18th World IMACS / MODSIM Congress, Cairns, Australia 13-17 July 2009
http://mssanz.org.au/modsim09

Outline of a framework for systematic decision analysis in flood risk management

H. Harvey[1], J. W. Hall[1] and R. Peppé[1]

[1] *School of Civil Engineering and Geosciences, Newcastle University, United Kingdom*
Email: hamish@hamishharvey.com

Abstract: Computational models have long been used to support decision making in water management. With growth in the availability of data and plummeting cost of computation, the way these models are deployed is changing. Single simulations for "best estimate" conditions are giving way to the use of multi-run computational experiments as exemplified by risk, uncertainty and sensitivity analysis. In this paper we look beyond these techniques to a new generation of tools. Building on simulation and risk analysis, these tools will enable systematic, quantitative *decision analysis*, allowing the merits of sets of options to be examined and compared, according to a variety of metrics and under a range of possible future conditions.

We outline a conceptual framework for decision analysis for strategic planning. The life cycle of flood defence assets lasts 100 to 150 years, and option appraisal must encompass this cycle. The many processes that will influence the state of the system on this time scale must be explicitly represented, allowing for the considerable uncertainties involved.

The approach described is novel in a number of respects, among which the following stand out.

- Processes of long term change, including the implementation of management interventions, are simulated. System states at future times are calculated conditional upon a simulated future, not manually assembled.
- The costing of interventions and options (timed sets of interventions) is integrated into the framework. Costs can be a function of system state at the time of implementation of the intervention, and can therefore be influenced by long term change processes and uncertainty.
- The assessment of option benefits and costs is integrated and the results combined into measures of option performance. This makes it possible to ensure that performance estimates are based on internally consistent scenarios and account for dependence between variables influencing costs and benefits.
- Computational experiments such as uncertainty and sensitivity analysis are applied to the overall performance analysis rather than its component parts, leading to information of greater value to decision makers.

We describe an implementation of this framework in the context of a simple flood risk management situation. The performance metric used is Net Present Value of reduction in Expected Annual Damage. Processes of long term change include increasing relative mean sea level and increasing damage potential resulting from economic growth. In both cases the rates of increase are uncertain, as are the costs of implementing the interventions from which options are constructed. These various uncertainties are expressed as probability density functions (PDFs) over the variables in question. They are propagated through to option performance by means of a Monte Carlo experiment, allowing a PDF on performance to be constructed for each option (figure 1). A range of additional diagnostic outputs can also be generated.

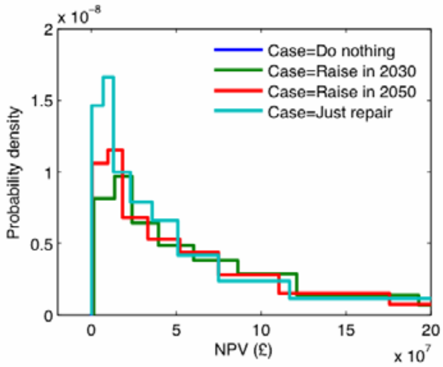

Figure 1 Probability density over Net Present Value of reduction in Expected Annual Damage for each option

Keywords: *Decision Analysis, Flood Risk Analysis, Long Term Change, Climate Change*

1. INTRODUCTION

The planning and implementation of strategic flood risk management interventions can take decades, and the investments involved must be justified using cost benefit analysis over an appraisal period that encompasses the design life of the resulting infrastructure. This leads to the desire to evaluate management options in terms of their return on investment over periods of a century or more. Climate, the economy and demography are but some of the more obvious sources of change that will influence estimates of flood risk on this time scale.

At the same time, ever more data are available to flood risk managers, at least in countries like the UK. Computing resources are similarly ever more abundant. The question then arises of how this “embarrassment of riches” might be converted into useful knowledge relevant to the problem of strategic decision making.

Some steps in this direction have already been taken. “Risk analysis” – which involves the estimation of the statistical expectation of annual damage, EAD, using a damage model and a joint probability density function over the inputs to that model – is now a standard part of the flood risk management toolbox in the UK (Hall et al., 2003; Dawson, 2003; Hall et al., 2005). Uncertainty propagation and sensitivity analysis methods have been applied to explore the influence of uncertainty on estimates of EAD (Environment Agency, 2009).

