Oxford to Cambridge Arc flood risk investment study

Adaptive approaches and optimisation technical report

Revision 3

25 April 2022

Environment Agency

Oxford to Cambridge Arc flood risk investment study

Project No:	B550C033
Document Title:	Adaptive approaches and optimisation technical report
Document No.:	B550C033/04
Revision:	Inception Phase
Document Status:	Final
Date:	25 April 2022
Client Name:	Environment Agency
Client No:	Client Reference
Project Manager:	Daniel Boyd
Author:	David Cobby, Joe Clarke
File Name:	OxCam economics study - adaptive approaches and optimisation technical report

Jacobs

The West Wing
1 Glass Wharf
Bristol, BS2 0EL
United Kingdom
T +44 (0)117 457 2500
www.jacobs.com

Document history and status

Revision	Date	Description	Author	Checked	Reviewed	Approved
1	27/11/2020	Inception Phase	D Cobby J Clarke	A Yeowell	A Yeowell	D Boyd
2	16/03/2022	Draft final	J Clarke	A Yeowell	A Yeowell	D Boyd
3	25/04/2022	Final	J Clarke	D Boyd	D Boyd	D Boyd

Contents

1. Introduction

The OxCam economic evidence study for investment in flood resilience and adaptation aims to find the economic optimum **level** and **timing** of investment in flood resilience across the Oxford-Cambridge (OxCam) Arc. Our high-level method is to use existing river level data and new 2D floodplain modelling combined with an assessment of impacts to represent the effect of a range of interventions over time under a series of climate change and development scenarios.

This technical report describes our approach to optimising investment options in the context of adaptive approaches, and it is one of three technical reports detailing elements of our approach. As shown in Figure 1.1, these technical reports are accompanied by a cross-cutting summary report which describes the overarching approach and findings.

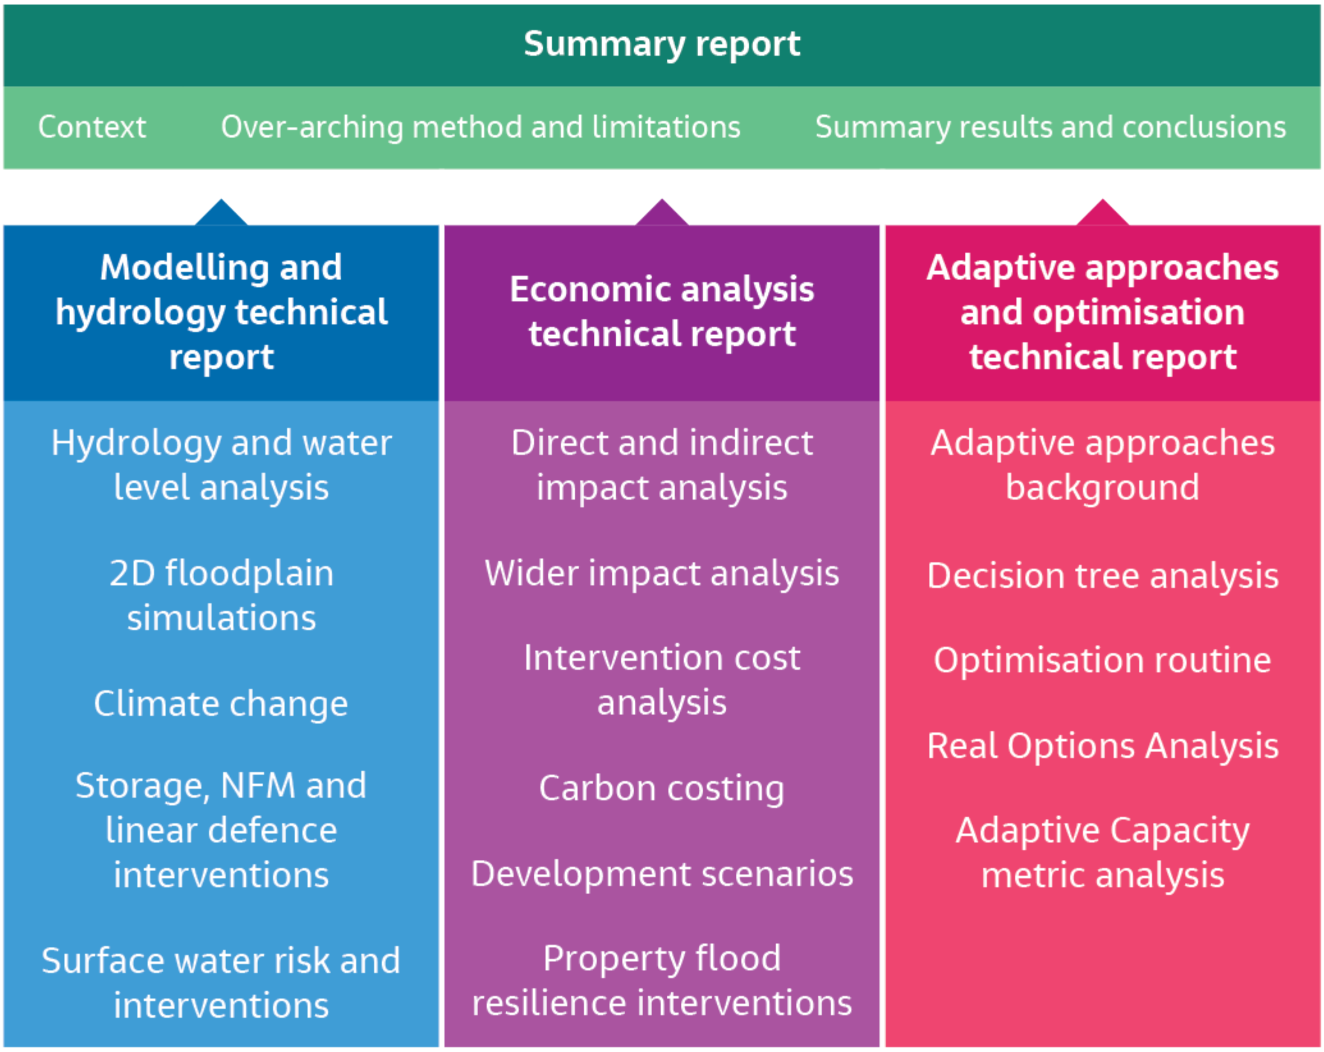

Figure 1.1. OxCam flood economics project reporting structure and content.

2. Background

2.1 Introduction

This section provides context to defining an approach to explore adaptive approaches in the context of finding the optimum level and timing of investment in flood risk and resilience interventions for the OxCam Arc. It explores adaptive approaches concepts, and background to the economic optimisation carried out by the *long-term investment scenarios* projects.

2.2 Adaptive approaches

A precautionary approach to managing flood risk seeks to build schemes or implement policies early, providing benefits over their useful life. In some cases, these can become unnecessarily conservative (e.g. large, high, expensive structures etc) creating unnecessary community concern about the implications of climate change and the urgency of response. This can lead to flood management activities being seen as a battle against nature which, in the past, has resulted in language such as "hold the line" and "flood defence". In economic terms, a precautionary approach may be more or less cost beneficial depending on how drivers of risk actually change in the future, so is dependent on factors in which there is uncertainty.

Conversely, an adaptive approach recognises that many aspects of an area are changing, and that responses may be most effective when they are implemented over time, within the context of wider change and not being fixated on a single future risk scenario.

For the OxCam Arc, changes in flood risk exposure is primarily envisaged through climate and development. Adaptive planning offers the ability to respond to new information in the face of change and uncertainty. It can involve identifying and evaluating alternative sequences of interventions (adaptation pathways). By not committing to a single long-term plan - which may only remain optimal as long as circumstances do not change – adaptive planning introduces the capacity to change from one sequence of actions to another, thus creating flexibility within a long-term strategy that can evolve as circumstances change over time. In this study, we represent the level of **adaptive capacity** through the adoption of a real options analysis approach, which quantifies the value of adopting adaptive options in the short term based on the potential value under different futures. The real options analysis also aligns well with common adaptive capacity metrics which measure the flexibility and robustness of investment (see Section 2.2.2).

Adaptation should be a planned process which aims to moderate, cope with and take advantage of the consequences of change. A key goal of adaptation is to become more resilient, able to cope with – and even thrive under – the change. An adaptive plan can be viewed as a series of actions over time, where the different pathways represent proactive but flexible planning over time for a future which may differ from that which is currently understood (Lawrence, et al., 2019).

Adaptive planning is transformative and proactive as noted by Defra in its Evidence Review of the Concept of Flood Resilience (Department for Environment Food and Rural Affairs, 2020). It applies an iterative 'learning' process of plan, test, design and intervene to developing strategic responses to flood risk management, noting that flood resilience can by created by:

- Considering a wide portfolio of structural and non-structural interventions.
- Acknowledging inherent and emergent resilience within natural systems and communities as well as within the built environment.
- Embedding flood resilience within a wider narrative of thriving communities and land use despite the current and future threats of flood events.

In the context of future OxCam development, we want to understand how we adapt our flood risk management approaches to the changing future patterns of both climate change and development. We are interested in *what* mix of flood resilience investment (embankments, natural flood management, storage etc) is best in a given area (although the actual mix is not the focus), *when* these interventions should be implemented and *what* the economic justification for them is, i.e. *how much*. In other words, we want to invest the right amount in flood risk management at the right time, that will generate maximum benefit. This information can be used to inform the high level case for investment in flood resilience across the Arc. The actual mix of interventions is less a focus of this high-level investment study and will instead be considered during more local-level detailed studies.

Planning for development in the OxCam Arc is an iterative process that:

- Identifies and evaluates interventions to provide sustainable forms of flood risk management that need to adapt to two different forms of change: changes in land use driven by the desire for social and economic development along the Arc; and changes in flood frequency and severity (to existing and new development) driven by climate (and associated environmental) change.
- Recognises that there may be irreducible uncertainties in the timing, nature and spatial extent of planned development (especially in the short to medium term when the majority of Ox-Cam development is anticipated), and associated with projecting the implications of climate change on flood hazards, especially in the longer term (beyond the middle of this century).
- Needs to reconcile thresholds triggering interventions based on: (a) economic efficiency as measured by Net Present Value (NPV); and (b) levels of flood risk management that will provide effective assurance to the planning authorities, developers and users of future houses, buildings and infrastructure that flood risks will be managed to a tolerable level (which may de different to that associated with the highest NPV) into the future.
- Makes decisions in the early stages of development that deliberately avoid irreversible interventions that could lock development into a pathway for which flood risk management may not be sustainable in the long-term. The process of decision-making therefore needs to value flexibility and adaptive capacity alongside direct measures of economic efficiency (i.e. costs and flood damages avoided that are used in estimating the NPV of flood risk management interventions).
- Involves reconciling bottom-up and top-down approaches to assessing risks, developing resilience and planning adaptation measures. In other words, making decisions starting from now and working towards the future can be regarded as bottom-up (or left to right) and allows the learning and adaptive response to how conditions are changing. However, having an appreciation of the anticipated outcomes in the future (i.e. top-down, or right to left) can be important to influence what decisions are taken initially. For example, if most pathways converge on a particular portfolio of flood risk management interventions or appropriate level of investment, does that influence what decisions we make in the near-term?

2.2.1 Principles for adaptation

The following 2016 overall framework (Prutsch, Grothmann, Schauser, Otto, & McCallum, 2010) for undertaking successful adaptation remains relevant:

- Any adaptation needs to be **sustainable.** This means that our responses should not add to climate change, or limit the ability of other parts of the natural environment, society or business to carry out adaptation elsewhere. Our responses must avoid any detrimental impacts on other parts of society, the economy or the natural environment.
- Actions should be **flexible.** Although there is still uncertainty over the future climate, we should consider options now and make decisions that maximise future flexibility – in many cases it is failure to take decisions that locks us into inflexible pathways.
- Action needs to be **evidence-based –** making full use of the latest research, data and practical experience so that decision-making is well-supported and informed.
- Our response to climate impacts should be **prioritised –** for example, by focusing more attention on policies, programmes and activities that are most affected by development and climate, those which have long-term lifetimes or implications, where significant investment is involved or high values are at stake, or where support for critical national infrastructure is involved.

- Adaptation measures need to be **effective** (reducing the risks from climate change without introducing perverse effects), **efficient** (the long-term benefits of adaptation actions should outweigh the costs), and **equitable** (the effects of the activity on different groups and where the costs should fall should be taken into account).

2.2.2 Adaptation pathways

A pathway comprises a series of interventions (shown in Figure 2.1 as the vertical drop in risk when a new intervention is introduced), each of which has a design life. As the end to the effectiveness of one intervention is anticipated, or because of some other external change, a decision point is reached. This decision point is an opportunity to review environmental and other conditions and to decide on the next intervention. Before each intervention is implemented, sufficient lead time is needed to prepare for implementation (Lumbroso & Ramsbottom, 2018). Figure 2.1 shows two approaches to investment:

- The dashed blue line illustrates a precautionary approach, which significant upfront investment.
- The solid blue line illustrates an adaptive approach, with risk mitigation delivered adaptively and illustrated by coloured 'saw toothed' lines.

The dashed black line illustrates the risk increasing without any mitigation activities.

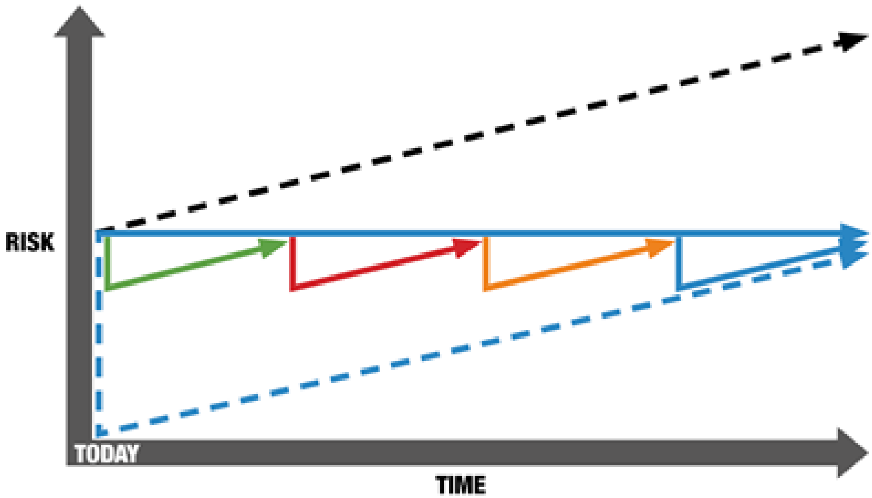

Figure 2.1. Schematic of an adaptive approach, where a series of interventions (represented by the 'zig zag' lines) achieve a target level of risk.

In this OxCam study, we do not seek to meet a pre-defined target level of risk – rather, we are aiming to identify the pathway of interventions which delivers the maximum net present value, from which we can identify the optimum investment (and as a by-product can identify the corresponding risk profile).

If multiple decision points exist through time, as is the case in an iterative adaptive approach, and there are a number of possible choices of interventions at each point, the range of possible pathways can be viewed as a decision tree (Figure 2.2). Each route through the decision tree represents a possible pathway. A 'map' of decision points and possible interventions can alternatively be visualised in the style of the London Underground map as illustrated in Figure 2.3.

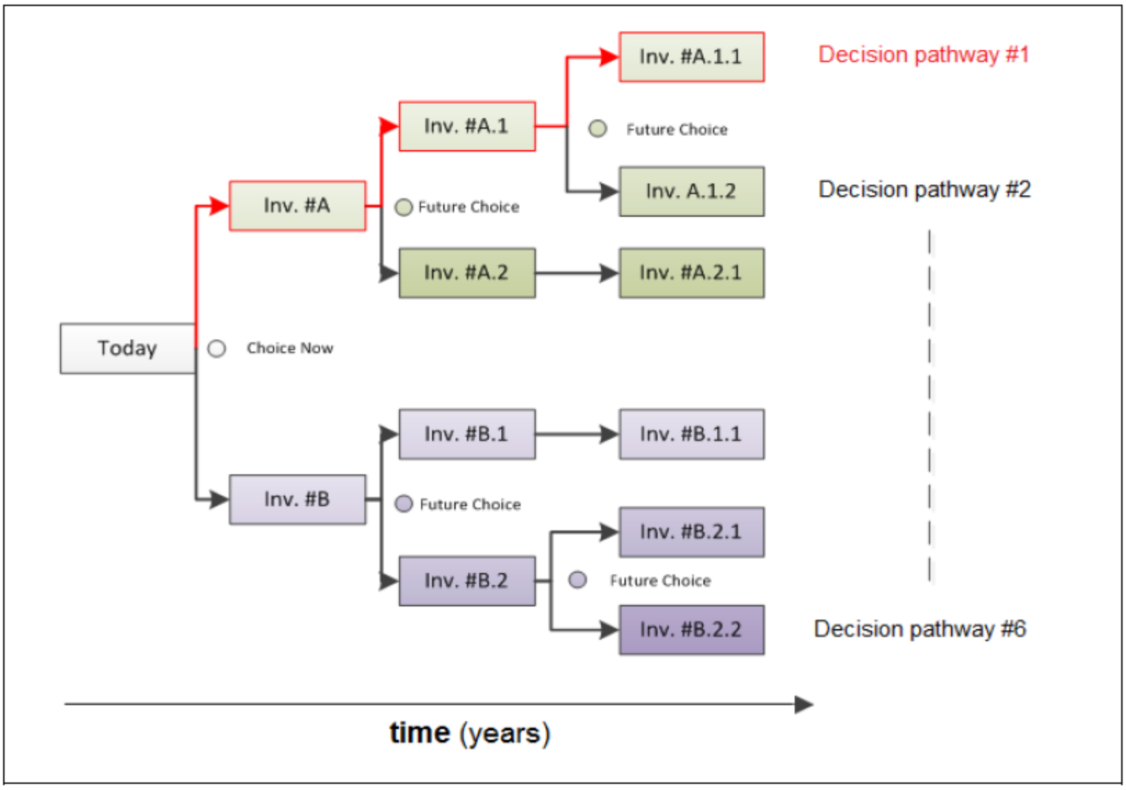

Figure 2.2. Illustrative decision tree representing six decision pathways and the associated choices that are available both now, and in the future (Environment Agency, 2018).

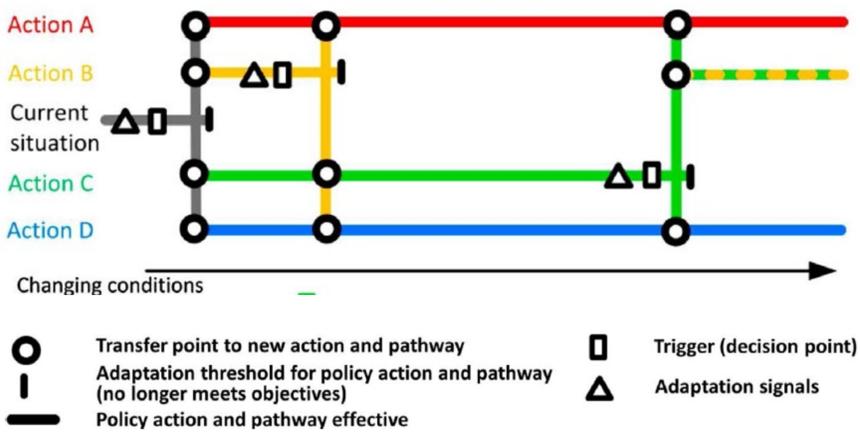

Figure 2.3. An example of an adaptation pathways map (Lawrence, Bell, Blackett, Stephens, & Allan, 2018).

As noted in the Environment Agency Adaptive Capacity Guidance (Environment Agency, 2018), pathways should be adaptive, but the individual interventions may not necessarily be viewed as adaptive. For example, raising flood embankments higher and higher is not adaptive because this is neither desirable nor practical. However, raising embankments for a period, during which time natural storage is developed behind the embankment and then used to replace the embankments, is adaptive when considered as part of the longer-term plan.

In addition to the more traditional annualised stream of costs and benefits used when evaluating interventions, the Environment Agency guidance introduces three measures which together could be used to measure the adaptive capacity of branches along the pathways in a decision tree:

- **Flexibility**: how many interventions remain open following an investment choice.
- **Robustness**: proportion of possible futures with this intervention as highest performer.

- **Opportunity lost**: difference between chosen intervention and best performing intervention.

It is worth noting that at an OxCam scale, these metrics may not be particularly meaningful – but we would expect the outcome of the **real options analysis** we have adopted to reflect the principles of these metrics: i.e. interventions recommended in year 0 are identified based on whether – in the balance of possible futures – they enable the highest whole life value pathways to be unlocked. This means that the interventions need to demonstrate **flexibility** and **robustness**.

Uncertainties affect assumptions around the type, rate and location of development. In addition, significant levels of uncertainty are associated with climate change projections. An ability to address the implications of uncertainty is one of the key advantages of adaptation decision-making, which recognises that:

- Not all uncertainties about the future can be eliminated
- Ignoring the implications of uncertainties could limit the ability to make adjustments in the future, resulting in situations that could have been avoided (e.g. lock-in to poorly adapted situations)
- Ignoring uncertainty could result in missed opportunities and lead to unsustainable plans and decisions based on them (e.g. path-dependency).

Critical to the success of an adaptive approach is to understand the (i) triggers which must be defined and identified which signal either a decision is required or interventions must be implemented and (ii) the lead times required for implementation of interventions. The relationship between triggers, lead times and implementation of interventions is illustrated in Figure 2.4.

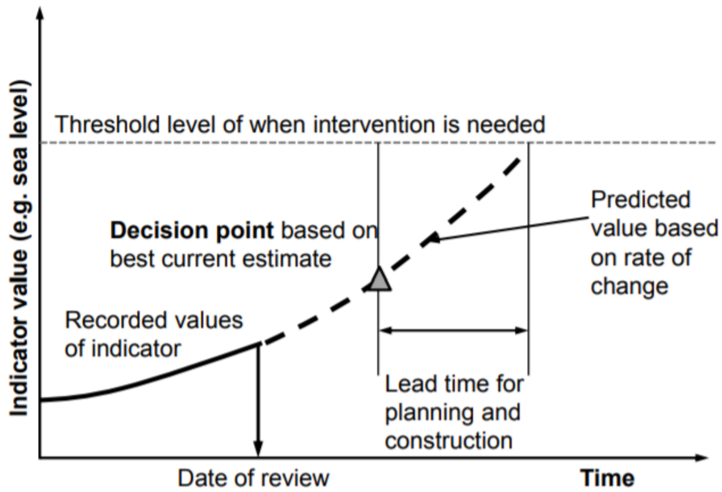

Figure 4 **Thresholds, lead times and decision points**

Figure 2.4. Thresholds, lead times and decision points (Lumbroso & Ramsbottom, 2018).

2.3 Economic optimisation

2.3.1 Long-term investment scenarios

The long term investment scenarios (LTIS) is an economic assessment of future flood and coastal erosion risk management in the period 2015 to 2065 (Environment Agency, 2014). The work assesses the consequences of investment choices to reduce the risks of flooding and coastal erosion, and explores the implications of different climate change scenarios and development scenarios.

At the heart of the LTIS projects is an economic optimisation that balances investment across England to find the highest net present value intervention policies. Because of the similarity in the goals of LTIS to this project, and the large-scale nature of the analysis, LTIS provides a strong foundation for the large-scale economic optimisation needed by the OxCam economic evidence study. As such, the approaches we have developed for

OxCam build on the LTIS approach. This section summarises the LTIS approach, as context to support the description of the approach taken in this project, which is introduced in Section 2.3.2.

LTIS defines the optimum level of investment under a given future scenario as the long-term level of investment that would be sufficient over the study period to fund all activity to manage flood and coastal risk where benefits are greater than costs, which generates the maximum net present value (NPV). It is based around the principle of applying a **policy** within each flood risk analysis unit (Halcrow Group Ltd, 2014):

1) Do nothing (no maintenance or replacement of assets).
2) Do minimum (maintain existing assets but don't replace).
3) Maintain crest level (maintain and replace assets to their existing crest levels).
4) Sustain current flood risk (raise defences and build new assets to keep pace with climate change).
5) Improve (raise defences and build new assets to reduce risk).
6) Improve+ (raise defences and build new assets to reduce risk further).

For a given future scenario, a set of rules are applied to identify the highest NPV policy in each a given flood risk analysis unit (referred to as 'Flood Risk Management Systems' (FRMS)), resulting in a 'basket' of selected policies and an associated investment profile (Neve, Smith, & Steel, 2016).

Most recently, LTIS 2019 expanded analysis in a number of ways, including the range of climate change scenarios explored, new development scenarios and wider impacts (including infrastructure impacts) (Environment Agency, 2019). While previous iterations of LTIS had focussed primarily on flood defence assets, LTIS 2019 explored approaches to quantifying the costs, benefits and therefore best value investment in a range of responses:

- Natural flood management.
- Temporary barriers.
- Property flood resilience.

As well as exploring investment in each of these interventions, LTIS 2019 also investigated how investment across a portfolio of interventions (combining the above with linear defences and flood warnings and incident management) could affect the overall optimum investment in flood risk management. Figure 2.5 shows the sequence of decision-making followed by LTIS 2019 – an example investment combination would be N2-R3-T1-W3-H4.

