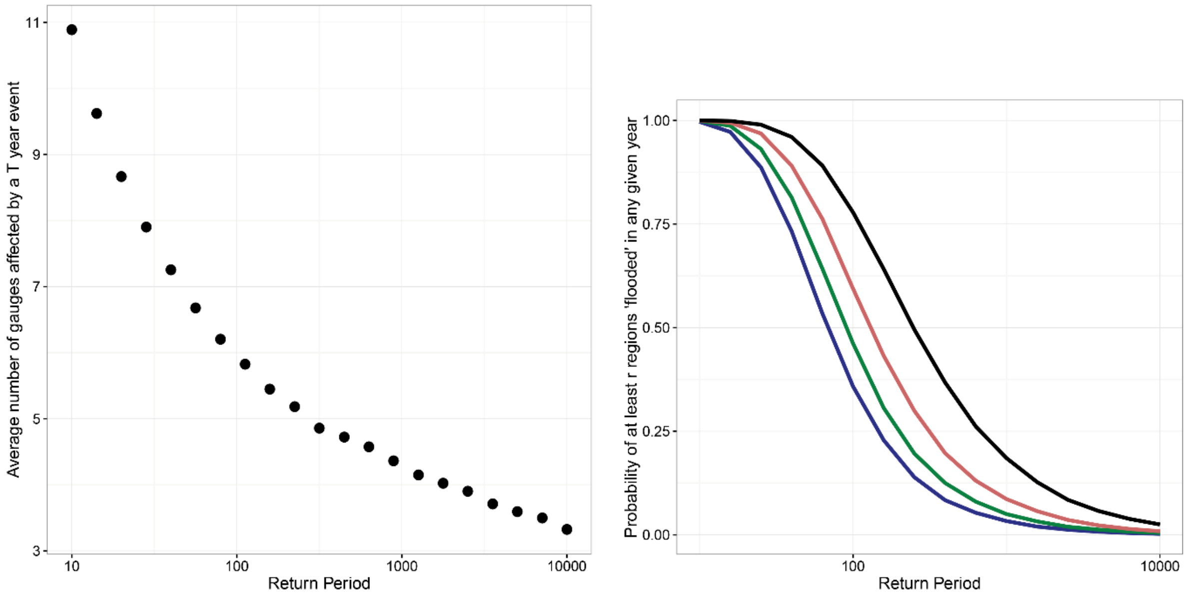

FIGURE 4 (left) The average number of gauges in a given event affected by a T year return period given that at least one of the gauges experienced an event with a T year return period. (right) Probability of observing at least r LRF regions above a T-year return period in any given year. The black, red, green and blue curves show the cases for when r = 1,2,3,4 respectively.

Simulations for planning

Every two years, the UK government carries out a national assessment of possible civil emergencies (such as extreme weather, accidents, disease and other risks). This National Risk Assessment (bit.ly/2RDL2zE) includes the development of scenarios that are intended to represent challenging, but not implausible, extreme events that could occur over the next five years. The scenarios form the basis for some of the emergency response planning that is carried out by various agencies to prepare for civil contingencies.

The modelling carried out to understand the probability of widespread flooding is used to develop these planning scenarios. Events simulated by the statistical model were used as the basis for the scenario development, ensuring that proposed scenarios would be drawn from an analysis that is consistent with observed patterns of river flow extremes, and based on a theoretically sound approach to extrapolate beyond the historical data. These two features allowed the generated scenarios to meet the requirement of being challenging, but not implausibly so. Here, a “challenging” event is one that would stretch emergency services and other agencies in terms of their capacity to respond. To understand the extent of that challenge, scenarios generated using the statistical model have to be translated into the real impacts on people and society, should they ever occur.

This translation requires a combination of complex and detailed models to predict the areas that would be flooded, given a particular event generated by the statistical modelling, and what the consequences of that flooding might be. The first step is to convert statistically-generated data about extreme river flows into predictions of the depth and speed of water flowing over floodplains. Based on these detailed flood predictions, the impacts of the scenario are captured in terms of 312 different impact metrics, covering the main categories of flood damages. The steps of this process are illustrated in Figure 5.