In this paper we argue that these developments represent first steps in a process, the logical conclusion of which is *systematic decision analysis*. This goes beyond the status quo of modelling and risk analysis in a number of respects, the most important of which are itemised in the abstract. We present a simple but general conceptual framework for decision analysis and illustrate its application to a simplified but representative strategic flood risk management decision.

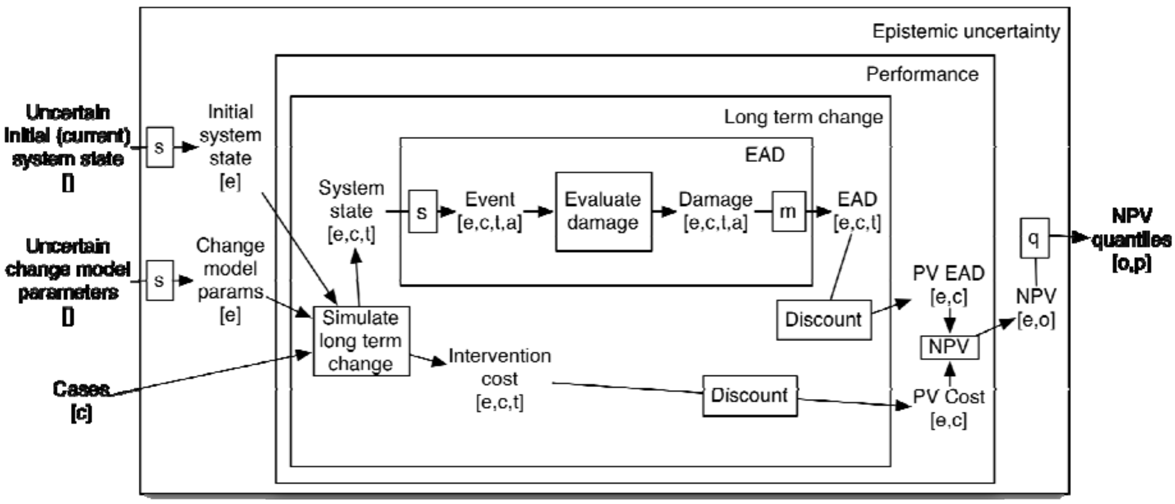

Figure 2 Overview of a conceptual framework for decision analysis. Larger boxes indicate the nested structure of the analysis. The computation itself follows the directed bipartite graph flowing from left to right. Unboxed nodes are data sets, while boxed nodes are transformations, where *s* stands for “sample”, *m* for “mean” and *q* for “estimate quantiles”. The dimensions of data sets are indicated by characters in square brackets beneath; these are *epistemic uncertainty* (e) *cases* (c), *options* (o), *time* (t), *aleatory uncertainty* (a), *cumulative probability level* (p).

2. CONCEPTUAL FRAMEWORK

Figure 2 sets out the conceptual framework. It consists of five nested layers of analysis, as indicated by the nested boxes, each of which is labelled in the upper right hand corner. One way of thinking about the analysis is to consider these boxes as representing the nested loops of a naïve, sequential implementation.

1. *Evaluate damage* using a deterministic model of the damage caused by a specific hydrological event.
2. *Estimate Expected Annual Damage* (EAD), commonly referred to as “risk”. This involves many invocations of the layer 1 damage estimator.
3. *Simulate long-term change*, generating system states and costs. For each state, estimate EAD by invoking level 2. A possible future is specified as an initial system state, an option (a timed

sequence of flood risk management interventions) and a set of parameters to a model of relevant long-term time-dependent processes (sea level rise, demographics, …).

4. *Evaluate option performance.* The layer 3 computation is conducted for each of a set of *cases*, including a "do nothing" base case and a number of "do something" *options*. For each option, a performance metric is evaluated by comparison of the base case EAD with the residual EAD associated with each option.
5. *Apply computational experiments to performance estimator*. Computational experiments can be used to explore the behaviour of the layer 4 performance estimator.

Layers 1 – 4 define a deterministic estimator of option performance. The integration of risk analysis and cost estimation in the context of simulation of long term change processes is novel and powerful. It is possible in this framework to ensure that costs and benefits are assessed coherently, for example in how they are influenced by uncertainty about future economic growth.

Given this estimator, some familiar techniques and algorithms take on a new significance. Parameter uncertainty can be propagated all the way through to the performance metrics of direct interest to decision makers, as described in section 3. Sensitivity analysis can be used to rank input parameters in terms of their contribution to uncertainty on performance of options. Thus we can explore the influence of uncertainty on our ability to rank options; our ability to *make a decision*.