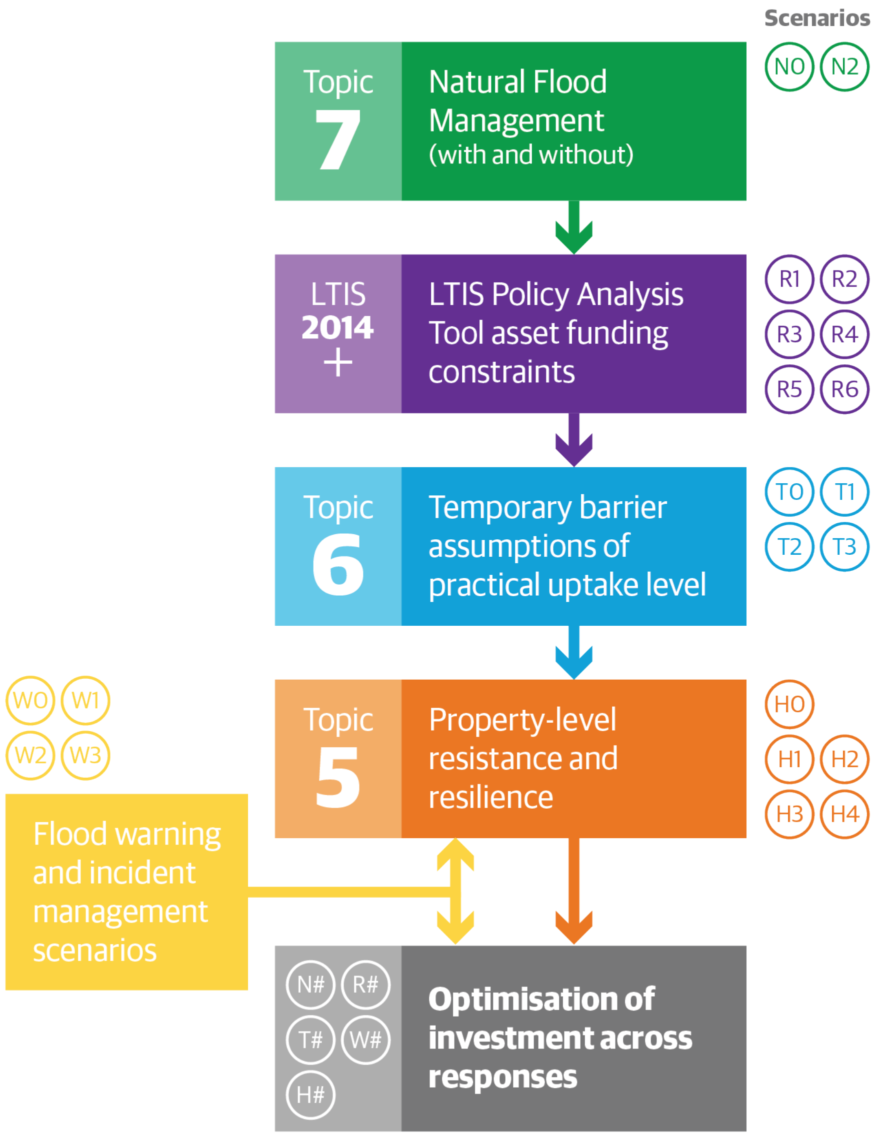

Figure 2.5. LTIS 2019 representation of a portfolio of interventions (Clarke, Wicks, Steel, Smith, & Neve, 2018).

Reflecting on the approach taken by LTIS, it is important to consider how the needs of the OxCam study is **different** to LTIS, as well as their similarities:

1) Both studies optimise the level of investment from an economic perspective, using net present value as the primary metric.

2) Both studies optimise investment across a portfolio of interventions which operate at different scales.

On the face of it, we could therefore consider applying the LTIS sequential method (shown in Figure 2.5) to the OxCam study area. However, unlike LTIS:

1) The purpose of the OxCam study extends to explicitly considering the **timing** of investment as well as the level. Although the analysis in LTIS allowed for different timings of investment, it was linked to pre-defined policies (as described by the numbered list above).

2) The OxCam study brings in the principles of adaptive planning, so needs to explore explicitly how the level and timing of investment might change in different futures – whilst maintaining a connection between those futures by recognising that the decisions we make **now** have to be adaptive to future changes in risk. This is quite different to LTIS, which only explores the optimum investment under each future in turn – meaning that decisions in year 0 could be different under each future.

This means that we had to move away from the LTIS optimisation approach, which was predicated on a) applying a set policy at each location and b) applying those set policies in sequence. Instead of thinking about interventions being optimised and applied in sequence, we needed to optimise the timing of investment and allow for investment in different interventions at different points in time. This meant developing an approach to optimisation which considered the portfolio of interventions – and the level and timing of investment in each intervention – holistically.

We could also consider how LTIS represented individual interventions. However, there was an expectation that, despite the large scale of the analysis (meaning that a local detailed approach would not be possible), methods to represent flood risk and interventions would reflect the catchment or local area's characteristics to a greater degree than LTIS. This meant that many of the simplifications that were justifiable at the national scale of LTIS cannot be applied to the OxCam study. For example, we have used flood frequency curves to represent the risk reduction benefits of storage and natural floiod management, and modifications to explicit depth-damage calculations to represent property flood resilience.

2.3.2 Defining the optimum

Optimisation, in general terms, is the process of finding the 'best' outcome – typically by maximising or minimising one or more metrics based on changes in variables, bounded by a series of constraints (Stanford University). These metrics are described as the 'objective function' of the optimisation. In the context of flood risk management, there are a number of possible objective functions that may be of interest. For example, we may be interested in:

a) Minimising the present value cost of achieving a certain outcome (such as a desired standard of protection, level of risk or other pre-determined flood risk management policy).

b) Maximising the net present value (i.e. present value benefit – present value cost).

c) Minimising the total present value cost (i.e. present value cost + present value residual risk).

d) Maximising the benefit cost ratio.

The flood and coastal erosion risk management appraisal guidance (Environment Agency, 2010) notes:

"The aim of the flood and coastal erosion risk management programme is to obtain best value for money for the whole programme, given the limited funds that are available. One way of ensuring it is to maximise the net present value (NPV) of all investment within the programme. An alternative solution is to maximise the benefit–cost ratio (BCR) of projects being funded informed by the incremental costs and benefits of doing more. This enables project level decisions to be informed by the efficiency aims expected of the wider programme. Generally, in flood and coastal erosion risk management, the latter concept is used."

Similarly, the Green Book describes the net present value and benefit cost ratio as the key economic metrics. It is also worth noting that the Green Book (HM Treasury, 2018) describes NPV as any discounted future values, and uses the term Net Present Social Value (NPSV) and Net Present Public Value (NPPV) to describe the stream of future costs and benefits to UK society that have been discounted using the appropriate Green Book rate. Therefore, when we use the term NPV consistently with FCERM appraisal guidance, we are referring to the NPSV as defined by the Green Book. The appraisal guidance also notes:

"Projects are unlikely to succeed unless both the ABCR and the IBCR are robustly greater than unity. Options showing a high ABCR will not usually succeed unless the IBCR is positive too as additional investments would deliver more elsewhere. By the same token, although in flood and coastal erosion risk management projects the do minimum option often exhibits by far the largest ABCR, this will not necessarily mean it is the best solution."

Table 2.1. Example calculations of target metrics for optimisation.

Options	Cost (A)	Baseline risk (B)	Residual risk (C)	Benefit (D = B - C)	Net present value (D - A)	Benefit cost ratio (D / A)	Incremental BCR [1]	Total cost (A + C)
0	£0m	£200m	£200m	£0m	£0m	0.00	0.00	£200m
1	£5m	£200m	£190m	£10m	£5m	2.00	2.00	£195m
2	£15m	£200m	£140m	£60m	£45m	**4.00**	**5.00**	£155m
3	£30m	£200m	£100m	£100m	**£70m**	3.33	2.67	**£130m**
4	£40m	£200m	£95m	£105m	£65m	2.63	0.50	£135m

Table 2.1 shows a series of hypothetical options, demonstrating how a) net present value, b) benefit cost ratio and c) total cost can be calculated from the same cost and risk data. The table also shows that different optimisation targets can lead to different outcomes (i.e. different options being selected as optimum). In particular, the maximum benefit cost ratio is likely to be at a lower investment level than the maximum net present value. However, noting the guidance above, it is worth highlighting that the maximum net present value option should be the highest cost option which delivers an incremental benefit cost ratio greater than 1. Building on this guidance, and aligning with the approach taken by the long-term investment scenarios (described in Section 2.3.1 above), this project adopted the **maximum net present value** as the objective function of the optimisation.

However, this project introduces an additional layer of complexity, in that we are interested in the optimum under multiple futures, acknowledging uncertainty in how risk (through climate change and development) may change over time. The Green Book provides guidance on a real options analysis which can be applied for this purpose, which is explored further in Section 2.3.2.

2.4 Adaptation in the context of optimisation

The fundamental question the project is trying to answer is: What is the optimum level and timing of investment from an economic perspective? Therefore, the primary metric which we measure along the pathways is the **net present value**, and the pathways identified as optimal (or 'most successful') are those with the highest NPV.

However, while one set of interventions might be optimal under a given future scenario, another might provide better value under a different future scenario. We are therefore interested not just in a single optimum, but also how that optimum varies under different futures. Furthermore, we are interested in ensuring that decisions made in the near term are flexible and robust under different futures. We therefore look beyond a suite of individual optimums, to a **real options analysis**, which is defined in the Green Book (HM Treasury, 2018). A real options analysis values flexibility to future uncertainty by weighting benefits by the likelihood of a particular future, and by considering options where investment is adaptive to that future uncertainty.

[1] Incremental benefit cost ratio is calculated based on the *additional* cost and *additional* benefit of an option compared with lower cost options. i.e. $(D_2 - D_1)/(A_2 - A_1)$.

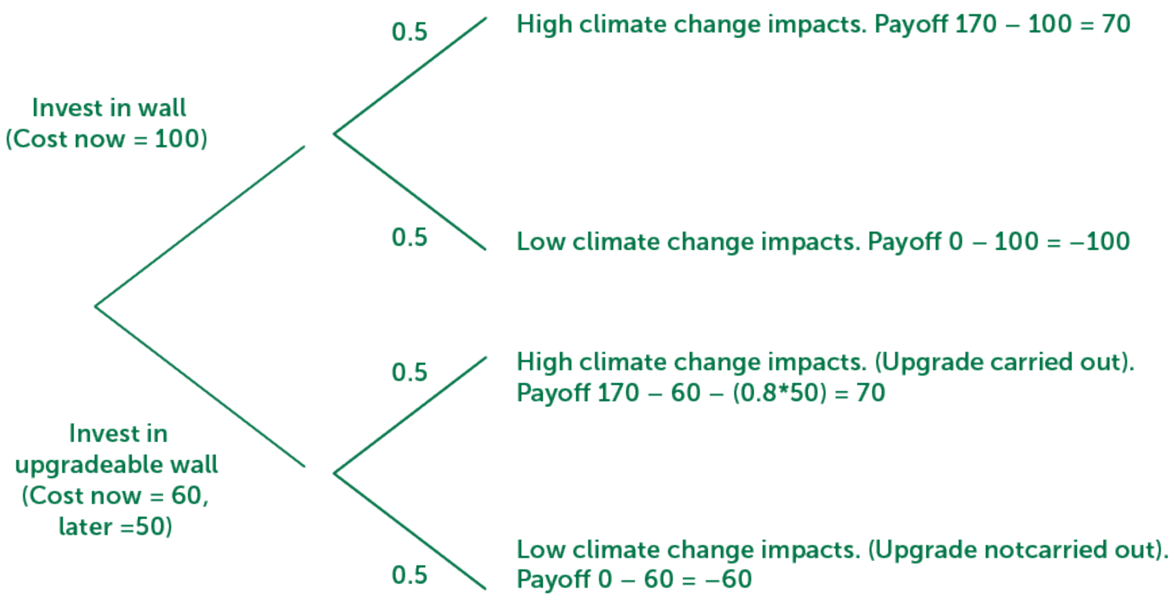

Figure 2.6. Example of a real options analysis, from HM Treasury Green Book Annex A5 Box 31.

In the simple example in Figure 2.6, an adaptive investment could include some upfront investment in an upgradeable intervention, then at a later date – if changes in risk mean it would be worthwhile – further investment to increase the scale of that intervention. As shown in the diagram, this can be compared with precautionary upfront investment – which under some futures may not deliver a return on investment.

Figure 2.6 can be compared with the decision tree shown in Figure 2.2, indicating how a real options analysis provides a route to optimisation that reflects adaptive decision-making and future uncertainty. Near-term investment which is *not* adaptive should inherently be 'rejected' by the real options analysis *if* the investment does not deliver better value for money than deferring some or all investment – this will be dependent on the nature of risk at a given location.

Could the optimal sequence of interventions be viewed as adaptive? The real options analysis implicitly accounts for the flexibility, robustness and lost opportunity of interventions, through the effect they have on the weighted net present value under all futures. For example, the net present value of a near-term intervention which only provides good value in half the futures will be 'dragged down' by a negative net present value in the futures in which it does not provide good value. By comparison, a near-term intervention which *enables* the best value interventions in each future (and, as in the 'invest in upgradeable wall' option in Figure 2.6, invests differently in different futures) should produce a higher net present value. However, negative net present values may still occur in individual futures, if the benefits are sufficient in other futures. Implementing the real options analysis in this context presents a challenge – we need to enable future decisions to be made at various points in time (we adopted 10 decision points at 10-yaer intervals), which adds a layer of complexity to the 'simple' case in Figure 2.6, which only has 'now' (when we don't know how the future will unfold) and 'the future' (when we do know). This therefore meant that we had to make an assumption about which of the 10 decision points should be 'fixed' under all futures (like the initial investment in Figure 2.6), and which would allow different investment decisions to be made in each future (like the future investment in Figure 2.6). Because knowledge of future risk is evolving quickly, in particular given the rapid development in OxCam, we adopted an approach where only the first decade of interventions (the decision point at year 0) were fixed in all futures – any future interventions were allowed to be different under different futures. In reality, to be truly adaptive, future decisions need to be built upon a revisited evidence base that reflects the evolving understanding of development and climate change.

This is also purely an economic analysis, although it has been developed to include wider economic impacts (including environmental impacts) and carbon costs, it does so in the context of maximising net present value. Carbon and environment impacts can therefore influence the result, but are not a hard constraint e.g. to provide environmental net gain or achieve 'net zero' carbon emissions.

The analysis could be extended to include other metrics of adaptive capacity, and also to test alternative optimisation targets. For example, what pathway of interventions achieved a 1% annual chance standard of protection (SoP) for the lowest cost, which is a more direct interpretation of Figure 2.1 and aligns with the National Planning Policy Framework (NPPF) requirement for the majority of development to be outside of Flood Zone 3a if possible. However, it is important to note that in the context of a portfolio of interventions, a so-called standard of protection for a given receptor may be delivered by a combination of interventions, and may not be a meaningful metric when considering property flood resilience.

3. Adaptive optimisation approach

3.1 Introduction

As a regional-scale economic optimisation, this study inherently diverges from a locally oriented adaption study, but must remain consistent with the underlying principles. A key difference is that a typical adaptive study is likely to *start* with local knowledge and build adaptive pathways based on that knowledge. The approach in this study could be considered as more objective, in that we start without any preconceptions of what the right combination or timing of investments might be, and instead use the data to guide us towards that combination. We also define the aim of the study differently – while the balance of investment across the portfolio of interventions is part of the results of the analysis, it is not the primary purpose, it is mainly a means to determining the optimum level and timing of investment.

Where a normal adaptive approach might set a locally-acceptable target for risk as the basis of choosing when and how to invest, we instead let the analysis tell us what is the most economically-beneficial level of protection that could be provided. From a practical adaptation perspective, its primary limitation is that we are focussed on **economics**, and are using the economic analysis to indicate the optimum approach to flood risk management and the robustness of interventions under future scenarios – therefore excluding the community aspects of adaptive management and risk tolerance.

There are a number of aspects of the optimisation which need to be defined:

- **Portfolio development**: How we define available portfolio(s) of interventions, and the combinations of interventions which can be made. This could be policy-based, as seen in LTIS, or intervention-led as seen in this study.
- **Pathway development**: How we define decision points, and the rules for sequencing of interventions and divergence of pathways under different futures.
- **Optimisation target function**: What we are optimising (net present value), and how it this measured under different futures (real options analysis).
- **Optimisation routine**: What computational approach we are taking to solve the complex problem.

The remainder of this section explores the first three, and Section 4 is dedicated to describing the optimisation routine.

3.2 Portfolio development

The flood risk management (or flood resilience management) system is made up of a portfolio of interventions which interact with the system (and each other) and a variety of ways. The scope of this project includes the representation of:

- **Catchment-scale storage**. We represent the equivalent of a catchment-wide system of storage interventions by a universal shift in annual exceedance probability, based on the total volume removed from the flood hydrograph's peak (allowing for a base flow). For the purpose of the optimisation, catchment-scale storage defines the largest spatial analysis scale. The cost model is based on the cost per storage volume.
- **Natural flood management (NFM)**. We represent the equivalent of a package of NFM interventions in a sub-catchment (tributary), based on the total volume removed from the flood hydrograph (therefore assuming NFM can be represented as an equivalent storage volume, and also that NFM in a tributary has a marginal effect on main river flows). NFM sub-catchments are an intermediate spatial scale, with a number in each (storage) catchment. The cost model is based on the cost per equivalent storage volume.
- **Linear flood defences**. These are represented explicitly – for a range of defence heights – in the flood simulations. Investment in linear flood defences is at the flood area scale, with each defence height scenario set to match river levels for a given annual exceedance probability event in the flood area. The cost model is based on the height-varying cost per unit length for different asset types.

- **Property flood resilience**. This is represented through modified depth-damage curves that represent a package of property flood resilience measures in a property, and are implemented at a flood area scale (to properties whose flood probability exceeds a minimum threshold). Property flood resilience is therefore represented as a change in impact as part of the economic analysis rather than as a change in flood hazard. The cost model is per typical residential property.
- **Surface water management** (sustainable urban drainage). We represent the equivalent of a portfolio of SuDS interventions in a "flood risk management system" (a spatial analysis unit with full coverage of England), based on the total volume removed from the floodplain, which assumes that SuDS can be represented as an equivalent storage volume, and that the volume of water on the floodplain represented in the surface water outputs is equal to the volume of rainfall in excess of the drainage capacity. The cost model is based on the cost per equivalent storage volume.

Because of the size of the study area, it is difficult to clearly visualise all the different spatial units associated with the interventions above together. Figure 3.1 shows the spatial units used by river flooding interventions: Catchments for storage, sub-catchments for natural flood management and flood areas for linear defences and property flood resilience. It shows a study area view, and a zoomed in view which better demonstrates how the flood areas and sub-catchments are related. Figure 3.2 shows the spatial units used by surface water interventions (flood risk management systems).

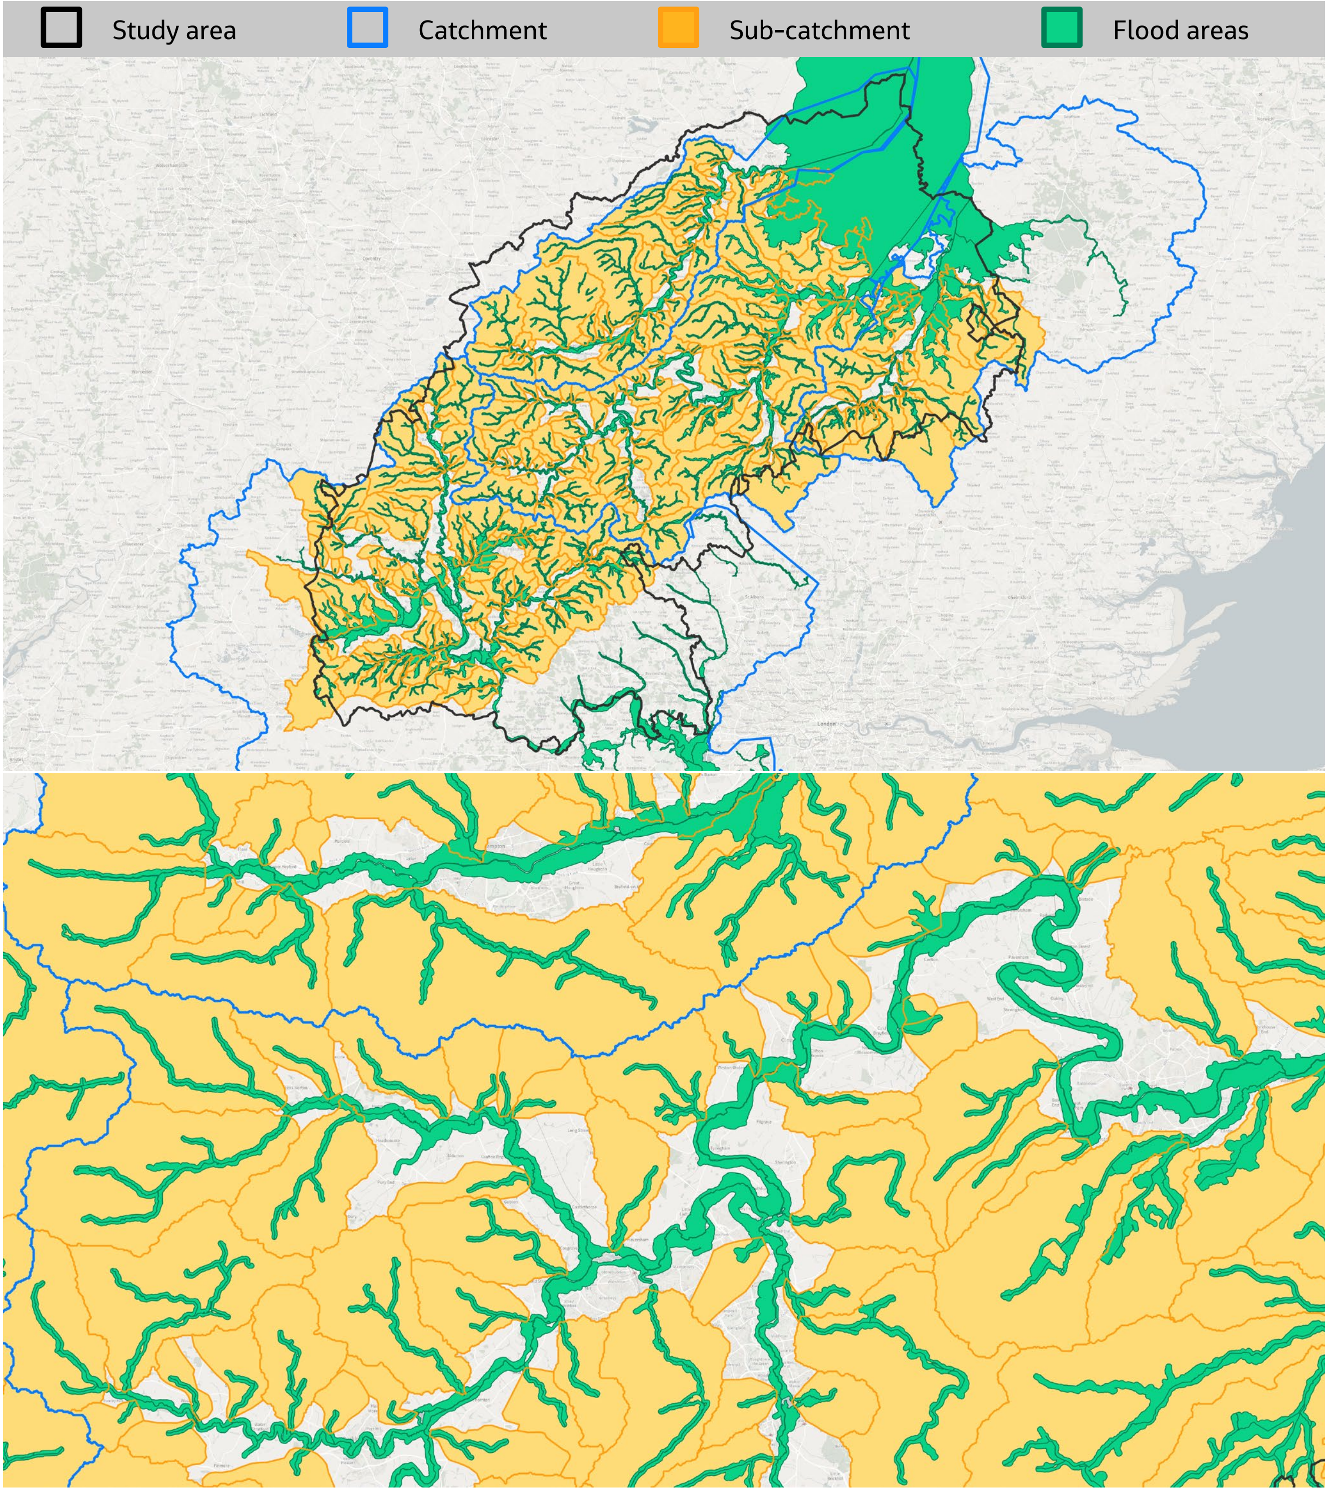

Figure 3.1. Spatial units used by river flooding interventions: Catchments, sub-catchments and flood areas. Top: study area view. Bottom: zoomed in view.