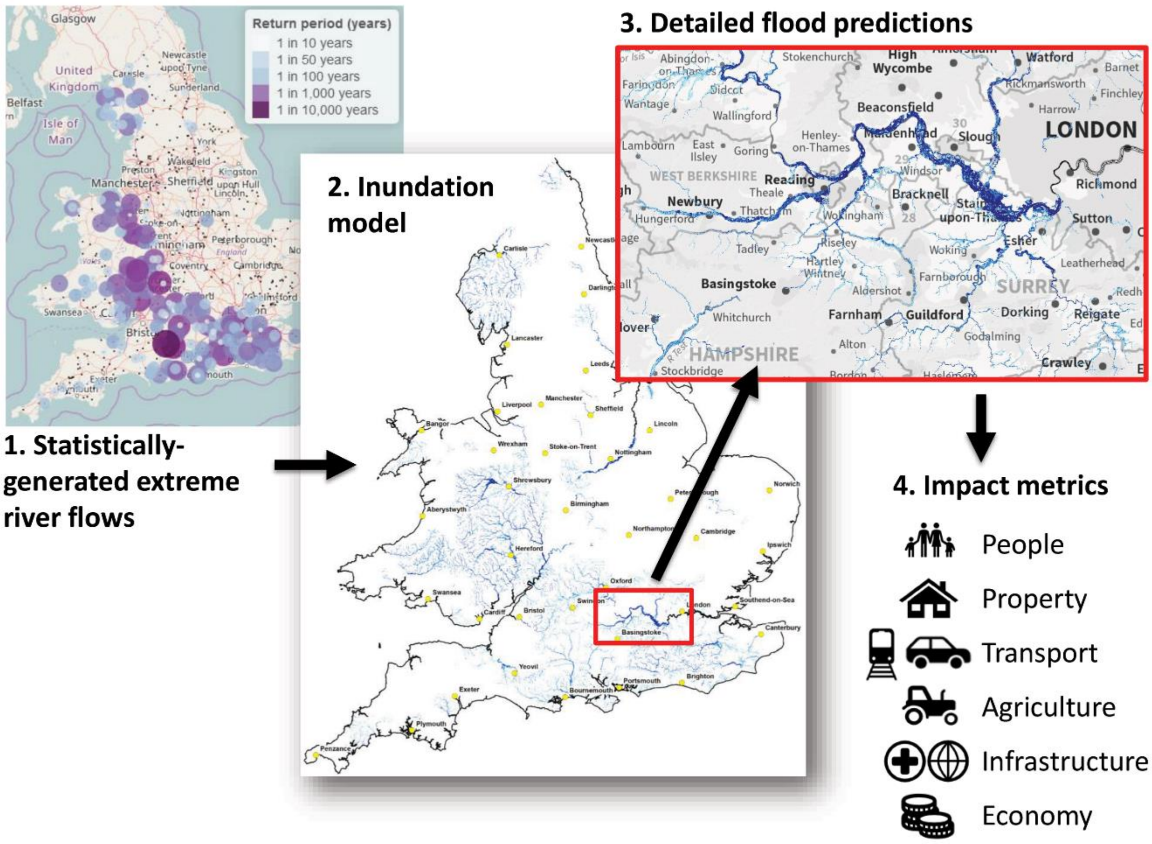

FIGURE 5 Process for translation of a statistically-modelled extreme river flow scenario (step 1) through numerical flood inundation models (step 2) to derive detailed flood scenario maps (step 3) and impacts (step 4).

What comes next?

Our analysis has shown that statistical models can help us understand widespread flood risk as well as aid planning for future scenarios. The statistical modelling of measurements of extreme river flow allows us to provide robust answers to national-scale questions by capturing the complex dependence structure of river flows across a large number of measurement gauges.

However, there is more work to be done. The analysis considered only locations where there are river flow gauges, but there is interest in obtaining estimates of flood risk for anywhere along the river network. Current research is addressing how we can utilise information about rainfall and the structure of the river network to inform this analysis and through the development of spatial extreme value methods.[1]

Ultimately, a comprehensive risk analysis should consider the full distribution of the impacts of flooding, both in terms of the statistical distribution of different impact measures, and the geographical patterns. Ideally, it should include all types of flooding, where they can occur, including flooding from rivers, the sea, urban drainage and groundwater. The methods we have shown for extreme river flows are an important step towards these goals. In particular, the ability to capture complex types of dependence among large numbers of variables is crucial to developing robust risk models.

Communication of risk needs to be informed by this kind of realistic analysis of flood events, which recognises their differing geographical patterns and provides assessments of the probability of events over space and time that are relevant for different decision-makers. For example, engineers designing a flood wall may still need to rely on the analysis of the return period of a flood level in a specific place. But planners, insurers and government policy-makers can benefit from analysis of scenarios that include both locally extreme and widespread flooding.

The answers to the questions posed by the NFRR highlighted that even extreme flooding events are more common than one might expect when looked at from a national perspective.

Author bios

Ross Towe is a Senior Research Associate on the Ensemble Project at Lancaster University.

Jonathan Tawn is a Distinguished Professor of Statistics at Lancaster University and is Director of the STOR-i Doctoral Training Centre, he won the Barnett Award in 2015.

Rob Lamb is Chief Scientist at JBA Group and Director of the JBA Trust, he is a Professor in Practice at Lancaster University.

References

1. Tawn, J. A., Shooter, R., Towe, R. P., & Lamb, R. (2018). Modelling spatial extreme events with environmental applications. *To appear in Spatial Statistics*.

2. Heffernan, J. E. & Tawn, J. A. (2004). A conditional approach for multivariate extreme values (with discussion). *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, *66*(3), 497-546.