We can also design experiments to answer less familiar questions. For example, Hall and Harvey (2009) present an experimental Info-Gap (Ben-Haim, 2006) *robustness analysis* structured according to the framework described here. An option has the property of robustness if its performance does not deteriorate rapidly as conditions deviate from those of the "best estimate". When a decision must be made under considerable uncertainty, a robust option may be preferred to one that is less robust even if the latter has considerably better performance under best estimate conditions.

The fact that options are represented explicitly and separately from other aspects of long term change opens the possibility of using search and optimisation algorithms to explore the space of possible options.

3. EXAMPLE MODEL

3.1. Flooding system and damage estimation

The decision analysis framework assumes the availability of an implementation in software of a damage estimating function $d(\mathbf{s},\mathbf{e})$, where the vector $\mathbf{s}$ is a description of system state and $\mathbf{e}$ of an event. This function d will usually be an integration of several component models. In the case of flood risk analysis, these include models of river hydraulics, defence reliability, flood inundation and damage generation.

Table 1 System state variables

Variable	**D**	**R**	**C**
Parameters to max sea level distribution		*	
Defence crest level	*		*
Ground level at defence	*		
Defence condition	*		*
Damage potential	*		
Upgrade cost factor			*
Repair cost factor			*

For the purposes of demonstration a simplistic damage model is used in this example, but one that nonetheless implements the basic structure of the Risk Assessment for Strategic Planning methodology used by the UK Environment Agency (see e.g. Hall et al., 2003; Gouldby et al., 2008). A hypothetical situation is considered, in which an area of low lying coastal land is separated from the sea by a single flood defence. Tidal events are defined by the peak surge tide water level, with a standard local hydrograph shape being assumed. An event may result in inundation by overtopping the defence or if the defence is breached. In either case it is assumed that the defence acts as a weir, so the total volume discharged into the floodplain during the tidal cycle can be estimated using the weir equation. Breach dimensions are estimated using simple rules. The volume discharged is used to estimate a water depth in the floodplain, from which damage is estimated from depth using a depth-damage curve capturing information about the type and density of buildings in the floodplain (Penning-Rowsell et al., 2005). The probability of breaching is estimated conditional on event maximum water level and used to find en event expected damage by taking a probability-weighted sum of damages associated with the breached and overtopped cases.

In the case of this simple example, the system state vector $\mathbf{s}$ contains the data indicated in table 1. The variables indicated in the column headed “D” are used in damage estimation.

3.2. Risk analysis

In the example analysis, an *event* is a tidal cycle, defined by the maximum water level reached at the defence, a standard hydrograph shape being assumed. The probability density function over annual maximum water level $\mathbf{e}$ is given by $p(\mathbf{s},\mathbf{e})$. The parameters to this distribution are also part of the system state vector (entries marked in column “R” in table 1).

Given this function and an assumption that only the damage generated by the most extreme tide in a given year is significant, the Expected Annual Damage (EAD) or “risk” $r(\mathbf{s})$ is given by equation 1, which also indicates that the risk can be estimated using Monte Carlo integration, where $\mathbf{e}_i$ is the *i*th member of an *n*-member sample from the distribution $p(\mathbf{s},\mathbf{e})$.

$$r(\mathbf{s}) = \int d(\mathbf{s},\mathbf{e})p(\mathbf{s},\mathbf{e})\,d\mathbf{e} \approx \frac{1}{N}\sum_{i=1}^{n} d(\mathbf{s},\mathbf{e}_i) \tag{1}$$

3.3. Option performance

We evaluate options according to the Net Present Value (NPV) of the reduction in EAD they generate relative to a “do nothing” base case as a performance measure. Allowing that the system state will change in time, we wish to compare options consisting of sequences of intentional modifications to the system. Considering an appraisal period of *m* years, the NPV of case *c*, g_c, is given by equation 2 in which $s_{c,y}$ is the system state in year *y* for case *c* and *d* is the discount factor. Case *c=0* is the base case.

$$g_c = \sum_{y=0}^{m-1} \frac{1}{(1+d)^y}\Big(r(s_{c,y}) - r(s_{0,y}) - c(s_{c,y})\Big) \tag{2}$$

4. LONG TERM CHANGE

In order to implement the analysis described in the previous section it is necessary to generate the system state vectors $s_{c,y}$. Even in a deterministic analysis taking no account of uncertainty there could be a large number of these. Constructing them by manual modification of a base system description is time consuming, error prone and cannot readily be repeated. One of the novelties of the approach described here is the generation of these state vectors by simulating the evolution of system state through appraisal time. In addition to the advantages noted, this approach means that the model of change captures explicitly and formally the assumptions about long-term change made in an analysis.