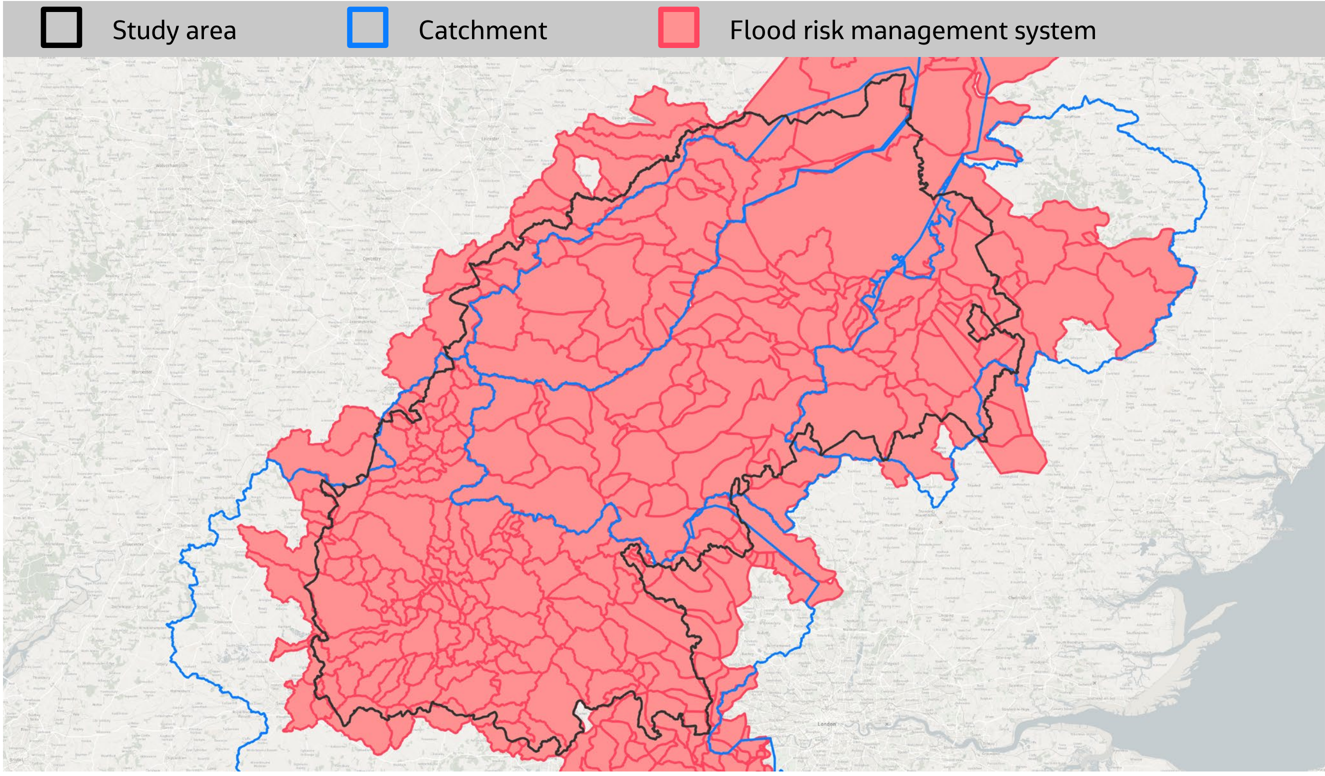

Figure 3.2. Spatial units used by surface water interventions (flood risk management systems).

Detailed descriptions of the modelling representation of each intervention can be found in the modelling technical report, and a detailed description of the cost model can be found in the economics technical report.

There are a number of ways in which these individual interventions could be pieced together to form portfolios, and the selection of approach has a profound effect on the nature of the optimisation:

1) **Policy-driven interventions**

   Defining the range of options for each intervention as a series of policies (such as 'do nothing', 'keep pace with climate change' and 'enhance standard of protection'). Each intervention would be explored in sequence to optimise investment in each intervention.

   The work in LTIS 2019 to optimise investment across the portfolio of interventions followed an approach similar to this, as described in Section 2.3.1 above. However, a number of the interventions were represented based on their physical representation rather than as policy-driven interventions.

   The primary limitation of this approach for this study is that the study includes a large number of futures, and each intervention needs a physical representation to match each policy in each future – including extreme options that are less likely to be practically feasible (such as enhancing level of protection in a H++ climate change scenario).

2) **System-level policies**

   a) Defining a series of system-level policies. A system could be defined as any single spatial unit – such as a catchment or a flood risk management system. A system-level policy could be one of the policies described above (such as 'do nothing', 'keep pace with climate change', etc.) – but achieved through the combination of interventions benefitting that system, rather than a single intervention.

   b) Finding the optimal way of achieving each system-level policy by optimising investment across the portfolio of interventions.

   c) Selecting the optimal balance of policies across systems. The core LTIS 2014 'policy analysis tool' optimises investment in this way.

The primary limitation of this approach for this study is that the system of interventions exists at a scale above individual local systems – with catchment scale storage and sub-catchment natural flood management, which would make defining system-level policies highly challenging.

3) Physically-derived interventions

a) Defining interventions based on locally-relevant physical representations, such as a representative range of storage volumes relevant to a catchment and realistic defence levels of protection (and broadly realistic defence heights – although using a consistent set of protection levels means there are always some outliers) to provide a range of available intervention options – e.g. 'small', 'medium' and 'large' storage interventions.

b) Relatively unconstrained catchment-level optimisation to select from the available sets of interventions.

Given that limitations and implementation challenges of 1) and 2), we have adopted option 3) for this study. An output of the analysis could be an assessment of the equivalent policy being adopted across the catchment and locally.

3.3 Pathway development

An adaptation pathway is built up of a number decision points that link future decision-making to potential trigger points. The aim of optimisation in this context is to navigate the near-term and future decisions to understand which deliver the highest whole life value, given that we face uncertainty in future risk due to climate change and development.

We can therefore consider the problem as a decision tree built by defining a series of decision points through our study period. These decision points are the moments in time at which investment decisions can be made in our analysis. Where an adaptation pathway would develop specific triggers that lead to decision points, our analysis assumes we have no preconceptions of when investment may need to be made. Indeed, identifying suitable triggers for a decision point is an important next step in translating the evidence emerging from this study into an adaptive plan.

At their most dense, decision points could be every year – but there would be significant computational implications of this if we are to analyse the costs and benefits of each branch of the decision tree (as the number of 'branches' increases exponentially with each additional decision point), so we have adopted decision points every 10 years.

In this analysis, we do not pre-define specific pathways or portfolios of interventions. Instead, we apply simple rules (as described in Section 3.3.4) to ensure that each pathway considered in the optimisation is legitimate and follows a logical progression of investment.

3.3.1 Available options at each decision point

Figure 3.3 shows the range of options available at each decision point for fluvial/tidal risk – surface water flood risk is considered separately to fluvial/tidal risk, but surface water interventions are considered in the same way in parallel (with fewer intervention combinations).

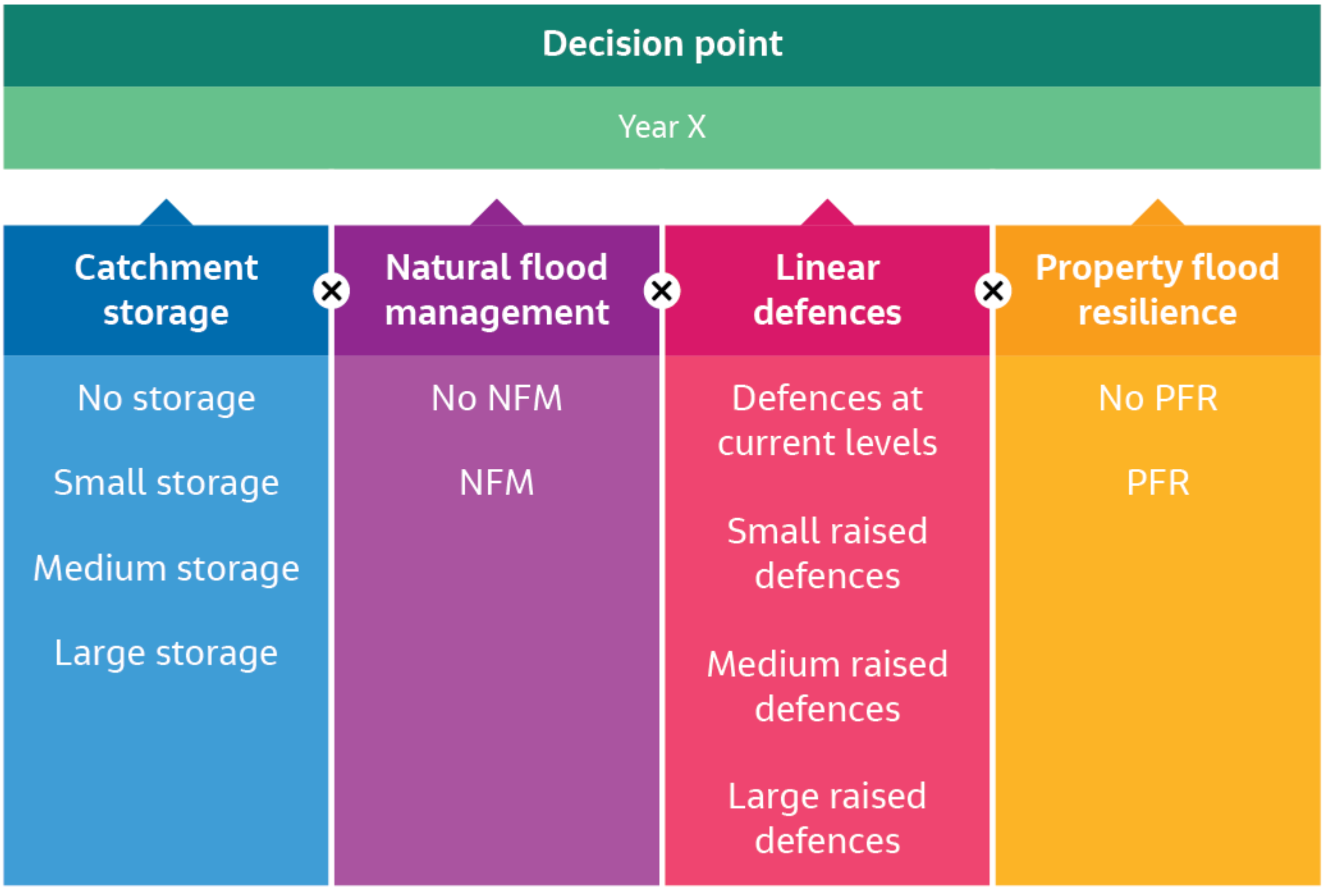

Figure 3.3. Range of options available at each decision point.

3.3.2 Construction of a pathway

This process is repeated at each decision point, to identify each possible investment pathway. Figure 3.4 shows the full set of fluvial/tidal intervention scenarios (represented either through hydraulic modelling or impact calculation modifications), including combinations of scenarios – each horizontal line represents a different scenario. The pink line shows an example investment pathway, 'stepping' between available results to calculate the change in impacts, and the different viable routes between these represent the branches of a decision tree. While Figure 3.4 represents the full set of intervention combinations that might be benefitting a given flood area, it does not however represent the added decision-making complexity because of the different scales the interventions operate at – meaning that the pathways at a flood area level are linked by interventions which span across a wider (catchment or sub-catchment) scale.

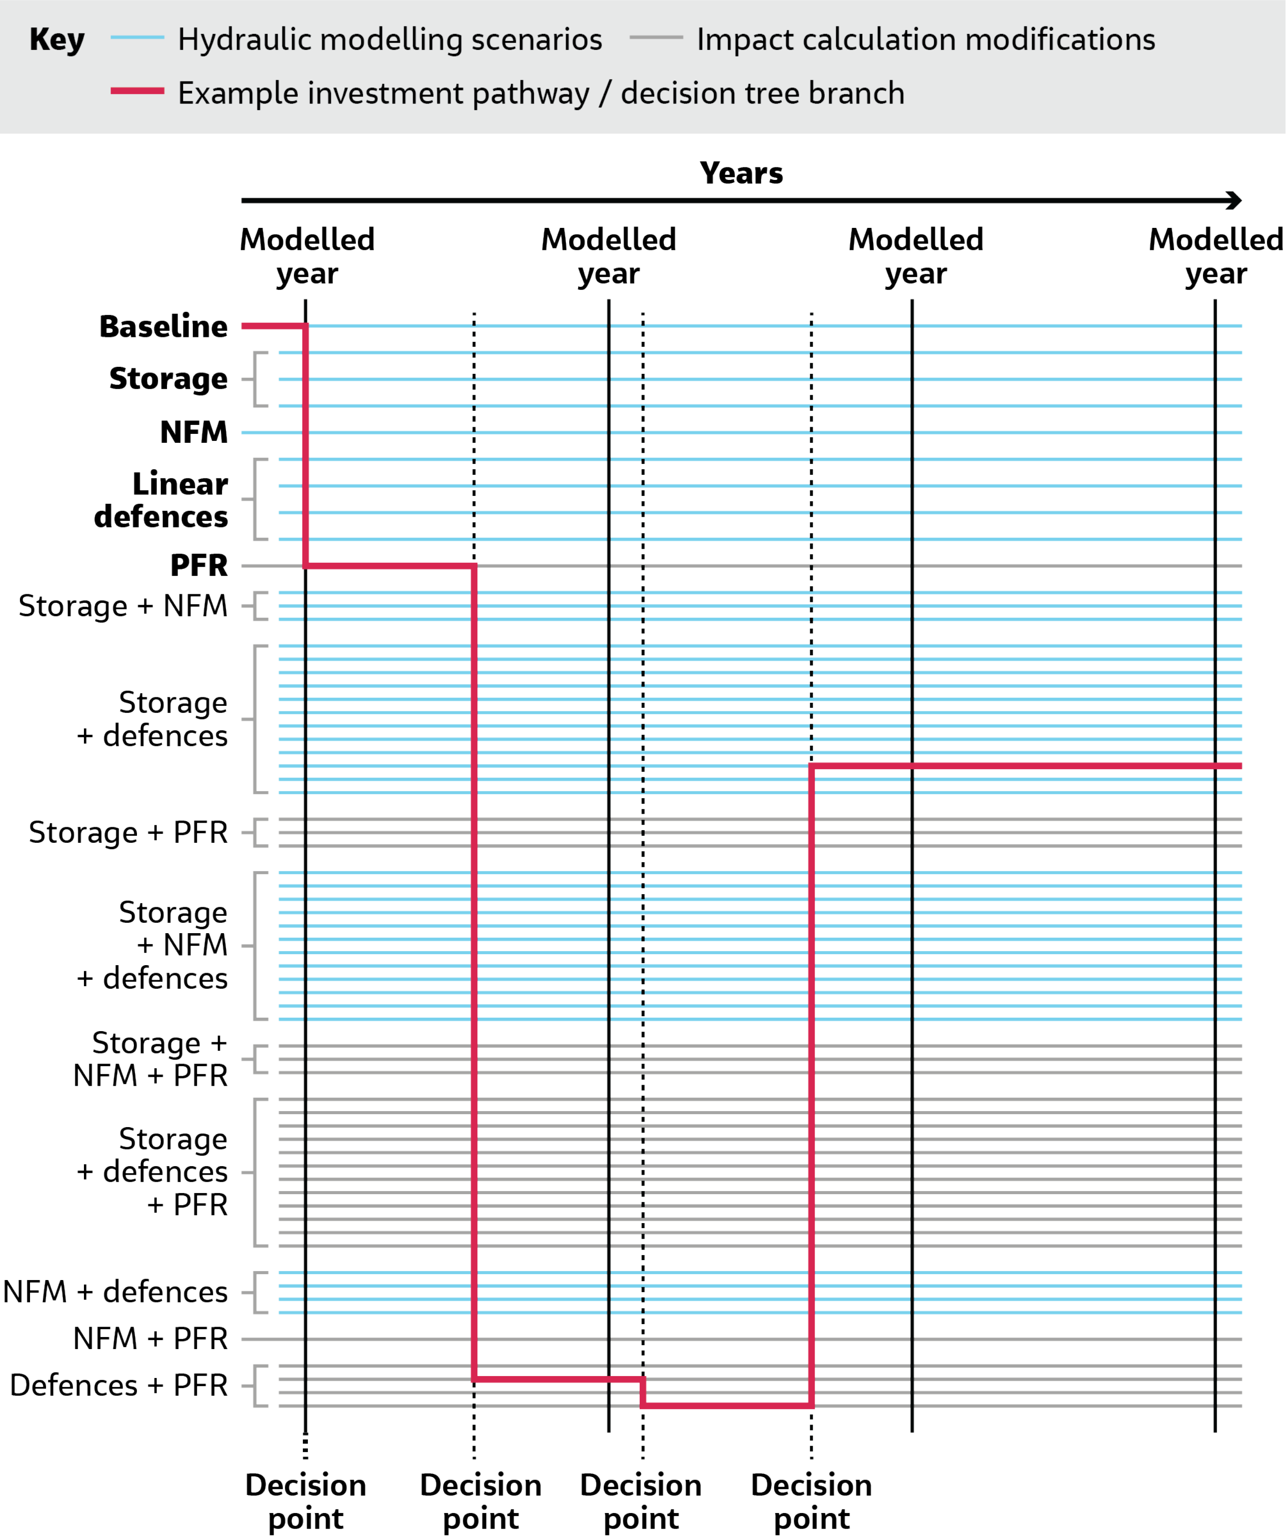

Figure 3.4. Implementation of fluvial/tidal scenarios through time, and an example investment pathway.

Stepping through the different pathway interventions creates:

- A shift in the annual average damage, overall producing a damage profile, which can be discounted to produce a present value impact, and compared with the baseline to produce a present value benefit.
- Cost information based on the year of investment (capital cost, annual maintenance costs and lifespan), which can be discounted to produce a present value cost.

3.3.3 Dealing with spatial scales

We must also consider the different spatial scales of fluvial/tidal interventions: Catchment scale storage, sub-catchment scale natural flood management, flood area scale linear defences and property-scale flood resilience. Alongside these we also consider surface water interventions. Our primary analysis unit for fluvial/tidal investment decision-making is at the flood area scale, and at the FRMS scale for surface water – shown in Figure 3.1 and Figure 3.2 in Section 3.2. However, decisions at the catchment level affect all flood areas within the catchment (just as decisions at a flood area level affect the cost effectiveness of catchment interventions), so the decision-making algorithm described in subsequent sections needs to ensure that intervention decision-making is applied consistently, and that the optimisation finds the best catchment-wide portfolio of interventions. For the purpose of the analysis, we have four catchments (Cam, Great Ouse, Nene and Thames) representing four possible catchment storage interventions. The Cam is a tributary of the Great Ouse (with the confluence South of Ely), so for reporting we present three catchments with the Cam and Great Ouse combined.

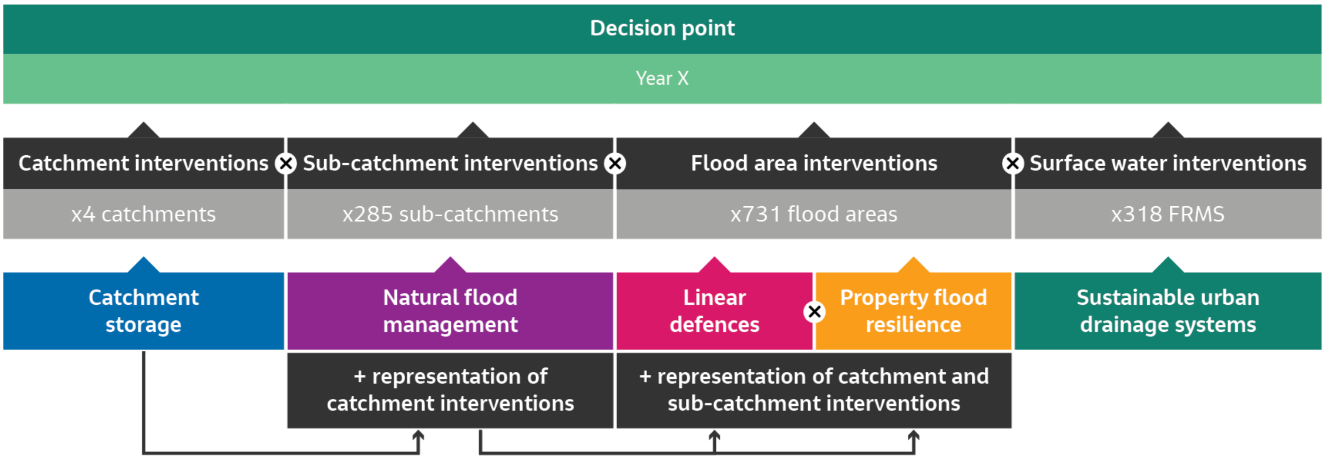

Figure 3.5. Analysis of parallel decision trees.

3.3.4 Precluding future options

We have applied simple logic whereby one decision at a flood area or FRMS scale precludes certain other interventions at future decision points. We do not wish to artificially exclude interventions that could be achieved, but we have adopted the following decision rules, designed to represent interventions that you could not feasibly implement, rather than interventions that we believe would be ruled out on economic or acceptability grounds:

- Catchment storage:
  - **Can** be added to in future to increase the total volume of storage in a catchment. This could be through investment in additional storage locations which contribute to the catchment-wide provision of storage, or by enlarging existing storage.
  - Cannot be removed to reduce the total volume of storage in future.
- Natural flood management:
  - Cannot be removed from a given sub-catchment once it has been invested in.
- Linear defences:
  - **Can** be raised in future (to allow for an adaptive rather than precautionary approach).
  - Cannot be removed or lowered.
- Property flood resilience:
  - Cannot be removed from a given flood area once it has been invested in.

The level and timing of investment at a location may change over time, and at different locations the level of protection may increase or decrease relative to climate change as a result of the changes in the portfolio of interventions. These changes may be dependent on future changes in risk – allowing the analysis as a whole to enable adaptive decisions. However, we assume that the *absolute* level (e.g. crest level, storage volume, etc.) of an individual intervention will not **reduce** over time – therefore excluding managed retreat from this analysis.

3.4 Optimisation target function

Sections 2.3 explores the range of potential optimisation target functions that could be explored. In summary, we have adopted **net present value** (NPV) as the primary economic metric, calculated as the present value whole life benefit (PV_B; over a define appraisal/study period) minus the present value whole life cost (PV_C; over that same period):

$$NPV = PV_B - PV_C$$

In-year cash costs (based on present day costs) are discounted using standard long-term discounting rates defined by the Green Book (HM Treasury, 2018). These economic aspects are described in more detail in the **economics technical report**.

The optimum level and timing of investment in a given scenario is defined as that which produced the maximum net present value. As noted in Section 2.3, adopting this metric also means that the benefit cost ratio is greater than 1 (i.e. each intervention is cost beneficial).

Section 2.4 introduces the **real options analysis** as a means to optimising investment (maximise net present value) in a way which reflects uncertainty in future changes in risk.

We have implemented an approach with aligns with the real options analysis, as follows:

- We consider 1 or more futures in a given scenario. The scenarios include:
  - All futures: Every combination of development rate (23k and 30k homes per year), development shape (new settlements, expansion of existing settlements, and hybrid) and climate change (central, upper central and high++) is considered, with the 18 futures assumed to have an equal probability of occurring. This scenario originally included a 43k homes per year option, but this was removed following discussion with the Environment Agency project team as it is not considered to be a realistic future scenario.
  - No development: Each climate change future (central, upper central and high++) is considered, but with no new development (existing properties only), with the 3 futures assumed to have an equal probability of occurring.
  - 27 individual future scenarios, representing each scenario combination, optimising investment in the single future. This provides a comparison between 'robust' investment under the 'all futures' scenario and the optimum investment in a single future. These single future scenarios are 'simple' net present value calculations.
- Interventions in year 0 (i.e. present day) are assumed to be the same under all futures, because these decisions have to be made based on our current understanding of risk.
- Interventions in subsequent years (i.e. at future decision points) can be different under each future, because these decisions can be made once we understand more about how future risk is unfolding.
- The total net present value (NPV) under the combined set of futures (f) in a scenario is equal to the sum of the net present values in each future (NPV_f) multiplied by the probability of that future ($P(f)$):

$$NPV = \sum^{f} NPV_f \times P(f)$$

- More precisely, this may be considered to mean the sum of the net present value of year 0 interventions (I_0), plus the sum of the net present values of the intervention options adopted in each future (I_f) multiplied by the probability of that future (*P(f)*):

$$NPV = I_0 + \sum^{f} I_f \times P(f)$$

A key challenge in implementing a real options analysis is estimating the likelihood of each future. Within the constraints of this project, we have not been able to explicitly estimate probabilities for each future. We have focussed instead on providing a breadth of futures, and assume an equal probability of each future occurring – and explore through sensitivity analysis the effect of this assumption.

3.4.1 Derivation of a "baseline" scenario

As described above, net present value is based on a whole life present value benefit minus a whole life present value cost. The present value benefit is in turn calculated by subtracting the whole life present value flood impacts with interventions from the present value impacts for a baseline scenario.