We separate the sources of changes to system state according to whether they are the result of intentional intervention in the interests of flood risk management or of processes, which we refer to as *exogenous*, over which the management system exerts no control. We consider the latter in this section. Changes resulting from intervention are discussed in the next section.

Table 2 Long-term change model

Variable	State variables influenced	Change model
Rate of sea level rise	Sea level distrib-ution parameters	Constant rate
Rate of economic growth	Damage potential	Constant rate, compound
Rate of deterioration	Defence condition	Constant rate, bounded

In the context of strategic flood risk management, typical exogenous change processes include increasing relative mean sea level, changes in the frequency of extreme rainfall, subsidence leading to reducing defence crest levels, defence deterioration, economic growth (or otherwise) as it affects the value of assets at risk from flooding, and land use change including population and demographic change.

Of these, in the example we consider sea level rise, defence deterioration and economic growth. **Error! Reference source not found.** presents the parameters to the overall model of long-term change, the model of change controlled by each parameter and the system state variables which are influenced by these change processes. Mean sea level is assumed to increase at a constant annual rate. Defence condition is indicated by a condition grade lying between 1 (perfect) and 5, and deterioration is defined in terms of the number of years it takes for the defence to drop a condition grade. Economic growth, which is assumed to lead to an increase in damage potential as the residents of the floodplain accumulate more valuable possessions, is taken to compound annually at a constant rate. The uncertainty in these assumptions can be analysed in the outer-most level of analysis.

5. INTERVENTIONS AND OPTIONS

In addition to exogenous change we consider *interventions*, changes made to the system intentionally. Examples of interventions for flood risk management include building new defences or raising crest levels of existing defences, repair and maintenance of infrastructure, construction of a tide-excluding barrier or flow-limiting barrage. These are examples of "hard engineering" interventions, designed to reduce the probability of flooding. Other interventions focus on reducing the impact of flooding that does occur: flood proofing of residential property, creation and enforcement of planning regulations and so on.

Table 3 Specification of cases as combinations of defence repairs (R) and upgrading of crest level (U)

Case	2010	2030	2050	2070	2090	2110
0						
1	R	U(5.5)	R	R	R	
2	R	R	R	U(5.5)	R	
3	R	R	R	R	R	

It is the purpose of the decision making process to generate a range of *options*, to evaluate these, and ultimately to select one for implementation. Options are sets of interventions, each applied at a particular point in time. At present options are specified manually, but, because they are explicitly represented, they could also be generated in software. Future development will explore the use of rules to generate options during simulation of long term change, for example allowing defence crest levels to be chosen within particular runs based on observed rates of sea level rise.

Interventions are modelled as pairs of functions. The first takes a system state vector describing the state of the system before the intervention is applied, and returns a state vector describing that system modified by the application of the intervention. The second takes the pre-intervention state vector and returns the cost of implementing the intervention in the context of that state. Intervention costs can in this way depend on any system state variables. These may denote physical characteristics of the system, such as the existing crest level of a defence being raised, or parameters to a cost model encapsulating the costs of materials and labour. Variables in table 1 marked "C" are used in intervention costing.

The example makes use of two intervention types. These are both applied to the single flood defence in the system. In both cases, the constant of proportionality in the cost model is part of the system state, and can be subject to long term change and uncertainty.

- Repair ("R") returns defence condition to its maximum (1), representing the effects of remediation works. The cost of repairing a defence is assumed to be proportional to the difference between maximum condition and actual condition prior to repair.
- Raise crest level (upgrade) by x m ("U(x)"). Note that this provides a parameterised family of interventions. It is assumed that the work involved in raising crest level brings the defence up to maximum condition as a side effect. The cost of raising defence crest level is taken as proportional to the square of the difference between new and old crest levels.