We have defined our baseline scenario as a scenario where existing linear defence assets remain in place. Maintenance and replacement of existing assets is therefore excluded, but we do allow for raising of existing defences – adopting an incremental capital and maintenance cost, to understand the *additional* investment needed in the study area over current continued maintenance of assets.

We originally envisaged creating a true 'do nothing' scenario where existing defences were left to deteriorate, or a 'do minimum' scenario where existing defences were maintained for the rest of their expected lifespan, but not replaced. However, it became clear that this presented practical issues – in particular because:

- We would have had to assume a fixed lifespan for all existing defences in a given flood area, rather than basing lifespan on the type of assets, their current condition and estimated deterioration rates, otherwise modelling the asset base within a flood area would have required numerous additional configurations of assets.
- The 'do nothing' scenario entails substantial maintenance costs.

3.4.2 Investment constraints

We have not included affordability constraints. Because we are seeking to find the optimum level and timing of investment, we are looking for the answer to the question of how much money needs to be spent and when, rather than to artificially steer the analysis by applying a constraint in terms of e.g. maximum annual investment. However, given that we apply investment at 10- year intervals, we have post-processed results to 'smooth' investment over the 10 -year period.

This also applies to funding. The optimisation focusses on finding how much money *should* be spent, not on how that investment is funded. The optimisation therefore does not take into account flood defence grant in aid or partnership funding; rather, it maximises net present value – which means any intervention with a benefit cost ratio greater than 1 could be included.

3.4.3 Study period

Given the aim of the project to optimise investment over the next 100 years, we explored running the optimisation over a 150 year study period, but with results reported for 100 years. However, this introduced challenges with how we extend the representation of climate change beyond the farthest epoch. While we could find technical solutions to these challenges, ultimately project constraints led us to adopt a 100 year study period.

We ran a sensitivity analysis (Table 3.1) over part of the study area to understand the implication of changing a) the study period (the period over which the optimisation is run), and b) the reporting period (the period for which the final results are calculated), through the effect on the optimum present value investment. It indicated that the results were not particularly sensitive to the change in study period (or reporting period) – which we assume to be a by-product of the level of discounting that would be applied to investment that far into the future.

Table 3.1. Exploration of study periods and reporting periods and their effect on the optimum present value investment.

Study period	150	150	100
Reporting period	150	100	100
Nene	£1,107 million	£1,095 million	£956 million
Ouse (part)	£206 million	£202 million	£229 million

4. Optimisation routine

During the project, we explored a number of different optimisation approaches as candidates for optimising the level and timing of investment in the OxCam Arc, including:

1) 'Guided' brute force decision tree exploration.
2) Dynamic programming.
3) Routing algorithm.
4) Genetic algorithm.

These are explored in Section 4.1. Ultimately, we adopted a bespoke optimisation routine that took elements of some of the above approaches, which is explored in Section 4.2.

This project presents a number of fundamental challenges – which differentiate the OxCam challenge from e.g. LTIS before it. In particular:

- We must account for a system of interventions that operate at different scales – catchment, sub-catchment and flood area. This means that we need to be able to make like-for-like comparisons between, say, a catchment storage intervention and a series of flood area linear defence interventions that both benefit the same areas. This also means that the policy-driven approach adopted by LTIS (e.g. "maintain crest level" or "keep pace with climate change") would not readily translate to this problem.
- We are exploring the **timing** of investment as well as the level, in a way which is consistent with adaptive planning. This means we need to allow for changes in interventions and changes in investment policy over time. This compares with the core LTIS 2014/2019 analysis which picked from a predefined set of 6 asset investment policies and adopted 1 policy per analysis area.
- We are exploring investment under multiple futures through a real options analysis optimisation that tests robustness under all futures, not just optimising investment under a single future.

As a pilot project, exploring the different optimisation routines has been a key part of this project, and we have learnt a great deal about the implications of the complexities introduced by the above challenges.

4.1 Considered approaches

4.1.1 Guided brute force decision tree exploration

In a brute force approach, we would seek to explore **every** branch of a decision tree to evaluate the net present value of all possible decisions in each future. We have inherently 'guided' our analysis in a number of ways by:

a) Creating a pre-defined set of interventions (such as small, medium and large defences, catchment storage, etc.).
b) Applying investment decisions at 10-year 'decision point' intervals – turning a continuous study period into discrete investment blocks.
c) Applying logic to the sequencing of interventions (for example, by assuming that small defences can be raised at subsequent decision points, but that they will not be lowered or removed).

The advantages of a brute force approach would be that a) it is conceptually simple (and therefore likely to be simple to implement), and b) it produces a complete database of results from which additional analysis can be derived (e.g. answering different questions without having to revisit the underlying optimisation process), with the constraints above being used to construct the decision tree, and prune branches which do not follow the required logic.

The available intervention combinations for a given flood area are made up of different levels of investment in the 4 main intervention groups – catchment storage, natural flood management, linear defences and property flood resilience. This results in up to 64 possible intervention combinations at each decision point (e.g. medium catchment storage + NFM + small linear defences + PFR). This in turn means that – for a single flood area – with 10 decision points there are up to 64^{10} intervention combinations, which would equate to 1,152,921,504,606,850,000 combinations. While some of these would be ruled out through the application of sequencing logic, the size of the problem space means that we do not believe such a brute force approach would be viable for this size of problem. For this approach to be feasible, it could be further constrained to create a more guided brute force approach. Care is needed to ensure that we could be confident that these constraints only exclude pathways which would not be feasible, rather than pathways which are realistic candidates that could form part of an optimum investment profile.

We could be more pragmatic, for example including a decision point every 10 years during the first 30 years of the study period (to enable investment to respond to evolving growth in the Arc), and then less frequently after that point. Or, for example, ruling out interventions which are not cost beneficial in isolation (assuming that in combination with other interventions they are also not be cost beneficial). However, ultimately, we believe a more intelligent exploration of the problem space would be necessary.

4.1.2 Iterative dynamic programming

"Dynamic programming is an optimisation approach that transforms a complex problem into a sequence of simpler problems; its essential characteristic is the multistage nature of the optimisation procedure." [2]

A computer science application of dynamic programming often involves breaking down an optimisation problem into simpler sub-problems which only need to be solved once and re-used[3]. A common example is the Dijkstra network routing algorithm, which is built on the principle that the shortest path between any 2 nodes only needs to be calculated once[4], and can form part of multiple larger solutions.

There are different ways that dynamic programming concepts may help us to solve our problem, in particular:

a) Approaches to simplify the solution through iterations which navigate the decision tree to find the optimum pathway.

b) Approaches to enable calculated results for part of the solution to be efficiently re-used. For example, by identifying future points where different branches of the decision tree converge.

To take an iterative dynamic programming approach, we could apply the same logical sequencing described above in Section 4.1.1, but adopt an approach which reduces the numbers of branches we need to explore and follow a more iterative path through the decision tree:

1) On the first pass, we consider all intervention combinations (64) at year 0. We look at each one in turn, and for each future decision point we simplify the decision-making so that the highest NPV combination is selected at each step. This generates an indicative NPV along the whole pathway starting with each of the 64 intervention combinations. From this, we take the combination from year 0 resulting in the highest indicative NPV forward to the second pass.

2) On the second pass, we take the year 0 intervention combination selected in the first pass, and explore all valid intervention combinations at year 10, again exploring each in turn using the simplified route through the decision tree. This provides an indicative NPV for each of the available combinations in year 10.

3) This process is repeated for each decision point – iteratively exploring each decision point through a simplified forward look, until all 10 decision points have been explored in detail. The result should be the highest NPV route through the decision tree.

The advantage of this approach is that it simplifies the problem space dramatically, whilst retaining the depth to which the highest NPV route is explored. This approach is derived from past experience of applying this dynamic

[2] http://web.mit.edu/15.053/www/AMP-Chapter-11.pdf
[3] https://www.freecodecamp.org/news/demystifying-dynamic-programming-3efafb8d4296/
[4] Dijkstra, E. W (1959) A note on two problems in connexion with graphs. Numerische Mathematik 1, 269 – 271. https://ir.cwi.nl/pub/9256/9256D.pdf

programming approach, in particular for optimisation of water utilities investment. We trialled the adaptation of an existing tool with a strong pedigree in that sector, but we believe it would be necessary to take that knowledge and implement new code to solve the problem, because there are a number of key differences between the problems:

- Water sector optimisation is oriented towards finding the least cost solution to meet specific consumption demands. For our flood resilience problem, the analogy would be setting a target 'standard of protection' and finding the most cost-effective way of meeting it. However, in this project we have a less constrained problem space, bounded only by the options we have considered and whether they have a positive net present value.
- Water investments are typically additive and independent. This means that the costs and benefits of each investment can be simply added together to derive a total cost and benefit. However, in flood resilience, the benefits of interventions in a given location are dependent. For example, investment in linear defences will affect the benefit of investment in property flood resilience. In addition to this, for OxCam we are considering dependent interventions at multiple scales: catchment, sub-catchment and FRMS. This adds further complexity, which is discussed below.

One concern of adopting this approach is that it may bias results towards bigger investment earlier, because at each step the highest NPV option is selected based on selecting the highest NPV at each timestep, which therefore negates the possibility of deferring investment. It would therefore be key to ensure that the real options analysis is implemented in a way that enables the each of the iterations to take into account the balanced NPVs under all futures – and select the highest NPV weighted by futures.

However, the primary issue in adopting this approach is the volume of analysis still required, primarily because of the inter-related catchment, sub-catchment and flood area scale interventions. We therefore explored taking a two-step approach within each time-step of the analysis, i.e. at each step, looping through all catchment intervention scenarios then all flood area level intervention scenarios.

4.1.3 Routing algorithm

The Dijkstra routing algorithm is a type of dynamic programming algorithm designed to solve the problem of finding the **shortest path** between two nodes. As shown in the example in Figure 4.1, each 'link' has a weight assigned to it – such as the travel time along the link – to support the optimisation.

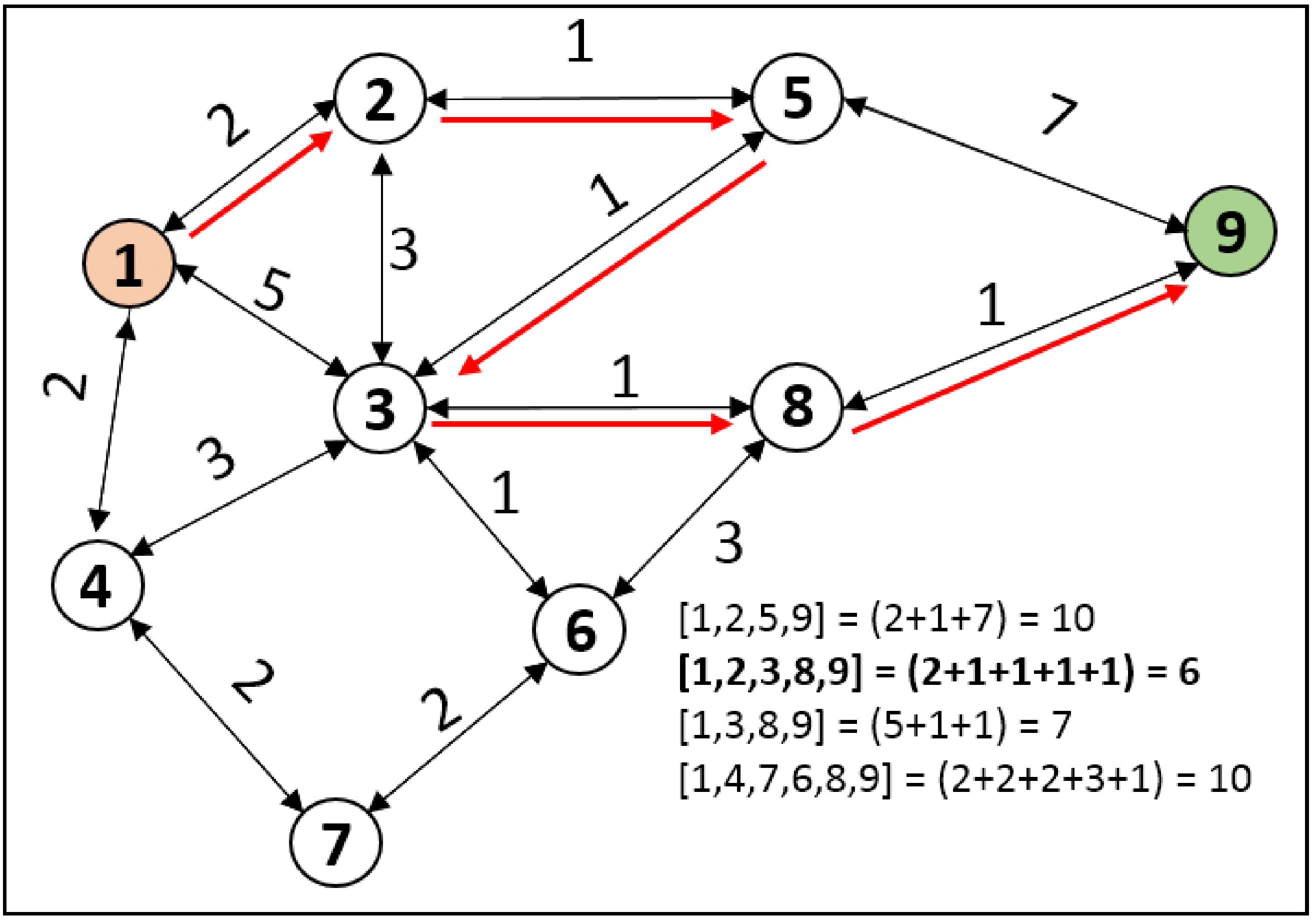

Figure 4.1. Illustration of Dijkstra's algorithm (Ur Rehman, Awuah-Offei, Baker, & Bristow, 2019).

Adoption of a routing algorithm could be as follows:

1) Treating the 64 different intervention combinations as system states, working through the 10 timesteps and identifying the highest NPV route to each system state at each timestep. Pathways which are 'orphaned' at a timestep (as they are not the highest NPV pathway to a particular system state) are rejected, as shown in Figure 4.2.

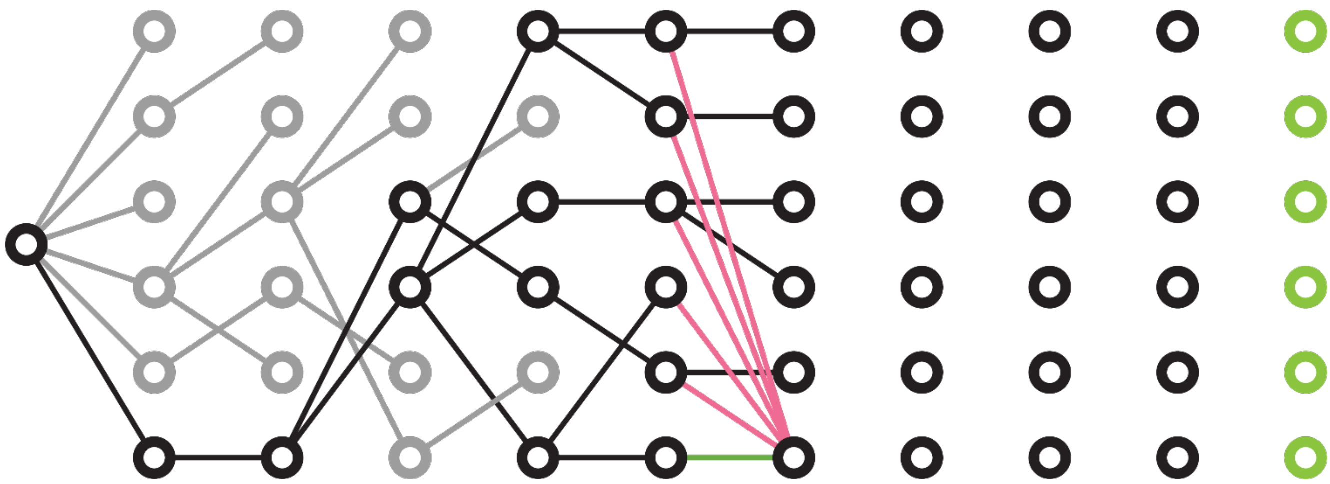

Figure 4.2. Routing algorithm: Iterating through timesteps.

2) The culmination of these iterations is the highest NPV route to each system state at the end of the study period, as shown in Figure 4.3.

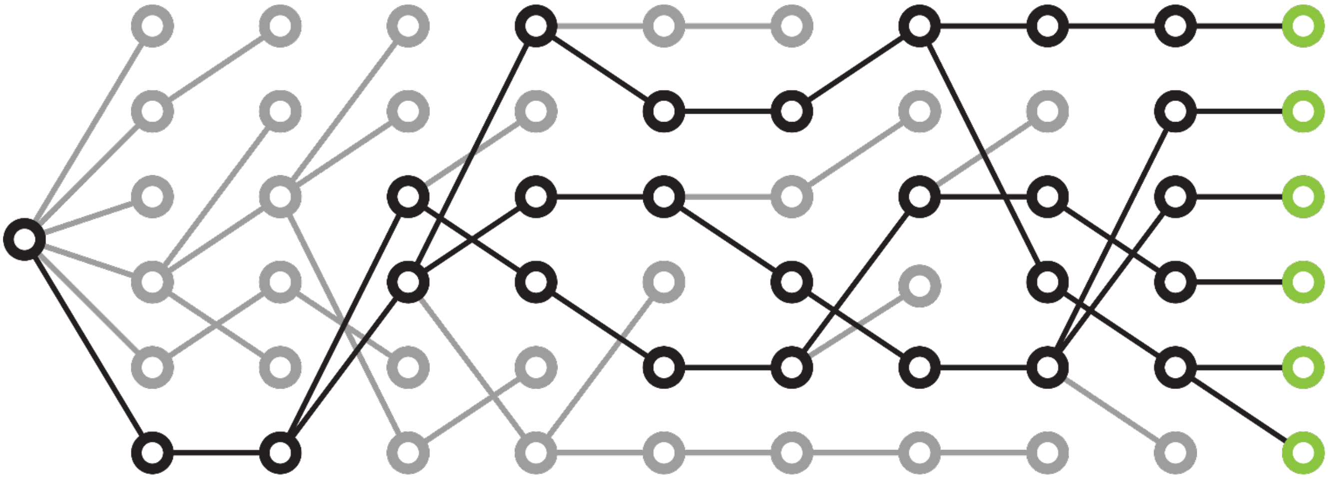

Figure 4.3. Routing algorithm: Highest NPV route to each system state.

3) From these, the highest overall NPV route can be identified.

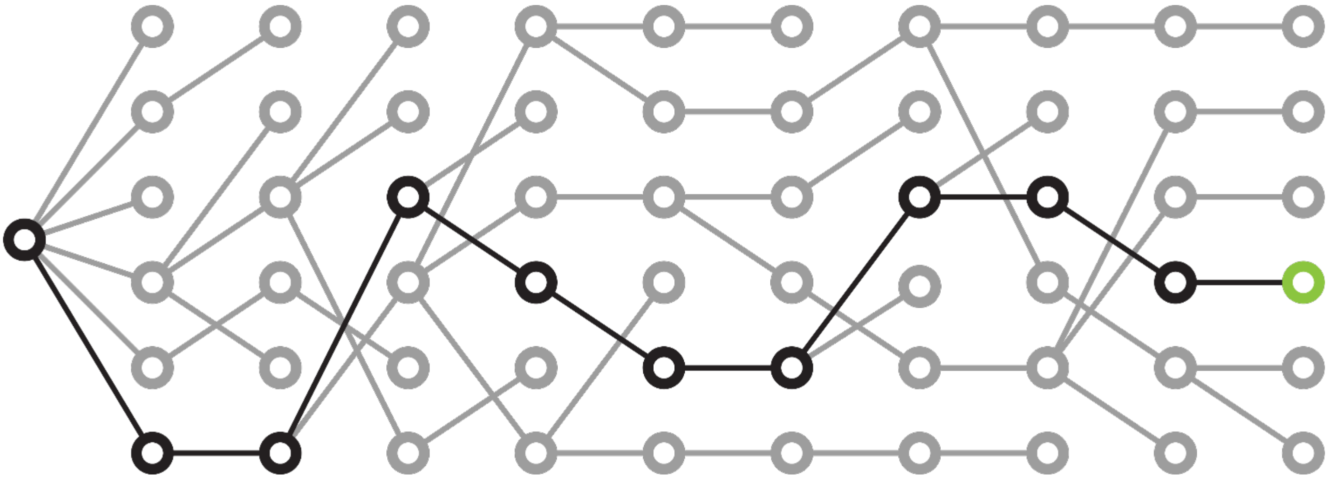

Figure 4.4. Routing algorithm: Highest NPV pathway.

The main issue with this approach is that a system state at a given point in time may not be reflective of the whole life value gained from that system state, for example because of different remaining lifespans when that system state is reached at different points in time. In our testing, we found the results from this routing algorithm were lower NPVs than from other approaches we explored, suggesting that this algorithm was not finding a true optimum.

4.1.4 Genetic algorithm

A genetic algorithm applies the principles of Darwinian natural selection (survival of the fittest) to finding the fittest (i.e. optimum) solution (Mallawaarachchi, 2017):

1) **Initial population**: Characterise each member of the population by a set of parameters known as 'genes' with a given set of values that form a 'chromosome'. For this project, parameters could be defined that describe the level of investment in each intervention, and the timing of investment in each intervention, for all the interventions in a catchment. For a real options analysis, parameters would be needed to describe the different investment choices that could be made under different futures.

2) **Fitness function**: This determines how 'fit' an individual is – which generates a fitness score for each member of the population. For OxCam, the fitness function would be the whole life net present value (under multiple futures for the real options analysis).

3) **Selection**: Pairs of members of the population are selected (Parent A and Parent B) – with high 'fitness' individuals being more likely to be selected for reproduction – to pass their 'genes' to the next generation. The population has a fixed size, so individuals with the least fitness (lowest net present value) are rejected.
4) **Crossover**: The crossover point dictates which genes are taken from Parent A and which are taken from Parent B. Children are created by exchanging the genes.
5) **Mutation**: At random, genes might be 'mutated' – randomly selecting a value that did not come from its parents. This aims to ensure that the algorithm doesn't converge on a local optimum by maintaining diversity in the population.
6) **Convergence**: After many iterations, the genetic algorithm process reaches a point where the best fitness score does not change significantly between generations.

We first explored the use of an off-the-shelf genetic algorithm solution, defining a series of variables representing the level and timing of investment, as described above. Because of the large numbers of variables required to describe the levels and timing (and different levels at different points in time) across the large number of interventions in a single catchment, we found that the genetic algorithm struggled to converge on a solution – and we believe it converged on a local optimum, i.e. a value from which individual gene changes do not improve net present value, but which is not a true optimum.

We therefore developed a more bespoke solution based on a genetic algorithm. Key differences included:

- Rather than using a simple list of variables, we developed a two-dimensional definition of the set of interventions that make up a pathway, i.e.
  - For each type of intervention (defences, storage, etc.):
    - For each location (relevant to the type of intervention):
      - Intervention level at year 0.
      - For each future scenario:
        - Intervention level at future year.
        - Future year at which intervention level applies.
- As part of the definition of interventions, we implemented code to ensure that the logical sequencing of interventions was always followed.
- Modified genetic algorithm logic, to constrain and guide the analysis based on our knowledge of the problem.

This exploration ultimately led us to the adopted optimisation routine, described below.

4.2 Adopted optimisation routine

Our adopted optimisation approach builds on the findings from exploring the approaches described above, and follows an iterative process. Steps 1 to 3 are designed to gradually converge on the optimal level and timing of investment by gradually increasing the complexity of calculation – at each point switching to a calculation which is lengthier, but from a starting point which is closer to the true optimum:

1) Explore each available location for each type of intervention in a catchment in turn:
   a) Using a fixed set of 3 intervention points (year 0, year 30 and year 60), identifying **all** combinations of investment that could be selected for this intervention at this location. For example, for linear defences at a given flood area, the options would be:

Year 0	Year 30	Year 60
No new defences	No new defences	No new defences
No new defences	No new defences	Small defences

No new defences	No new defences	Medium defences
No new defences	No new defences	Large defences
No new defences	Small defences	Small defences
No new defences	Small defences	Medium defences
No new defences	Small defences	Large defences
No new defences	Medium defences	Medium defences
No new defences	Medium defences	Large defences
No new defences	Large defences	Large defences
Small defences	Small defences	Small defences
Small defences	Small defences	Medium defences
Small defences	Small defences	Large defences
Small defences	Medium defences	Medium defences
Small defences	Medium defences	Large defences
Medium defences	Medium defences	Medium defences
Medium defences	Medium defences	Large defences
Medium defences	Large defences	Large defences
Large defences	Large defences	Large defences

b) Quantifying the net present value of each pathway for each intervention, and selecting the highest net present value option.

The outputs of step 1 are an optimal level and timing of investment in each intervention and location **in isolation**. This serves as an **upper limit** of investment – we assume that when the full portfolio of interventions is considered, investment in some interventions will no longer be cost beneficial (or that, at most, the same level of investment is cost beneficial). For example, with storage in place, the benefits (and therefore net present value) of individual linear defences is assumed to be reduced.

2) Iteratively explore the complete catchment system, calculating present value benefits **without capping and write-off**. The capping and write-off calculation requires property-level impact data to be retained throughout the present value benefit calculation.

a) Year 0 investment:

- For each type of intervention, test whether reducing investment in year 0 increases the net present value.
- For each type of intervention, add together the combined benefits of reducing investment (in locations where reducing investment increases the net present value).
- Compare the total catchment-scale net present value improvement across the different types of intervention, and 'accept' the intervention with the greatest improvement.
- Repeat this process over a number of 'mini iterations' – gradually finding additional net present value improvements.

b) Future years investment **levels**:

- For each type of intervention, for each future, test whether reducing investment in each future year increases the net present value.
- As above, add together and accept the best improvement, and repeat the process.

c) Future years investment **timing**:

- For each type of intervention, for each future and for each future decision point, test whether bringing investment forward by 1 decision point or moving investment back by 1 decision point, increases the net present value.

d) Repeat this full process until no more improvement can be found.

Capping and write-off is assumed to **always** reduce (or have no effect on) present value do nothing impacts – and therefore always reduce benefits. Outputs from step 3 are therefore assumed to serve as an **upper limit** of investment – we assume that when capping and write-off is introduced and benefits decrease, investment in some interventions will no longer be cost beneficial (or that, at most, the same level of investment is cost beneficial).

3) Iteratively explore the complete catchment system, calculating present value benefits **with capping and write-off**, following steps a) to d) described above.

4) Finally, extracting the full results profile (with capping and write-off) to write the final results.

This optimisation routine is repeated:

1) For each catchment.

2) For each set of futures considered:

   a) 'No development' real options analysis, using 3 climate change futures but with no future development.

   b) 'Full' real options analysis, using 18 futures (2 development rates x 3 development 'shapes' x 3 climate change futures).

   c) 27 individual future scenarios, with investment optimised 'per future' rather than the real options analysis approach described above.

The process is also repeated for river flooding and surface water flooding – with the surface water analysis optimising investment in each flood risk management system, rather than the flood area <> sub-catchment <> catchment hierarchy used for river flooding.

5. Method implementation

5.1 Scenarios

As described in the modelling and hydrology technical report and economics technical report, we have explored a range of scenarios:

- 3 climate change scenarios (central, upper and H++).
- 3 development rates (23k, 30k and 43k homes per year) and 3 development 'shapes' (expansion of existing settlements, creation of new settlements and a hybrid of the two), plus a no development scenario.

We have combined these scenarios in different configuration to form sets of futures:

- A real options analysis, optimising investment across 18 futures, providing the core set of results (described as the 'all futures' analysis), using 3 climate change scenarios, 3 development shapes and 2 development rates (excluding the 43k homes per year as it was deemed to be unrealistic). Each of the 18 futures is assumed to be equally likely.
- A real options analysis, optimising investment across 3 no development futures (for the 3 climate change futures), to understand the level of investment that would be needed to manage risk to existing properties. Each of the 3 futures is assumed to be equally likely.
- 27 separate optimisation outputs for investment across all the individual futures (not including the no development scenarios).

These form 29 separate 'runs' of the optimisation analysis.

5.2 Implementation in code

We implemented the method described above as a Python script that can be called with a catchment ID (1 to 4) to run the analysis for all sets of future scenarios in turn:

1) Define the sets of future scenarios to be explored.
2) Getting the 'best pathway' for each intervention in turn. The appropriate scale for each type of intervention (catchment for storage, sub-catchment for natural flood management, flood area for linear defences and property flood resilience) is hard-coded, but the spatial relationships between catchments, sub-catchments and flood areas is defined in the database, allowing for flexibility in implementation.
3) Building a data object describing the 'best pathway' for each intervention at each location, with the pathway (initially with interventions at year 0, 30 and 60) defined for each location.
4) Starting the full iterative process described above. A single 'iterate' function describes the routine – with 'pass 1' excluding capping and write-off and 'pass 2' including it. Within in iteration, various shifts in the pathway for each intervention are explored – as described above, it looks first at interventions in year 0, then at the level of investment in interventions in future years, then at the timing of investment in future years. This function operates a) at a full catchment scale, and b) based on each location for each intervention – it is therefore deliberately unaware of the flood areas which benefit from a given intervention – this is obtained from the 'get_npv' function described below.
5) The net present value of each explored shift in investment is calculated by a 'get_npv' function. This function – like the steps describes above – operates at a full catchment scale. The 'get_npv' function:
    a) Splits out the pathways for each intervention to describe the portfolio of interventions adopted in every flood area.
    b) Triggers a multi-core process that iterates through all the flood areas to calculate the present value benefits of the proposed interventions **under each future**. This is the most computationally expensive part of the algorithm, so we spent time:

- Improving the processes within the underling 'CustomScenario' present value impact calculation (see below).
- Developing an approach to storing pre-calculated impact values for a given pathway at a given flood area for a given future, checking for existing values and using them when they are available. This element was particularly important because of the nature of the iterative calculation – for example, when exploring whether linear intervention investment at one particular location should increase or decrease, the impacts in all but one flood area will be unchanged and can simply be looked up.
- Developing the code in a way that it worked across multiple processing threads, and that it switched between single-thread and multi-thread processing at the right times. The main complexity here is that there is a ~7 second delay when 'spinning out' additional processing threads – so without careful management, it could actually *increase* processing time.

c) Brings together the results of each individual flood area to obtain a catchment-wide present value benefit a) in each future and b) weighted across all futures.

d) Loops through each intervention at each location **for each future** and calls the appropriate cost function to obtain the present value cost for each intervention, and then brings them together to obtain a catchment-wide present value cost a) in each future and b) weighted across all futures.

e) Returns the net present value by subtracting the present value cost from the present value benefit for each future and weighted across all futures.

The net present value of each explored variation is compared with the previous iteration result to determine whether the change has a positive, negative or neutral effect on net present value.

5.3 Relationship with economic calculations

The optimisation routine acts as a 'wrapper' around the so-called 'CustomScenario' Python code we developed to extract annual average and present value impacts. This code is described in more detail in Chapter 5 of the economics technical report. The optimisation code then provides the logic to explore different pathways – calling the CustomScenario code to obtain the present value impacts for a candidate pathway and the baseline present value impacts to calculate the present value benefit of the pathway. The iterations described above are achieved through a series of parameters passed to the method – in particular to dictate whether capping and write-off is applied, but also whether to return the full profile of impacts.

5.4 Costs

The optimisation routine also includes an implementation of the cost functions (described in Chapter 2 of the economics technical report). Each intervention has a cost function which accepts the relevant parameters as inputs. For linear interventions, costs are calculated based on the defence height or defence raising required **per asset** combined with the asset length and cost model. For other interventions, costs are calculated at the relevant scale (number of properties for property flood resilience, volume of equivalent storage in a sub-catchment for natural flood management, volume of storage in a catchment for storage). The cost functions take a general form:

For each year in the study:

1) Is there a new (or larger) intervention? If so, calculate the capital cost and capital carbon cost, based on the **additional** capacity provided. For example, moving from 'small' to 'medium' storage incurs a capital cost based on the cost of a storage scheme that provides the **additional** storage volume offered by 'medium' over 'small' storage in a given catchment. Because the cost functions are not generally linear, this means that the capital cost of building a 'medium' storage may be different to the capital cost of building a 'small' storage, then increasing storage capacity to match the 'medium'.
2) Have the existing interventions reached their end-of-life? If so, calculate the capital replacement cost.

3) Calculate annual maintenance cost and maintenance carbon cost, based on the current total scale of the interventions (of each type).

4) Apply the relevant discount rate for this year to add the year's costs to the total present value cost.

5.5 Distribution of computing

As well as parallelisation within the code, which distribute – to calculate present value impacts for each flood area as part of the 'get_npv' calculation described above – we also distributed the optimisation across multiple processing machines, using each machine to run the analysis for a given catchment. The main reason we distributed the analysis in this way was that different sets of futures share common sets of results, and the analysis stores the calculated impacts for each future – so distributing the analysis in this way enables the analysis to re-use these stored outputs.

The modelling automation used our Flood Platform computing cluster to distribute the flood simulations across our suite of modelling machines. But Flood Platform is designed around a centralised database. This centralised database formed the basis of the OxCam project results database, including the pre-calculated impacts associated with each simulation. Extracting the full body of data for a given flood area requires extensive data transfer – in particular because we have to retain property-level impact information to enable the capping and write-off calculation. To enable the distributed optimisation, we mirrored the full database to each processing machine to minimise the data transfer and reduce the load on a single central database.

6. Results

6.1 Baseline risk

6.1.1 Introduction

The results in the baseline risk section describe the level of flood risk from river flooding and surface water flooding **without** new investment in interventions, under different scenarios – and exploring present value risk and risk profiles.

6.1.2 Present value baseline risk

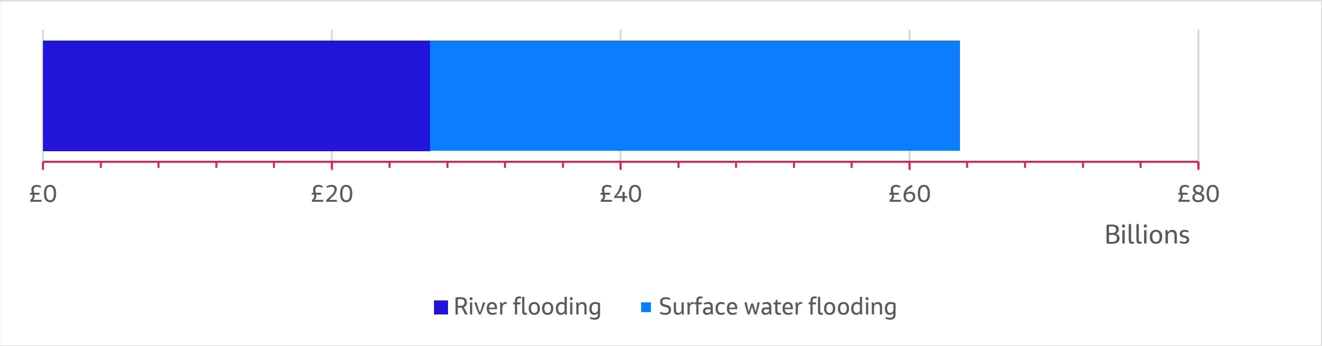

Figure 6.1. Present value baseline risk.

As indicated by Figure 6.1, the present value baseline risk (over 100 years) for the OxCam Arc is £63.5 billion[5]. River flooding contributes 42% (£26.8 billion), and surface water contributes 58% (£36.7 billion).

Figure 6.2 shows the breakdown of present value baseline risk between catchments, and Figure 6.3 breaks those values down further – showing baseline risk by catchment, divided by source.

[5] Thousand million.

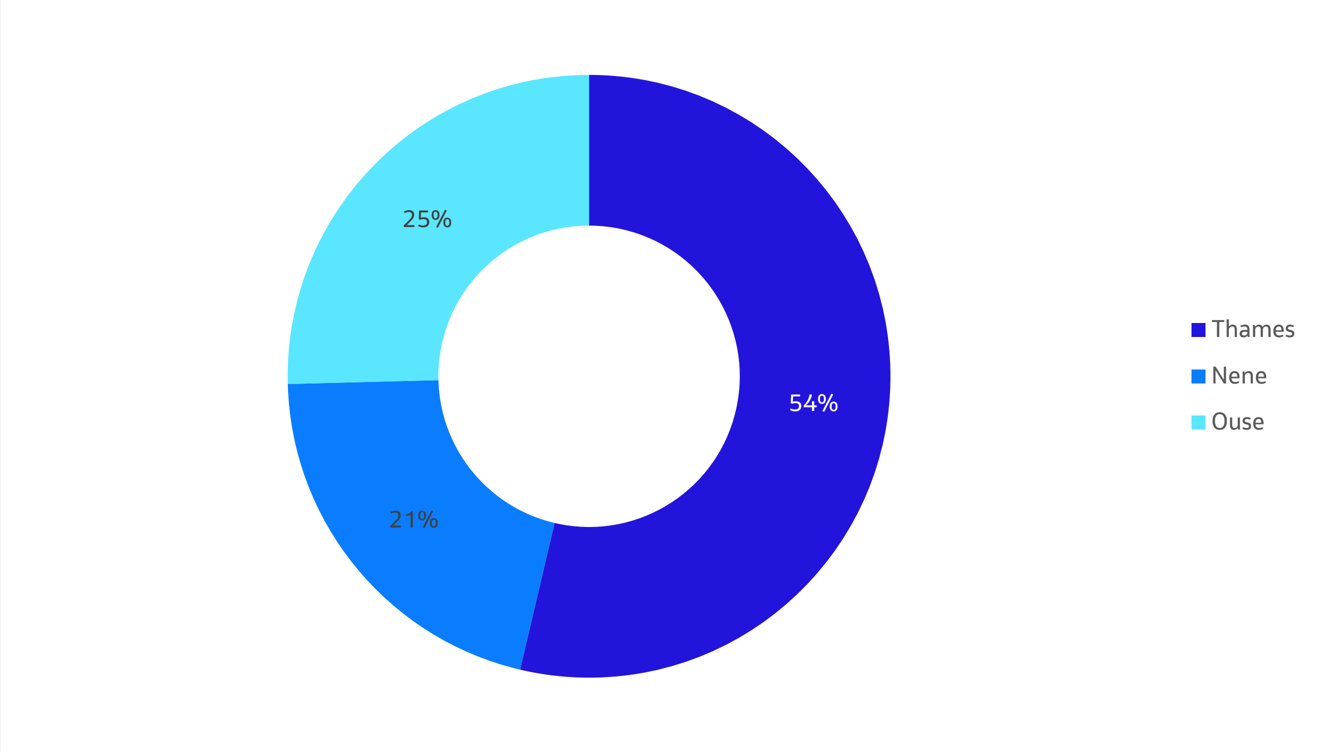

Figure 6.2. Breakdown of present value baseline risk between catchments.

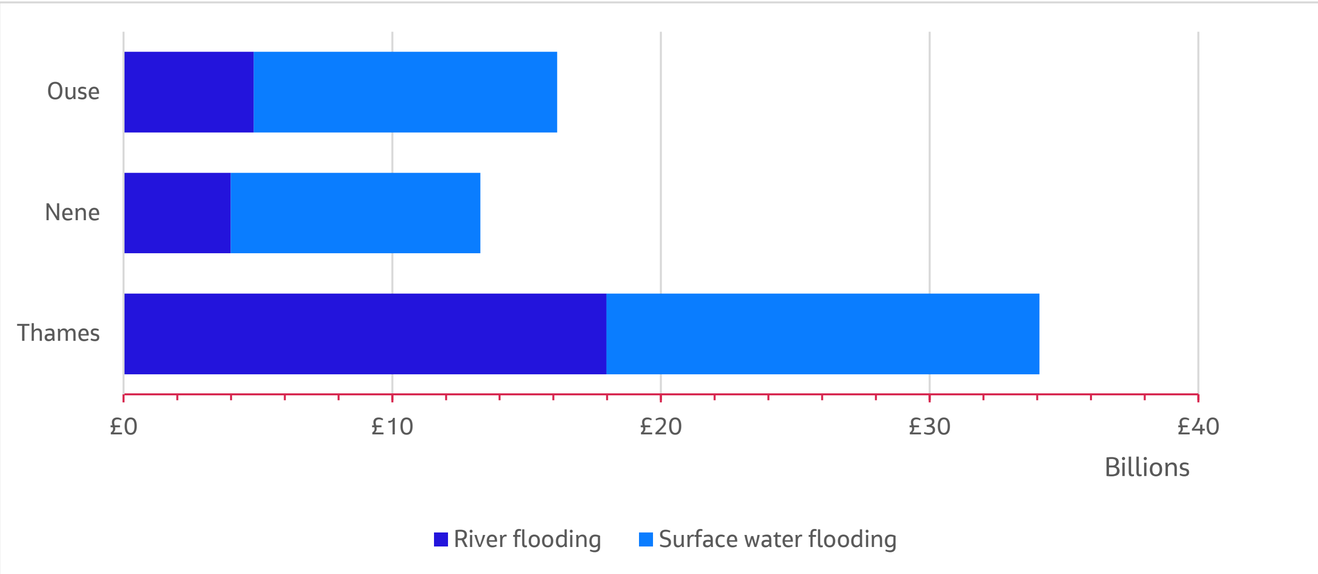

Figure 6.3. Breakdown of present value baseline risk between catchments and sources.

6.1.3 Risk in different futures

Figure 6.4 shows how the baseline present value risk varies under different futures. Each bar on the chart shows a different development scenario, and the different 'chunks' of the bar show a) the level of risk in the central climate change scenario, and b) the additional risk for other climate change scenarios.

Figure 6.5 also shows how the baseline present value risk varies under different futures, but each bar shows a different climate change scenario plus shape of development. The different chunks of the bar show a) the level of risk for the 23k homes per year scenario, and b) the additional risk for more rapid rates of development.

Comparing the two figures indicates that there is more variability between climate change scenarios than there is between rates of development.

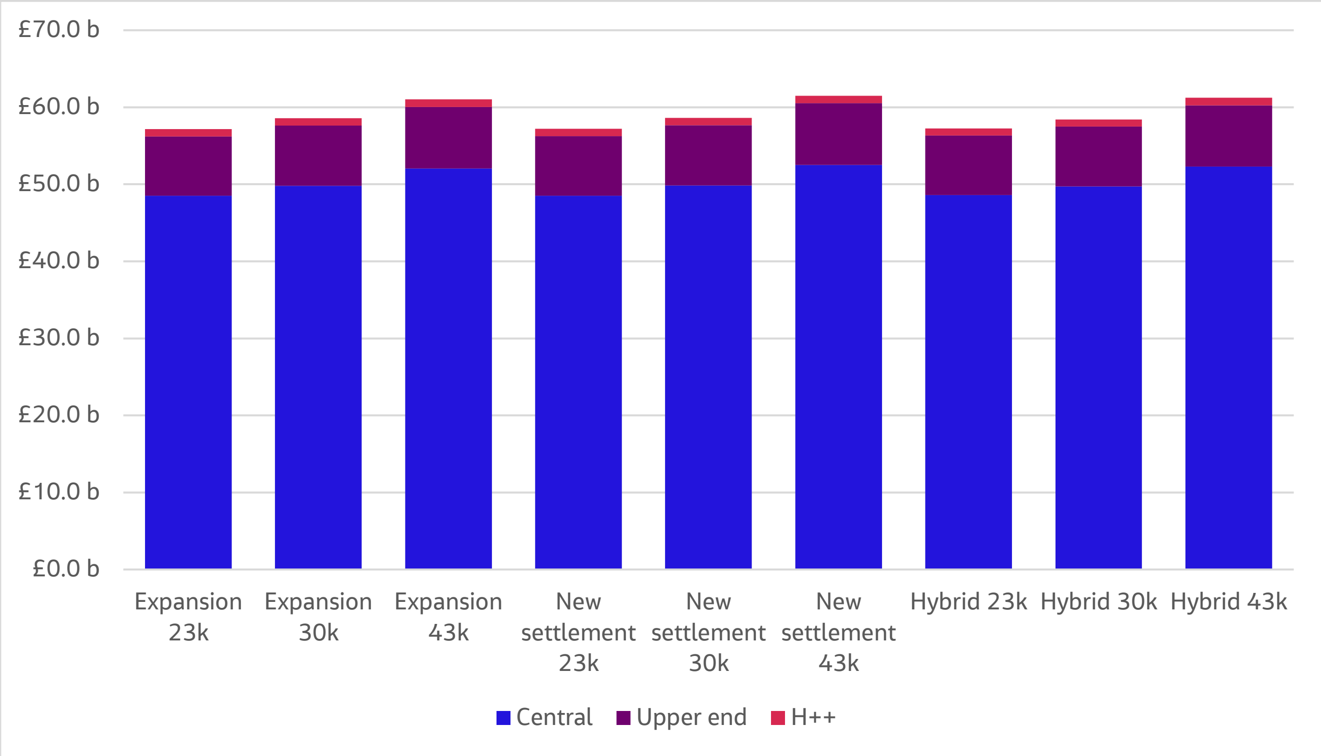

Figure 6.4. Baseline present value risk under different futures, colour-coded by climate change scenario.

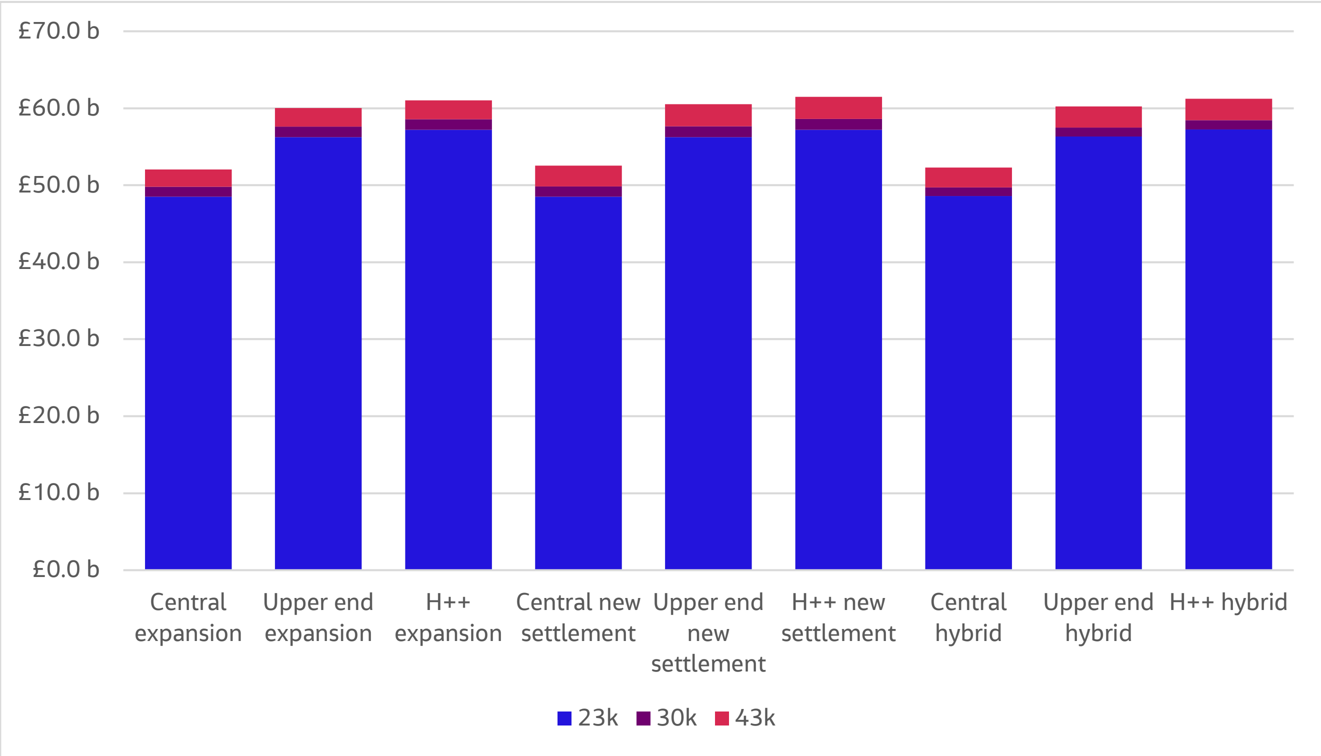

Figure 6.5. Baseline present value risk under different futures, colour-coded by rate of development.

6.1.4 Baseline risk profile

Figure 6.6 shows how the range of possible risk profiles under different future scenarios. The coloured bands show different parts of the distribution of results – with purple representing results above the median, and pink below the median. The grey line shows the mean risk profile. Figure 6.7 shows the same profile but for **all** futures.