The cases considered in the example are shown in table 3. In the example, interventions are applied on decadal boundaries. The appraisal period runs from 2010 to 2110. Case 0 is a "do nothing" base case, while cases 1–3 are the options we wish to consider.

6. EXAMPLE EXPERIMENT: UNCERTAINTY ANALYSIS

A traditional approach to engineering decision making would assume stationary climate and vulnerability and compute Net Present Value based on best estimate values for state variables. This can readily be recreated in terms of the estimator set out above by setting the rates of change for the exogenous change model to 0. Having implemented the analysis steps described above in a software device in such a way that the calculation of NPV for a set of options can be conducted from input data fully automatically (that is, without further manual intervention), however, it becomes possible to subject this device to all manner of computational experiments.

Such experiments have the general form that some, possibly large, set of input parameter vectors are generated, performance of each option is estimated for each member of that set, and the results are reduced to summary statistics. These are then usually presented graphically. For this example, a simple Monte Carlo uncertainty propagation experiment was applied. We assume that input uncertainties are specified as probability density functions (PDFs), which are sampled and propagated. We can then estimate a PDF on performance of each option, indicating the quality of our estimate.

7. IMPLEMENTATION

A naïve implementation in a sequential programming language of the decision analysis calculation described would translate the layers described in section 2 into nested loops. Only as many intermediate results would be stored as necessary to complete the generation of the final result. We take a different approach. Our implementation is based on the Reframe metamodel (Harvey et al., 2008), wherein analysis is specified as a sequence of transformations of multi-dimensional data sets with named dimensions. From this perspective, the calculation proceeds as indicated in the directed bipartite graph passing from left to right in figure 2.

The labels of data sets and transformations are specific to the example described above, particularly in respect of the performance metric (NPV) and associated intermediate data sets. The conceptual framework is general, however. Our aim is to construct a software tool that allows a variety of performance metrics to be defined, including the vector-valued metrics required for multi-criteria analysis.

The overall calculation is divided into three distinct phases:

1. Based on the experiment specification, the inputs for each required run of the underlying damage model are generated. These inputs are stored in a multi-dimensional array.
2. Each element of this run specification is then fed through the damage model, the results being accumulated into another multi-dimensional array sharing dimensions with the input array.
3. Inputs and results are processed to generate summary measures and visualisations.

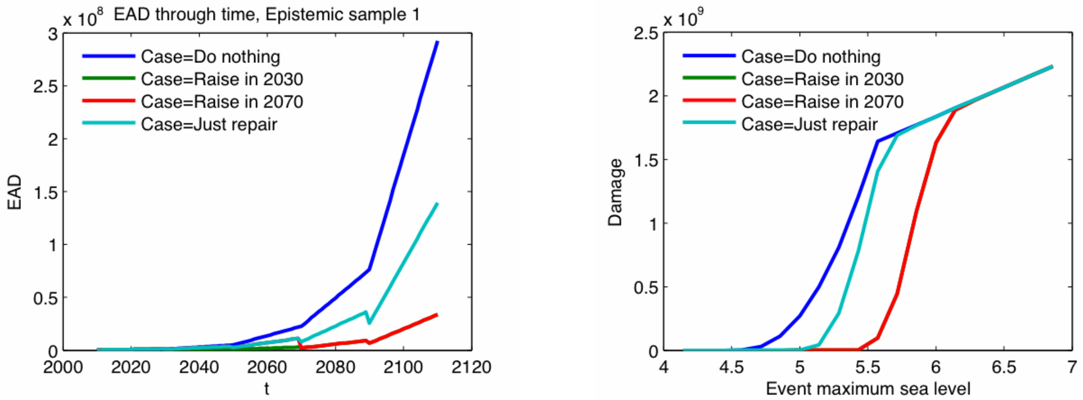

(a) (b)

Figure 3 Examples of additional outputs possible after damage model runs have completed

8. RESULTS AND VISUALISATION

This arrangement has a number of advantages over the naïve approach. It enables distributed computing facilities to be used for the execution of the damage model, which characteristically represents by far the largest portion of the analysis CPU time. It also means that a complete set of inputs and outputs are available for the analyst and other parties to explore, analyse, and visualise as they see fit.

The example experiment was designed to estimate, given a joint probability density function over the input parameters, a distribution over performance for each of a number of options. This is a simple form of uncertainty propagation. Figure 1 shows the primary result, a probability density function over Net Present Value for each option.