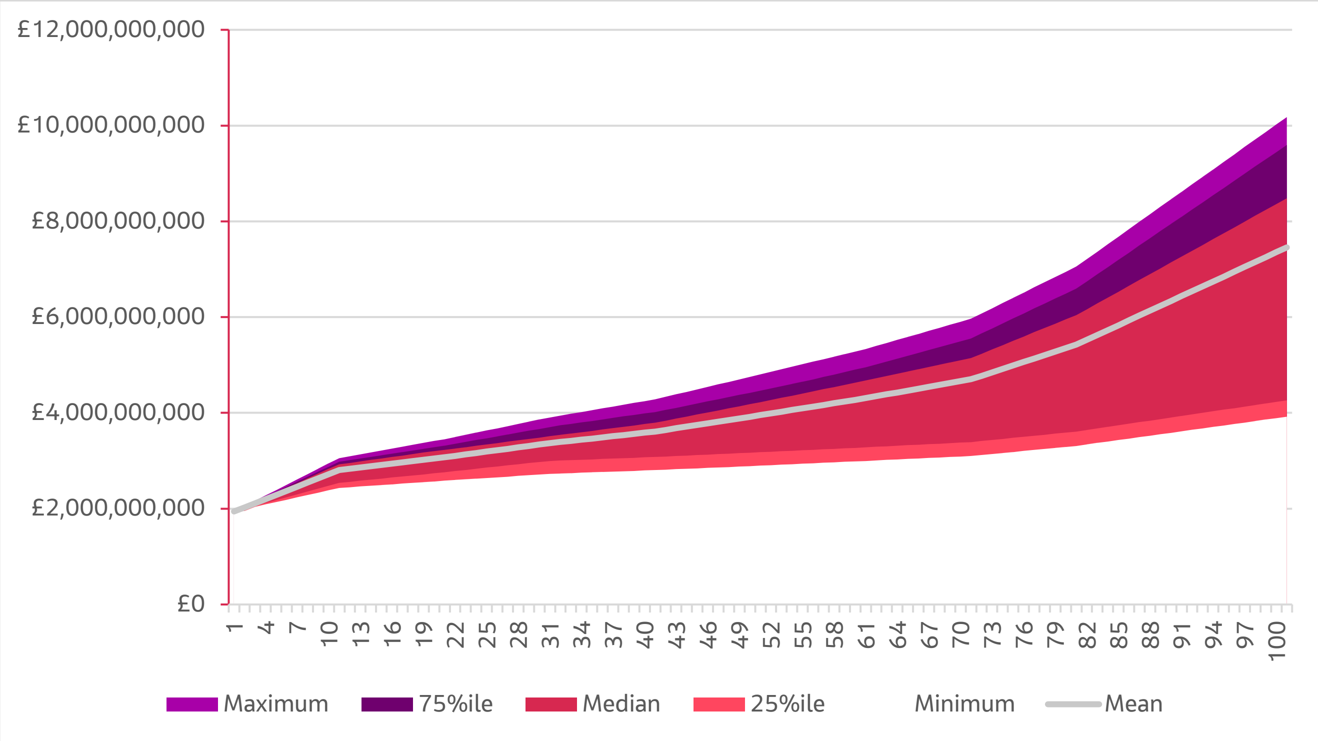

Figure 6.6. Risk profile range under different future scenarios.

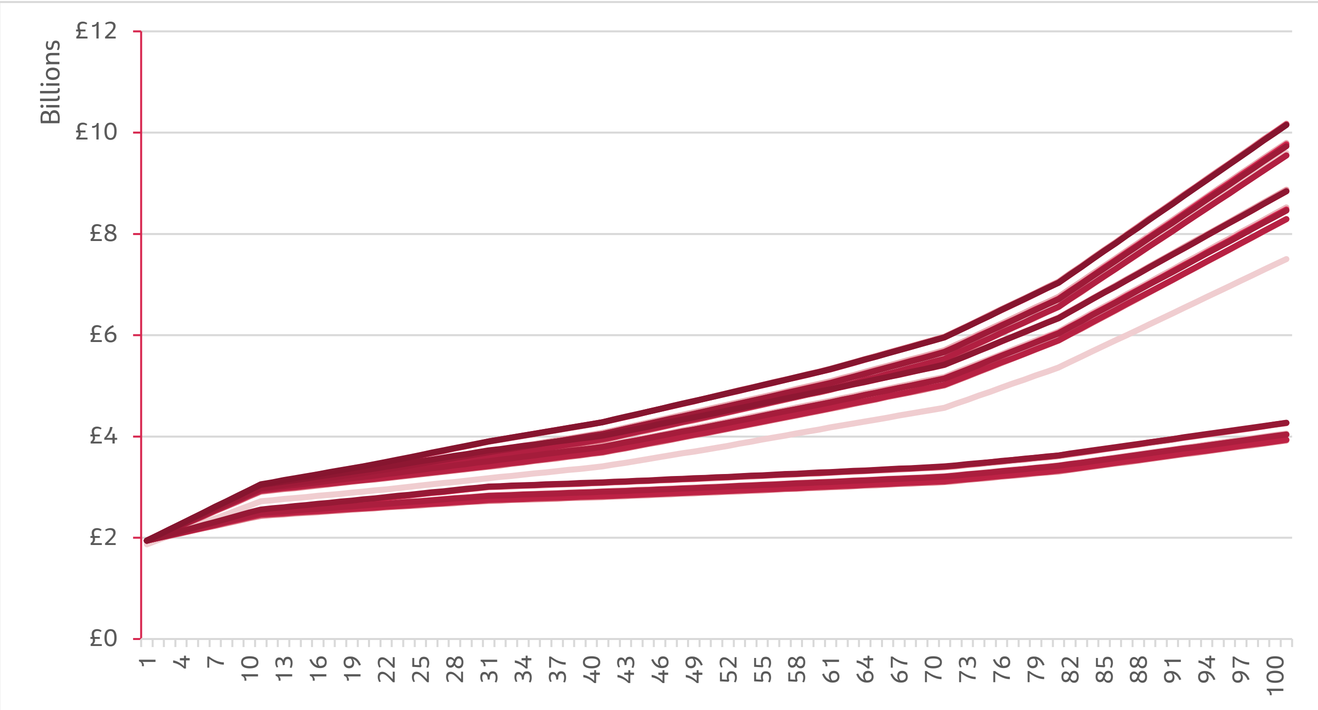

Figure 6.7. Risk profiles under different future scenarios.

6.1.5 Baseline risk by impact category

Figure 6.8 shows the breakdown of baseline risk by impact category. Table 6.1 breaks down these results further by listing baseline risk by impact category and by source.

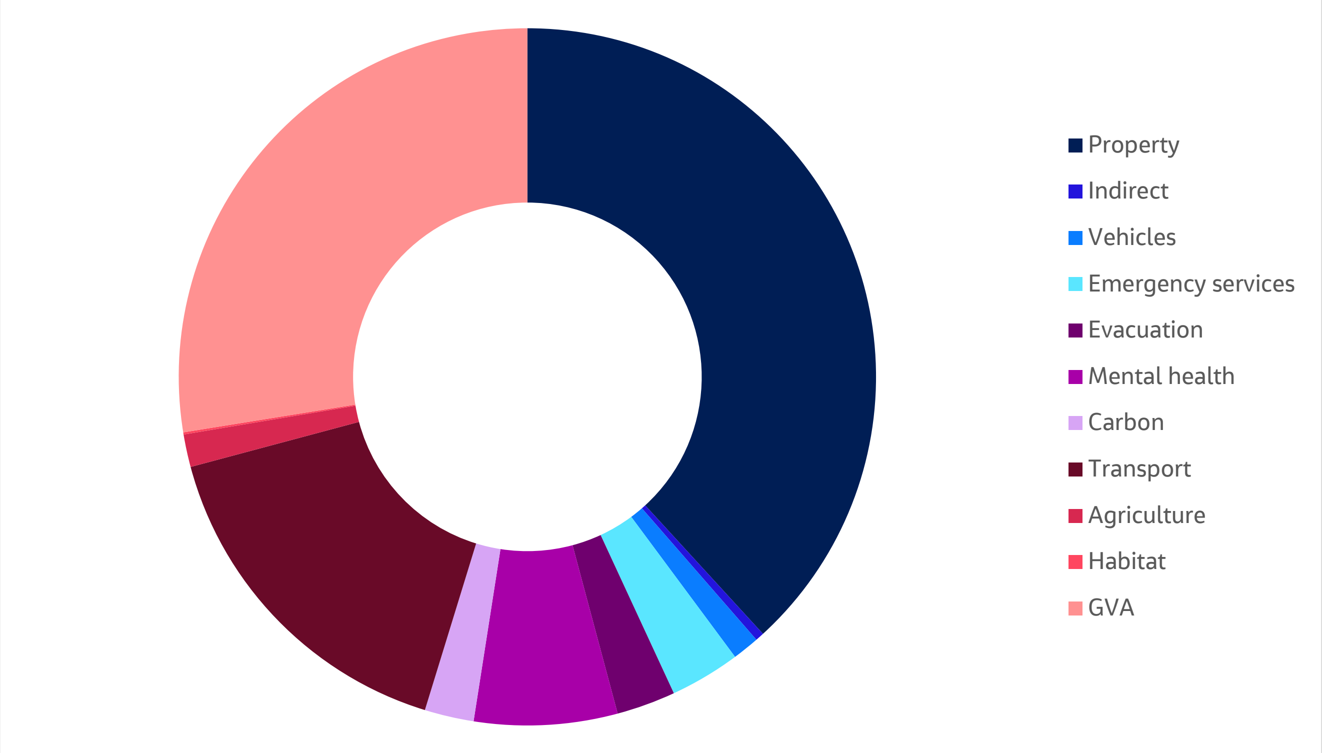

Figure 6.8. Present value baseline risk by impact category.

Table 6.1. Present value baseline risk by impact category and source.

	Rivers	Surface water	Total
Property	£15,364 m	£11,297 m	£26,661 m
Indirect	£139 m	£143 m	£282 m
Vehicles	£408 m	£464 m	£872 m
Emergency services	£1,193 m	£1,081 m	£2,273 m
Evacuation	£847 m	£1,060 m	£1,906 m
Mental health	£1,723 m	£2,909 m	£4,633 m
Carbon	£738 m	£856 m	£1,594 m
Transport	£2,261 m	£8,970 m	£11,231 m
Agriculture	£1,003 m	£55 m	£1,058 m
Habitat	£14 m	£56 m	£70 m
GVA	£5,913 m	£13,316 m	£19,229 m
Total	£29,603 m	£40,206 m	£26,661 m

6.1.7 Spatial distribution of risk

Visualising results across the OxCam Arc is challenging – with large differences in absolute values (of damage and benefits) between analysis units (flood areas for river flooding and flood risk management systems for surface water flooding). We have therefore focussed on **ranking** of analysis units and on **relative** changes in risk – however, ranking areas based on absolute risk values will still produce results which are biased by analysis

Figure 6.9 and Figure 6.10 show the spatial distribution of baseline risk for river flooding (by flood area) and surface water flooding (by flood risk management system) respectively, where 1 equals the highest absolute baseline risk value.

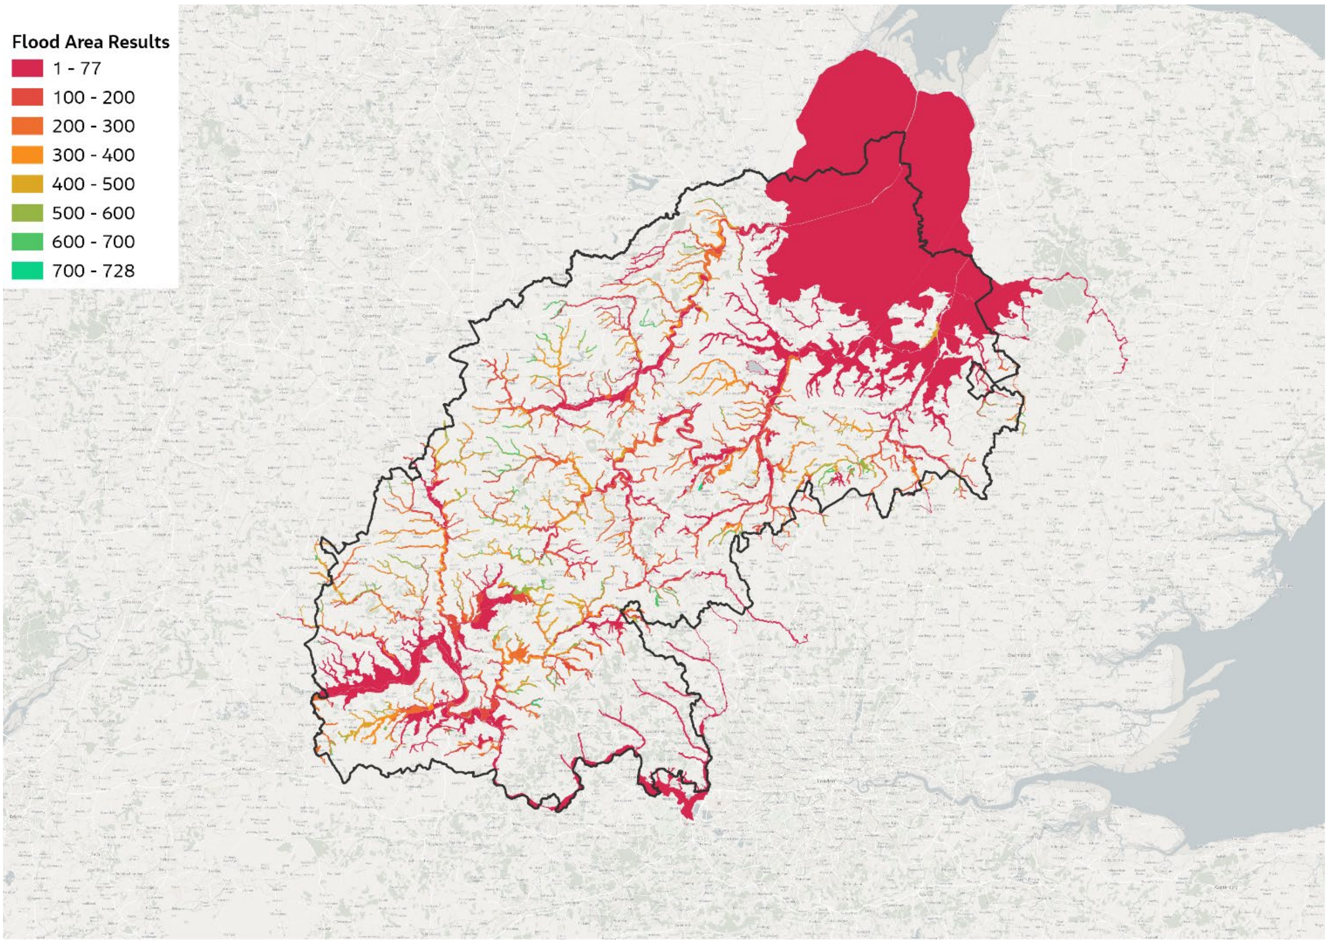

Figure 6.9. Baseline risk from river flooding by flood area, ranked from highest (1) to lowest (728).

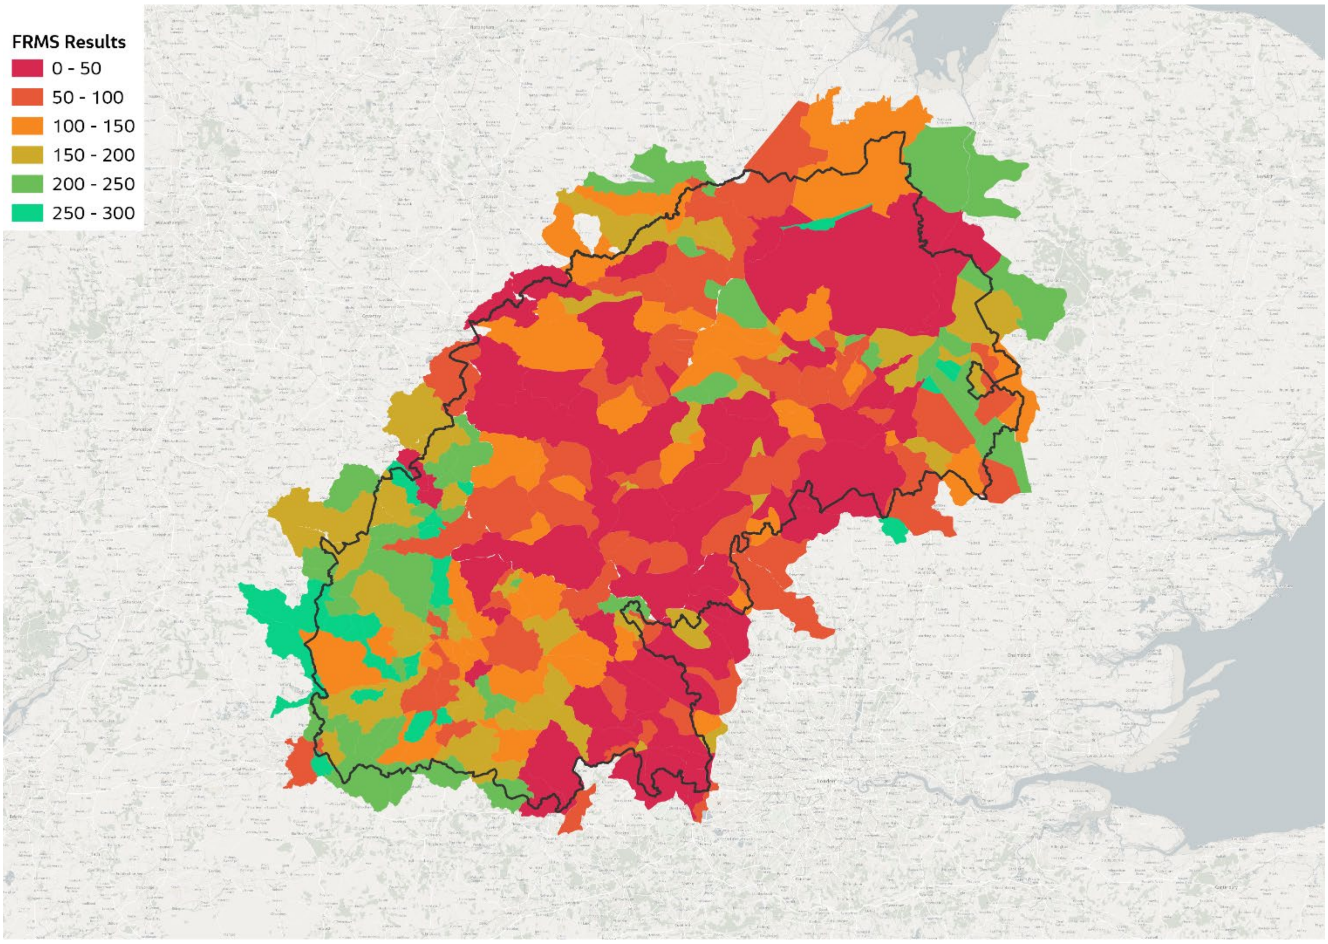

Figure 6.10. Baseline risk from surface water flooding by flood risk management system, ranked from highest (1) to lowest (280).

6.2 All futures analysis

6.2.1 Introduction

The results in this section explore the 'all futures' analysis. This is the full real options analysis described in Section 5.1, including all permutations of 3 sets of climate change scenarios (central, upper and H++), 3 shapes of development (expansion of existing settlements, development of new settlements, and a hybrid of the two) and 2 rates of development (23,000 homes per year and 30,000 homes per year).

6.2.2 Summary results

The headline results of this analysis are the present value 'whole life' figures, taken as the expected outcome (averaged across the 18 possible futures) for the baseline risk, the optimum level of investment, the benefits of that investment, residual risk and the net present value. Table 6.2 and Figure 6.11 show the summary present value results for the all futures analysis.

Table 6.2. Summary present value results.

	Cost (£m PV)	Baseline risk (£m PV)	Benefit (£m PV)	Residual risk (£m PV)	Net present value (£m PV)
Rivers	£4,784m	£26,816m	£22,845m	£3,971m	£18,060m
Surface water	£849m	£36,684m	£5,364m	£31,320m	£4,515m
Total	**£5,634m**	**£63,500m**	**£28,209m**	**£35,291m**	**£22,575m**

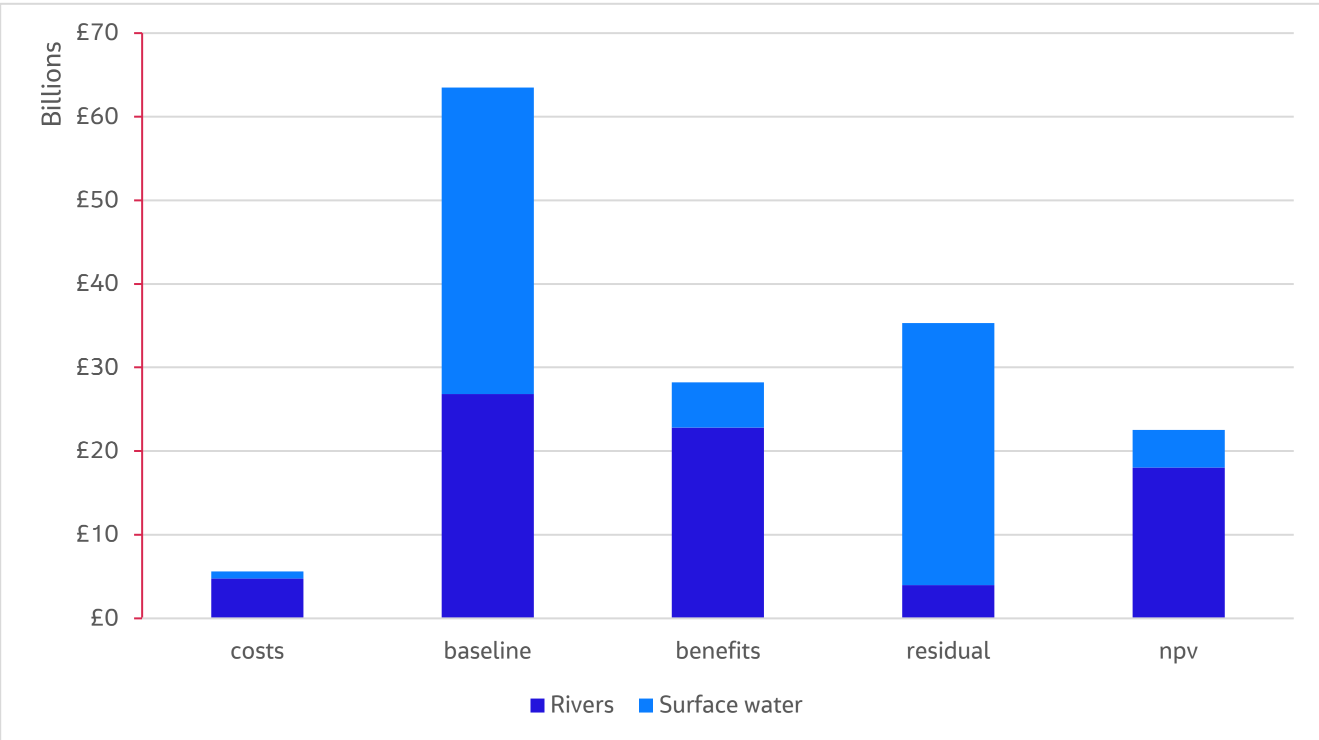

Figure 6.11. Summary present value results.

6.2.3 Investment by future scenario

The results in this section explore how results vary across the range of future scenarios. In this case, because the all futures analysis is a real options analysis, the action taken in year 0 is the same irrespective of the future decisions that are made as a result of emerging future risk – so differences in investment and benefits are the result of decisions made in future years.

Table 6.3. Distribution of present vale results across the range of future scenarios.

	Costs	Baseline	Benefits	Residual	NPV
Minimum	£4,629 m	£48,506 m	£16,956 m	£31,022 m	£13,202 m
25%ile	£4,753 m	£49,780 m	£17,522 m	£32,102 m	£13,657 m
Median	£6,027 m	£56,752 m	£23,152 m	£33,493 m	£18,958 m
75%ile	£6,110 m	£57,635 m	£23,766 m	£34,151 m	£19,517 m
Maximum	£6,195 m	£58,635 m	£24,468 m	£34,670 m	£20,136 m
Mean	£5,634 m	£54,662 m	£21,499 m	£33,163 m	£17,398 m

Table 6.3 shows the distribution of results across the range of future scenarios that make up the all futures analysis, and Figure 6.12 shows the range of present value investment across the futures – also broken down by source.

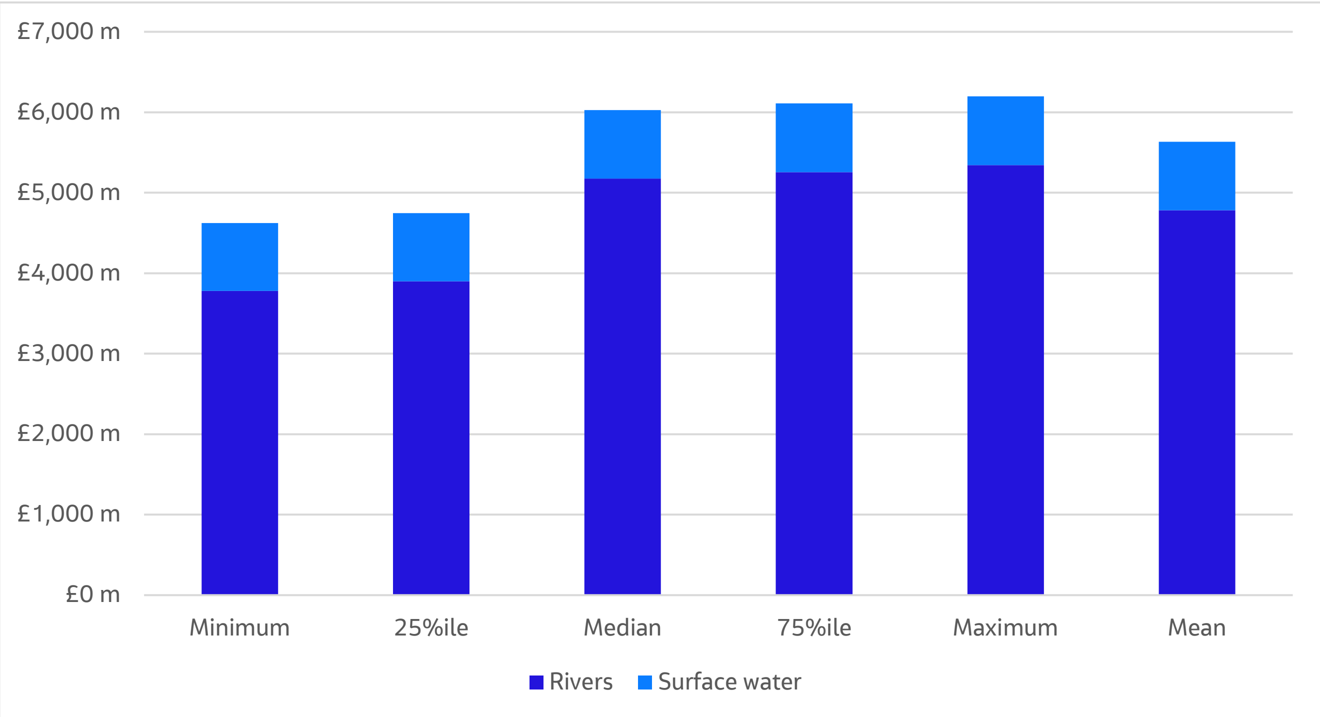

Figure 6.12. Distribution of present value investment across the range of future scenarios, broken down by source.

6.2.5 Investment by catchment

The analysis produced results for 4 separate areas of the OxCam Arc, which align with the Thames, Nene and 2 parts of the Ouse, which are based on the scale at which we have considered catchment-scale storage options (with sub-catchments for NFM and flood areas for linear defences and PFR being sub-units within those 4 areas). The 2 parts of the Ouse have been combined to present results for the 3 catchments.

Table 6.4. Present value investment by catchment.

	River flooding	Surface water flooding	Total investment
Thames	£17,978 million	£16,106 million	£34,084 million
Nene	£3,994 million	£9,287 million	£13,281 million
Ouse	£4,844 million	£11,291 million	£16,135 million

Table 6.4 and Figure 6.13 show the breakdown of investment between sources and between catchments, showing that 54% of investment is in the Thames catchment, 21% in the Nene and 25% in the Ouse.

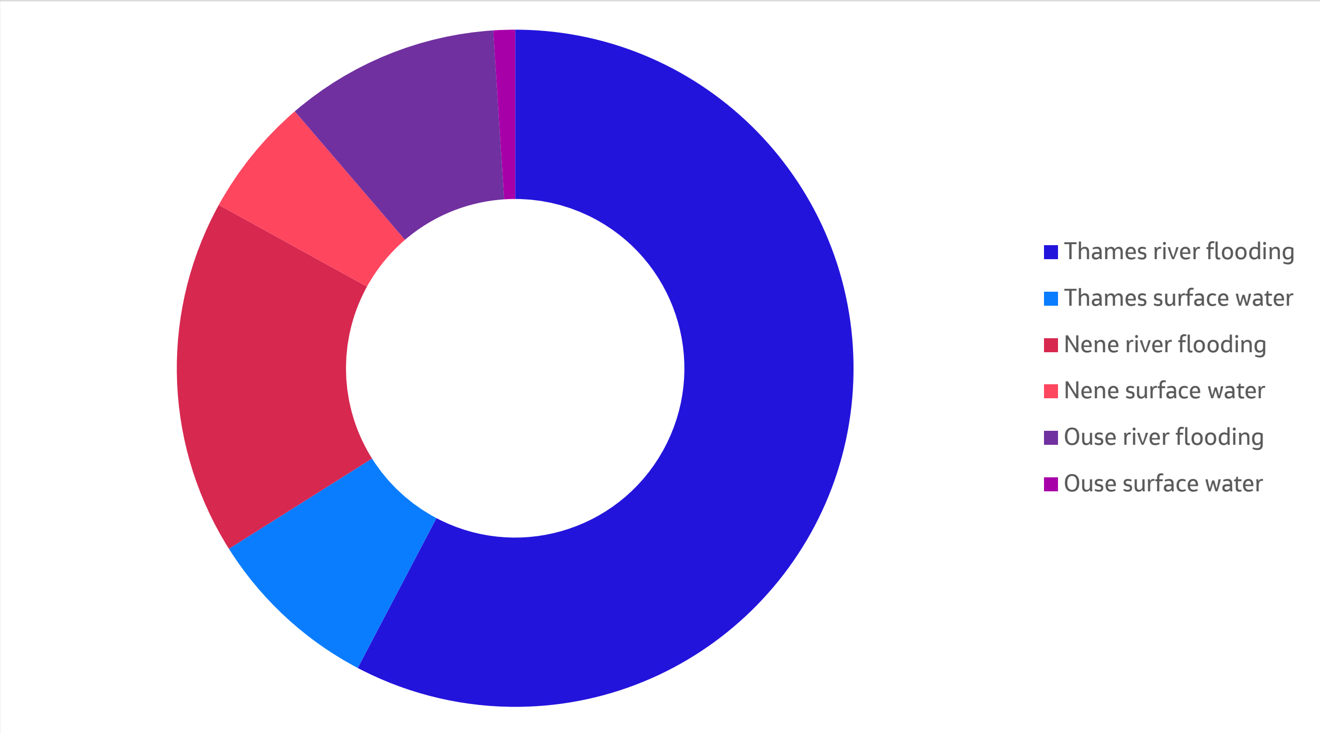

Figure 6.13. Present value investment by catchment.

6.2.7 Investment by intervention

The definition of the portfolio of interventions for this project is an important part of the optimisation analysis – providing us with a useful selection of potential intervention options of different kinds. However, their primary purpose is to give us sufficient confidence in the overall optimum level of investment, rather than serving as a guide for which interventions are likely to be most viable in certain locations – because that is dependent on many local considerations that cannot be explored at this scale.

With that context in mind, it is still useful to explore the results in this detailed technical report – and from these results draw high-level conclusions (which are explored in the summary report). Figure 6.14 shows the present value level of investment by intervention. Table 6.5 breaks these results down further to show investment by intervention in each catchment.

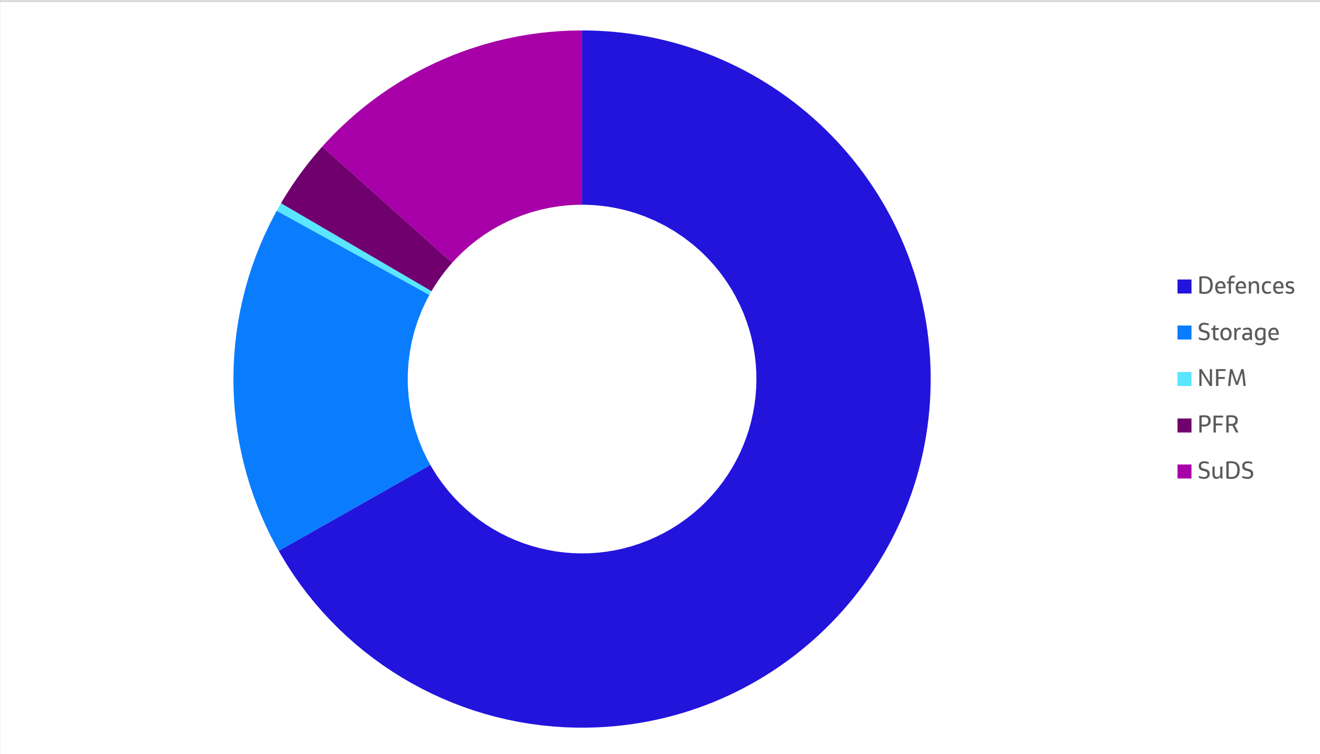

Figure 6.14. Present value investment by intervention.

Table 6.5. Present value investment by intervention and by catchment.

	Defences	Storage	NFM	PFR	SuDS
Thames	£3,182 m	£0 m	£21 m	£120 m	£399 m
Nene	£311 m	£642 m	£0 m	£28 m	£296 m
Ouse	£16 m	£210 m	£0 m	£21 m	£8 m
Total	**£3,509 m**	**£852 m**	**£21 m**	**£168 m**	**£703 m**

6.2.8 Contribution to costs

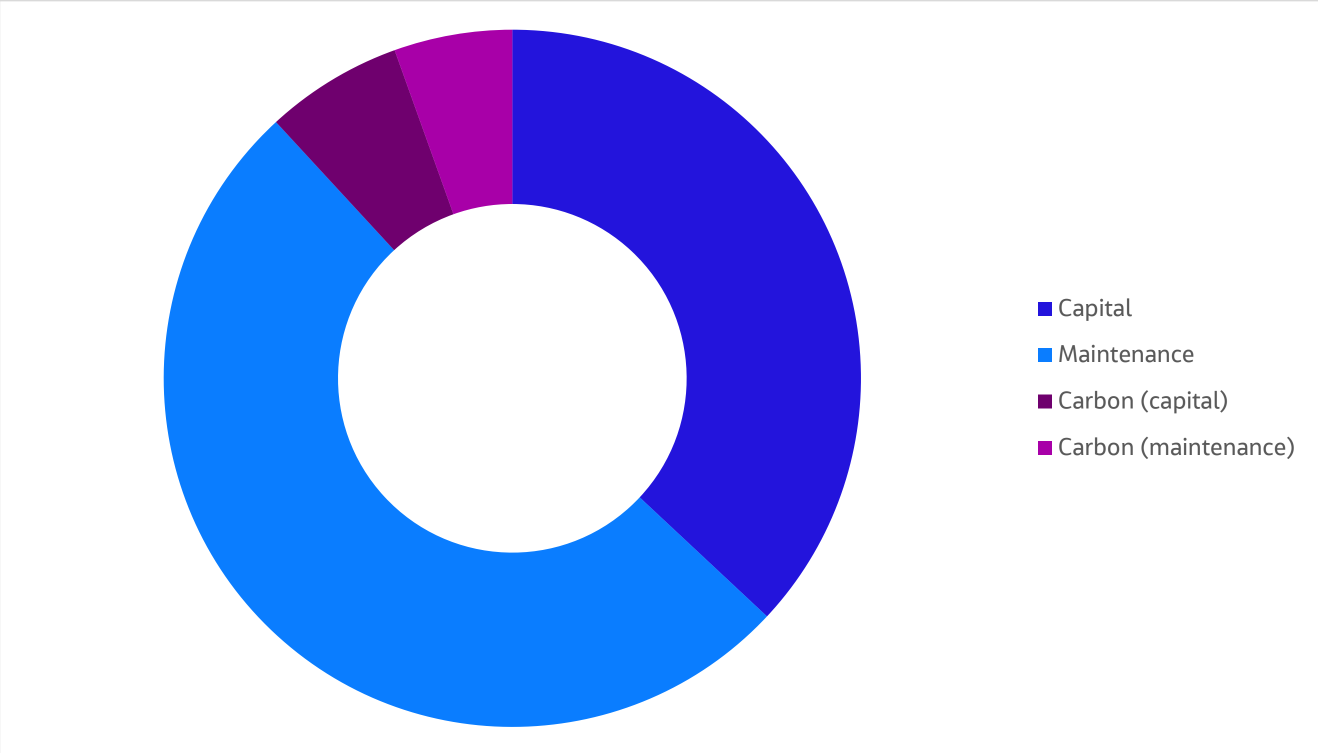

Figure 6.15. Contribution to costs.

Figure 6.15 shows the breakdown of present value costs into capital, maintenance and carbon (capital and maintenance). Table 6.6 shows how this breakdown varies by source.

Table 6.6. Contribution to costs by source.

	Capital	Maintenance	Carbon (capital)	Carbon (maintenance)
Rivers	£1,851 m	£2,738 m	£372 m	£54 m
Surface water	£326 m	£276 m	£4 m	£269 m
Total	**£2,177 m**	**£3,014 m**	**£376 m**	**£323 m**

6.2.10 Investment profile

The analysis has 10 decision points – starting with investment in year 0 – which is deemed to be cost beneficial irrespective of future changes in risk – followed by investment at 10 year intervals, which is explored under each future. The by-product of this approach is that the raw investment profile is made up of large peaks representing capital investment at each 10 year investment point, interspersed with annual maintenance investment. Figure 6.16 shows these raw results – showing the expected investment (averaged across the range of future scenarios).

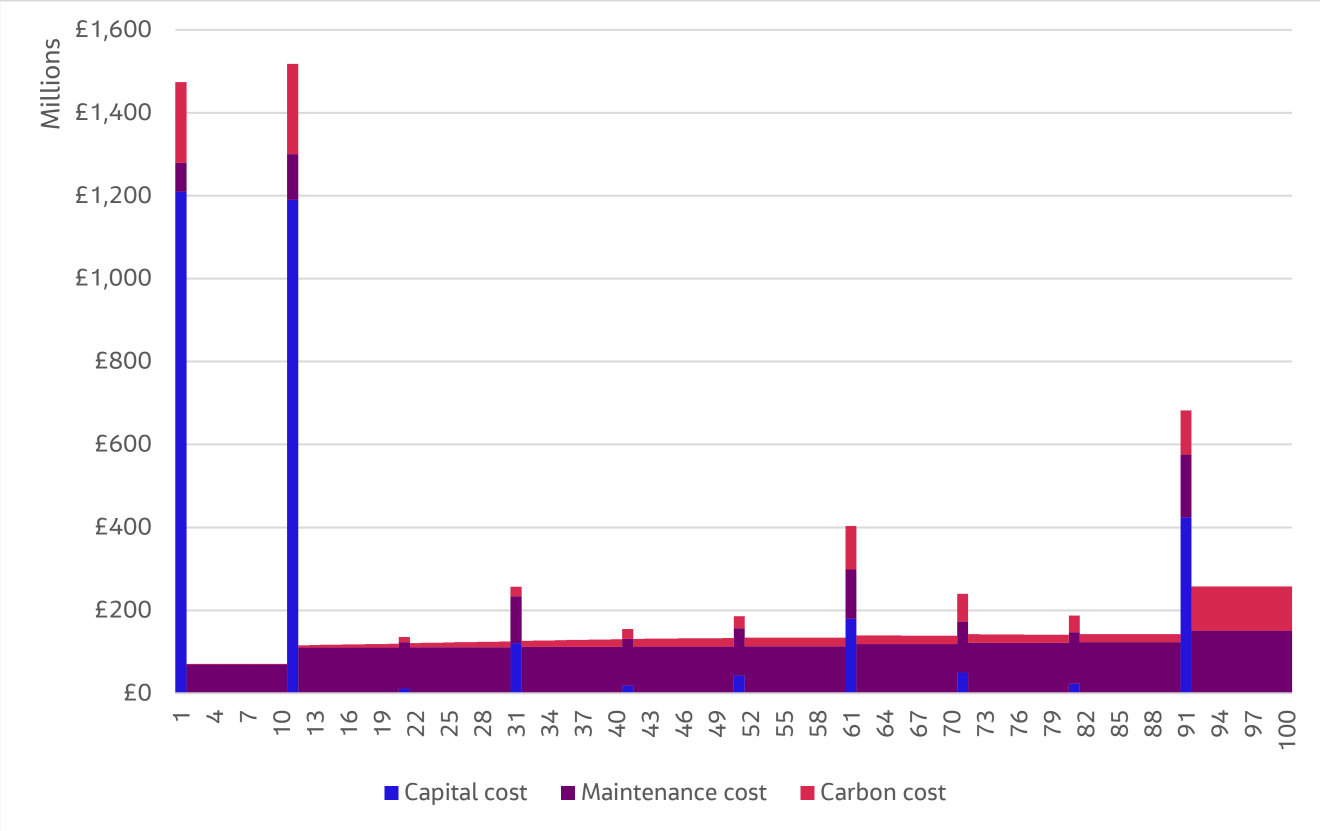

Figure 6.16. Raw annual investment through time.

In practice, investment is likely to be spread more evenly through time, for example, guided by annual or spending review period funding. We made the decision to exclude these annual constraints from the analysis, because some of the large interventions (in particular catchment storage) would not readily fit into that constrained model.

However, we have post-processed the raw investment profiles to produce a smoothed investment profile – treating each decision point as investment that would in reality be spread over the following 10-year period. Figure 6.17 shows this redistributed investment profile. As with Figure 6.16, these results are averaged across the range of future scenarios, and Figure 6.18 goes on to show how the results differ across the range of futures (the yellow band shows the minimum and maximum investment, providing a range in which we would expect the future investment profile to follow).

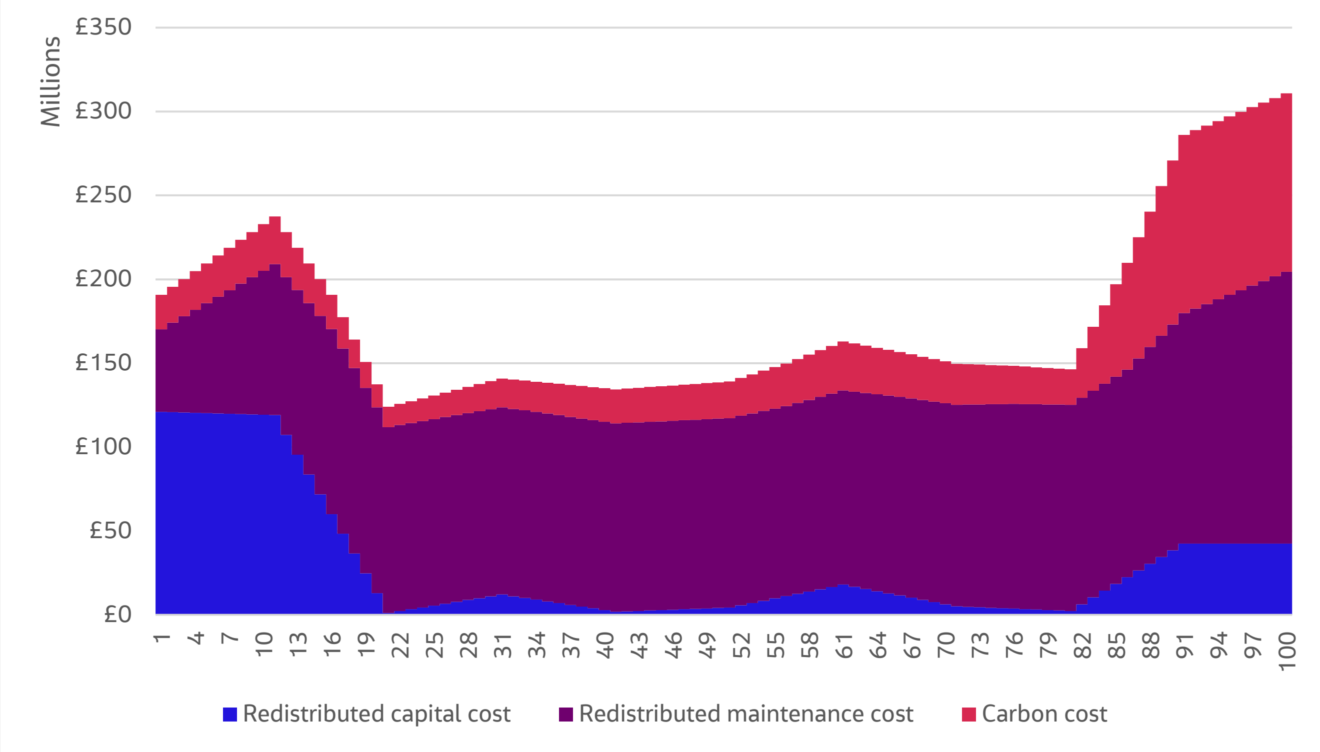

Figure 6.17. Redistributed annual investment through time.

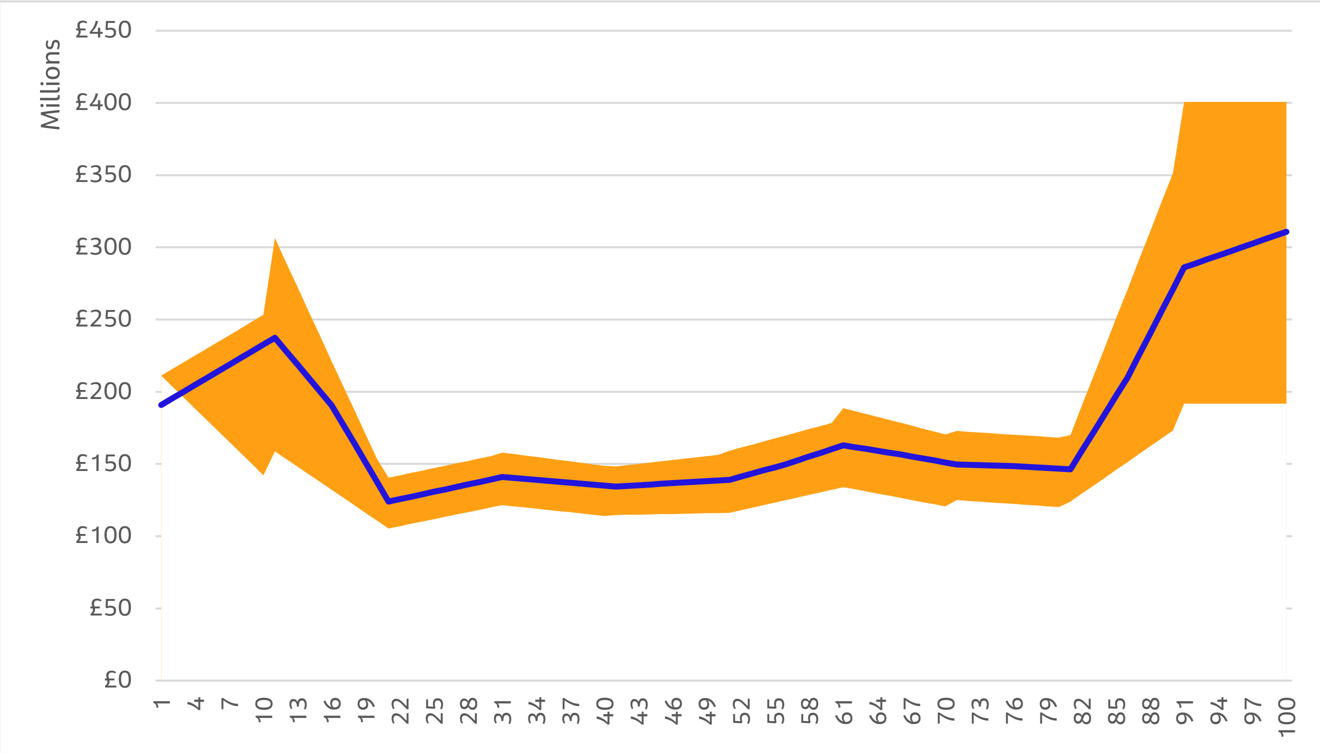

Figure 6.18. Range of annual investment under different futures.

6.2.12 Residual risk

Figure 6.19 shows the present value damage avoided (green) and residual risk (pink) for river and surface water flooding. They are stacked so the total length of the bar shows the total baseline risk. The graph shows that with the optimum level of investment, a total of £28.2 billion (£22.8 billion rivers, £5.6 billion surface water) of the £63.5 billion baseline (£26.8 billion rivers, £36.7 billion surface water) present value risk is avoided, leaving a residual risk of £35.3 billion (£4.0 billion rivers, £31.3 billion surface water).