When conducting this kind of analysis it is not sufficient to compute and present the primary result. A variety of visualisations of intermediate steps in the analysis are required if the behaviour of the model is to be understood and trusted by its developers and by stakeholders in the decision it is to inform. These may be required at any time, including well after the selected option has been implemented, for example if the quality of the modelling on which the decision was based is challenged in legal proceedings. In this situation it is likely that nothing short of access to the results of individual simulations will be acceptable.

By way of example, two useful auxiliary visualisations are presented in figure 3. Subfigure (a) shows the development of EAD through time for all cases (base case and options) for a given sample of epistemic uncertainty. Subfigure (b) displays the variation of damage with event maximum sea level for a particular system state. Such views can be extremely useful in understanding features of aggregated results.

9. CONCLUSIONS

We have introduced the concept of quantitative decision analysis as the logical extension of the recent implementation of risk analysis in flood risk management. A conceptual framework is proposed which captures the essential structure of such analyses in a general form. The conceptual framework defines five layers of analysis. The first four – *damage simulation*, *risk analysis*, *simulation of long term change including management intervention* and *performance assessment* – together constitute a deterministic model of option performance. The fifth layer involves applying some form of *computational experiment* to this model.

It is anticipated that a given performance model may be subjected to numerous such experiments, each providing modellers or decision makers with different information. Some possible experiments include uncertainty analysis, sensitivity analysis, search (for example for good options) and optimisation, and robustness analysis (Hall and Harvey, 2009). The value of uncertainty and sensitivity analysis are considerably increased when applied to a fully integrated performance analysis than to one or another component of that analysis. Robustness analysis generates new information of a kind previously unavailable to decision makers, allowing them to examine how flood risk management options will perform if future conditions deviate from current best estimate predictions.

An example analysis has been implemented in terms of this framework as a first stage in proving its usefulness and generality. A prototype web-based user interface is being developed which will allow the user to specify and run experiments on this example and explore and visualise the results from these experiments.

The framework has been designed to allow the implementation of a generic software tool: a model- and even domain-independent decision analysis tool which assists with the process of assembling models of physical processes and their social impacts, representations of interventions as modifications to such models, cost models for such interventions, models of long term change, and management options. This generality will be realised by iterative refinement from the model-specific prototype.

ACKNOWLEDGEMENTS

This work was funded by the UK Engineering and Physical Sciences Research Council grant EP/F020511/1.

REFERENCES

Ben-Haim, Y. (2006), *Info-Gap Decision Theory: Decisions Under Severe Uncertainty*, 2nd ed., Academic Press.

Dawson, R. (2003), Performance-Based Management of Flood Defence Systems, PhD Thesis, University of Bristol.

Environment Agency (2009), *TE2100 Phase 3(i) Topic 2.2 Task 2 Technical Working Paper 2: Uncertainty and sensitivity analysis of risk calculations*, Environment Agency.

Gouldby, B., P. Sayers, J. Mulet-Marti, M.A.A.M. Hassan, and D. Benwell (2008), A methodology for regional-scale flood risk assessment, *Proceedings of the Institution of Civil Engineers Water Management* 161(WM3), 169-182, doi:10.1680/wama.2008.161.3.169.

Hall, J.W., R.J. Dawson, P.B. Sayers, C. Rosu, J.B. Chatterton, and R. Deakin (2003), A methodology for national-scale flood risk assessment, *Proceedings of ICE, Water & Maritime Engineering* 156(3), 235–247, doi:10.1680/maen.156.3.235.37976.

Hall, J.W., and H. Harvey (2009), Decision making under severe uncertainties for flood risk management: a case study of info-gap robustness analysis, In Proc. 8th International Conference on Hydroinformatics.

Hall, J.W., P.B. Sayers, and R.J. Dawson (2005), National-scale Assessment of Current and Future Flood Risk in England and Wales, *Natural Hazards* 36(1-2), 147–164, doi:10.1007/s11069-004-4546-7.

Harvey, H., R. Peppé, and J. Hall (2008), Reframe: a software system supporting flood risk analysis, *Journal of River Basin Management* 6(2), 163-174.

Penning-Rowsell, E., C. Johnson, S. Tunstall, S. Tapsell, J. Morris, J. Chatterton, and C. Green (2005), *The Benefits of Flood and Coastal Risk Management: A Manual of Assessment Techniques*, Middlesex University Flood Hazard Research Centre.