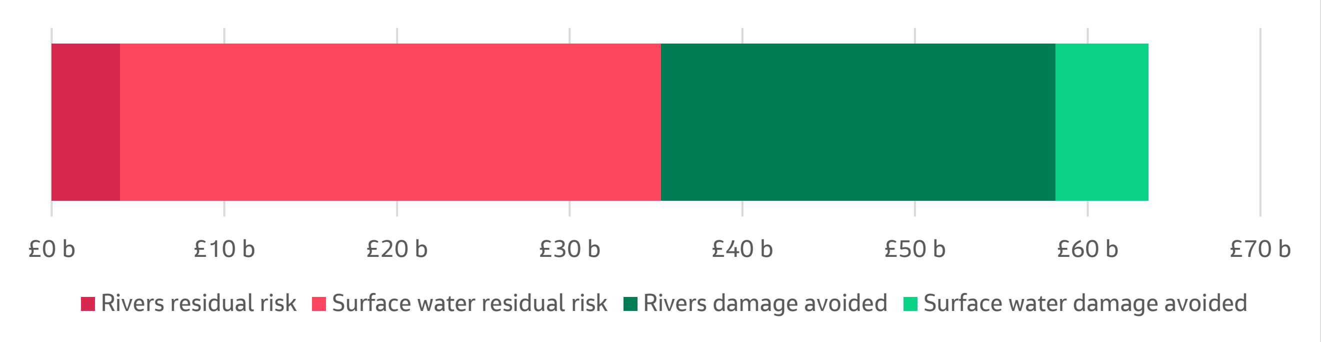

Figure 6.19. Present value damage avoided and residual risk for each source.

Figure 6.20 shows the risk profile under the optimum investment scenario – as in Figure 6.19, showing the annual average damage avoided (green) and residual annual average damage (pink). Table 6.7 pulls out annual average damage snapshots for year 0 and a number of future years.

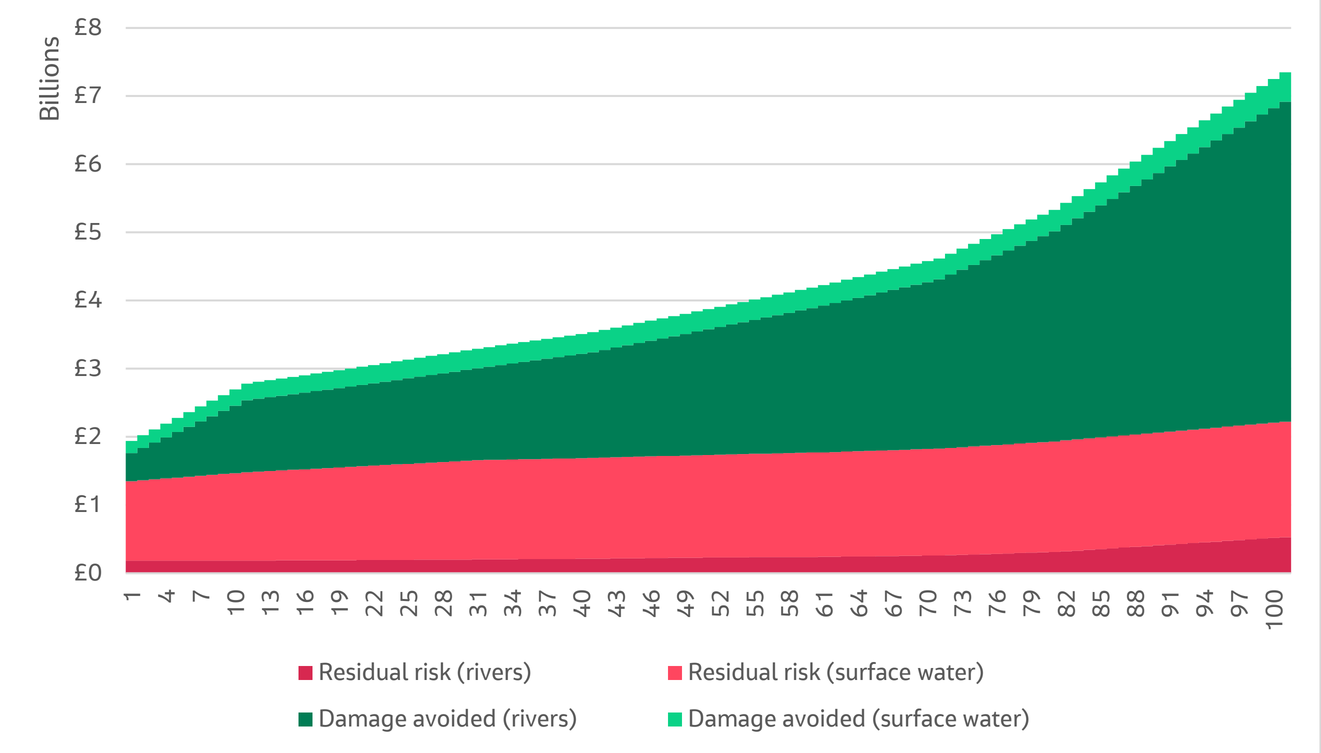

Figure 6.20. Annual average damage avoided and residual risk profile for each source.

Table 6.7. Annual average damage avoided and residual risk profile.

Year	0	10	25	50	100
Residual risk (rivers)	£183 m	£185 m	£197 m	£231 m	£529 m
Residual risk (surface water)	£1,167 m	£1,296 m	£1,416 m	£1,505 m	£1,693 m
Damage avoided (rivers)	£410 m	£1,056 m	£1,270 m	£1,842 m	£4,696 m
Damage avoided (surface water)	£182 m	£245 m	£278 m	£297 m	£432 m

Figure 6.21 shows the distribution of baseline and residual risk profiles under the range of futures. Figure 6.22 shows the full set of residual risk profiles.

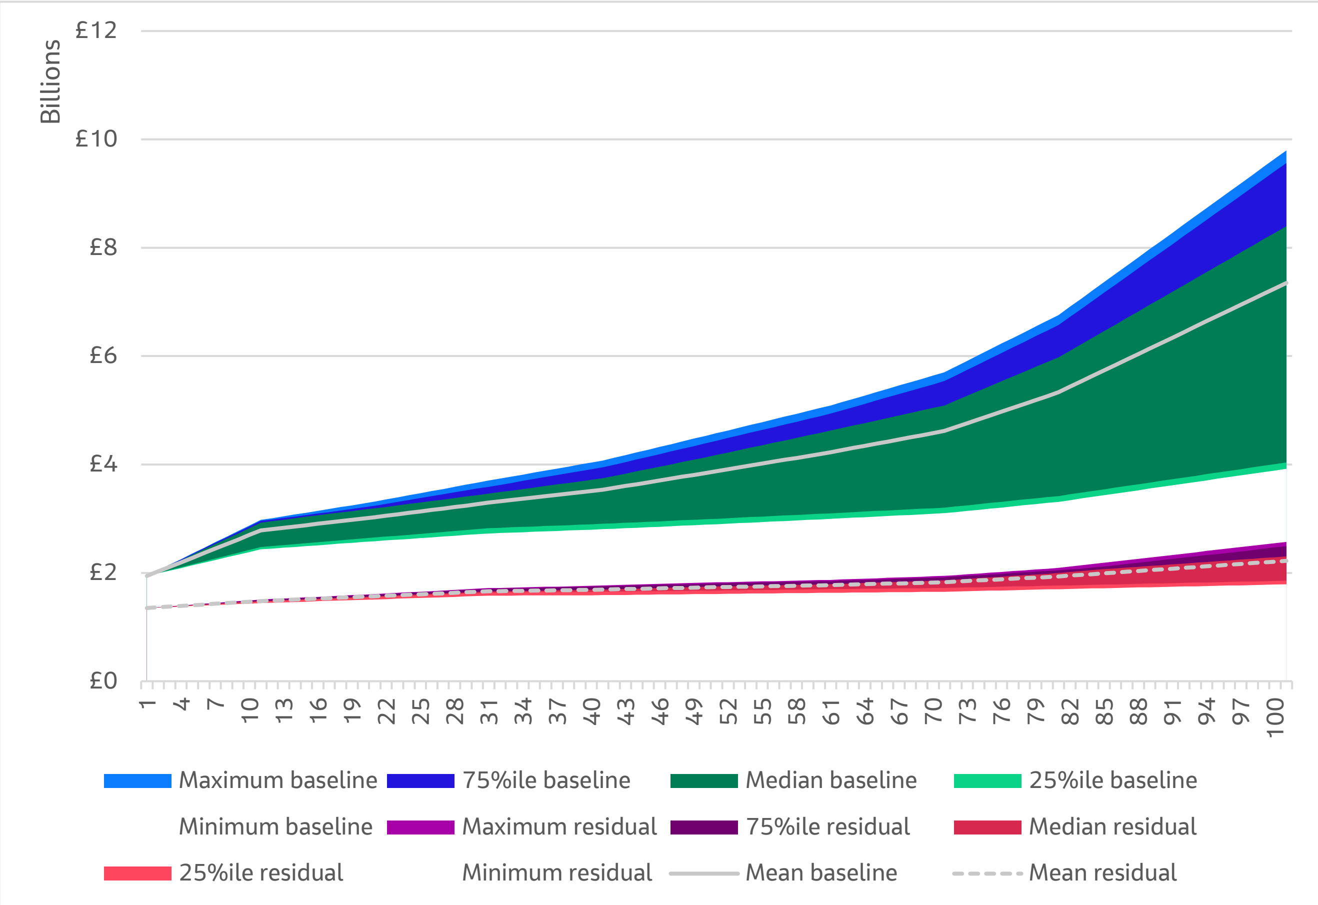

Figure 6.21. Baseline and residual risk profile range under different futures.

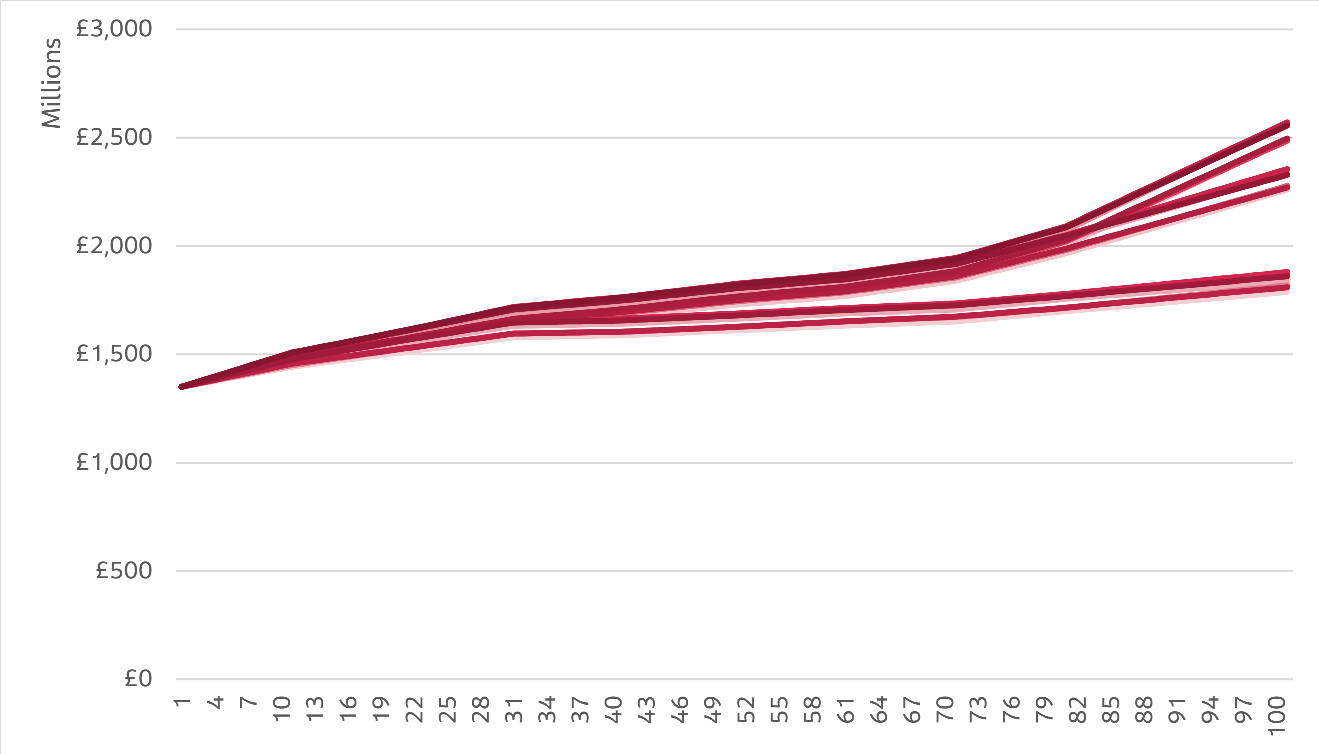

Figure 6.22. Residual risk profiles under different future scenarios.

6.2.14 Spatial results

Figure 6.23 and Figure 6.24 show the ranking of flood areas (river flooding) and flood risk management systems (surface water flooding) respectively – ranking the analysis units by the absolute reduction in present value risk under the optimum investment, where 1 is the area which has the greatest reduction in risk.

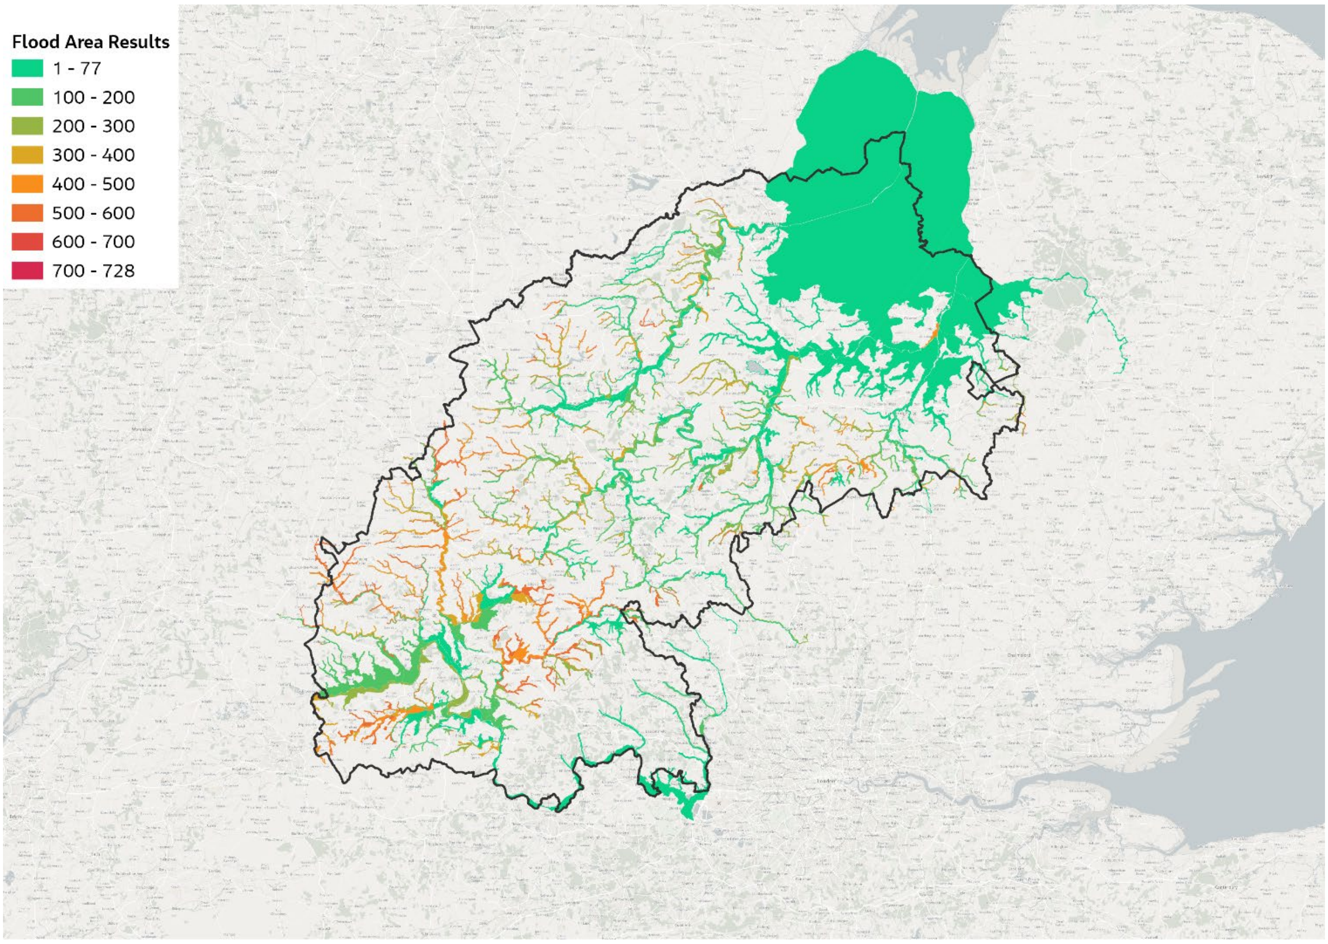

Figure 6.23. Benefits of the optimum level of investment in river flood risk management interventions by flood area, ranked from highest (1) to lowest (728).

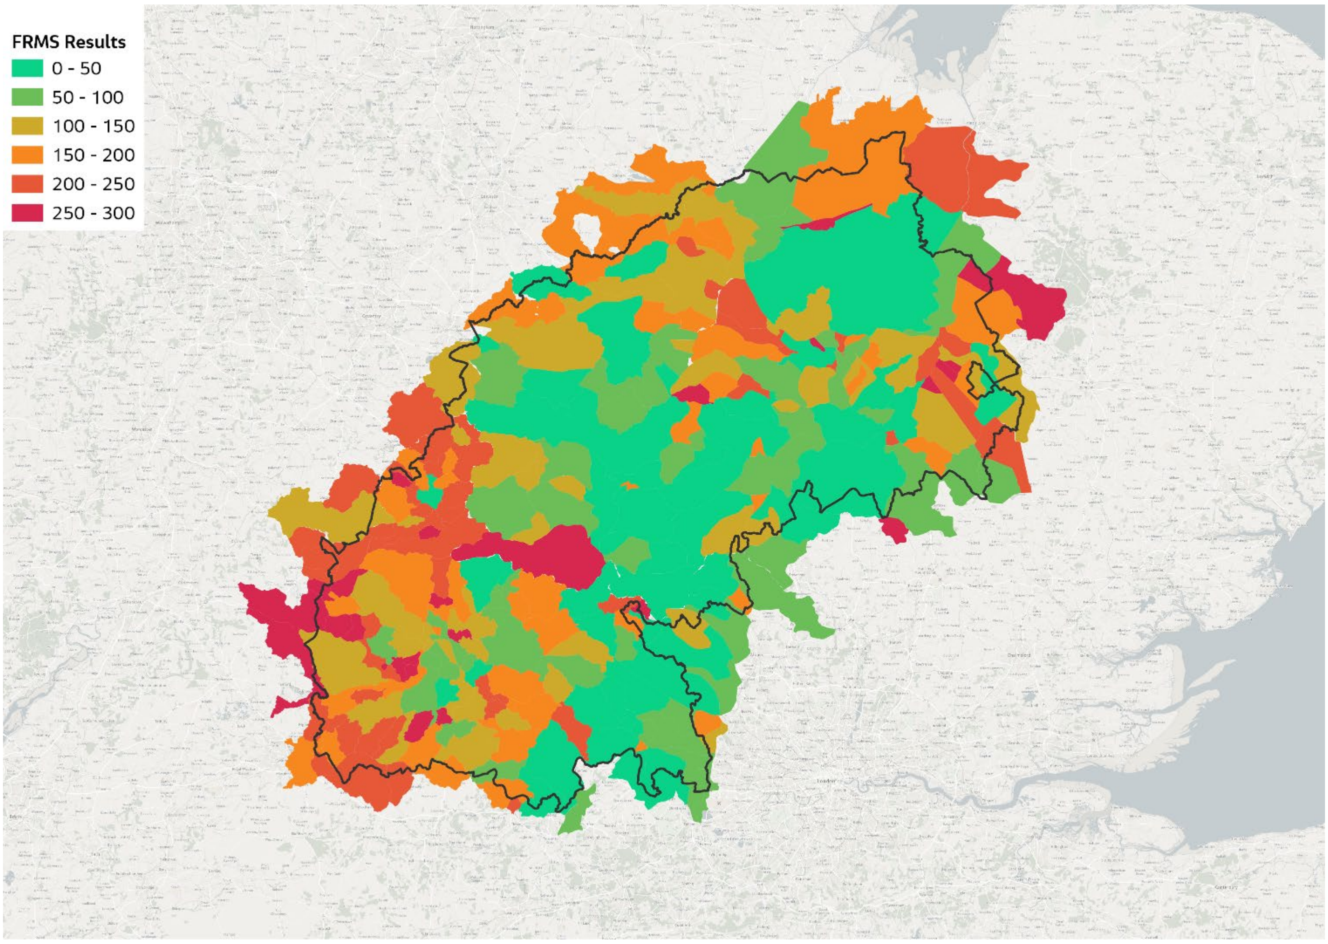

Figure 6.24. Benefits of the optimum level of investment in surface water flood risk management interventions by flood risk management system, ranked from highest (1) to lowest (280).

Figure 6.25 and Figure 6.26 show the percentage damage avoided by the optimum level of investment in river flood risk management interventions and surface water flood risk management interventions respectively. As indicated by the summary results for this analysis, there is significant residual surface water flood risk, which is why the map in Figure 6.26 indicates relatively low percentage damage avoided figures for surface water, especially compared with Figure 6.25, which shows that a large proportion of the risk of flooding from rivers is avoided in the optimum investment scenario.

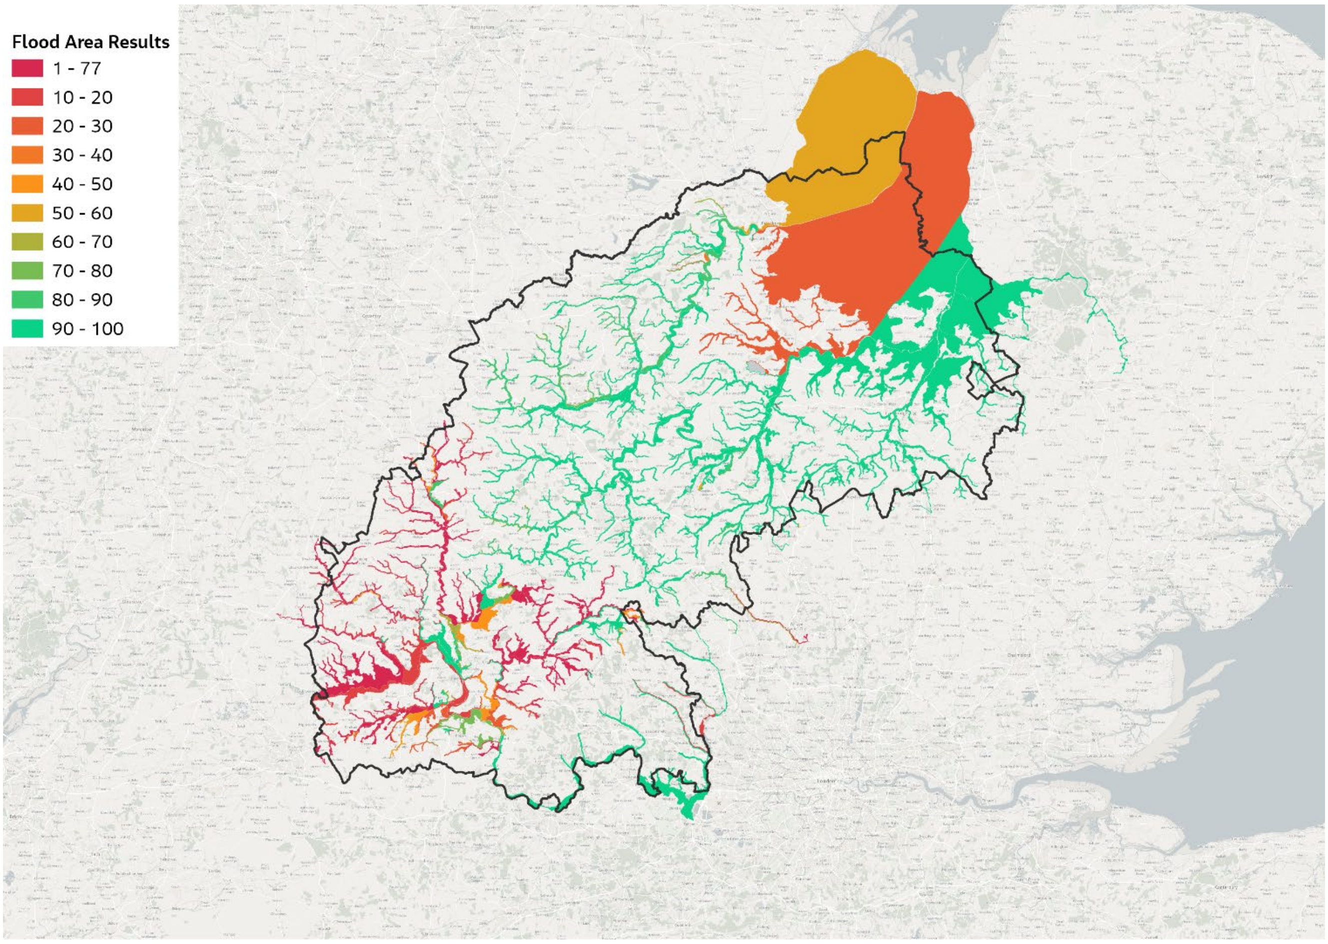

Figure 6.25. Percentage damage avoided by the optimum level of investment in river flood risk management interventions by flood area.

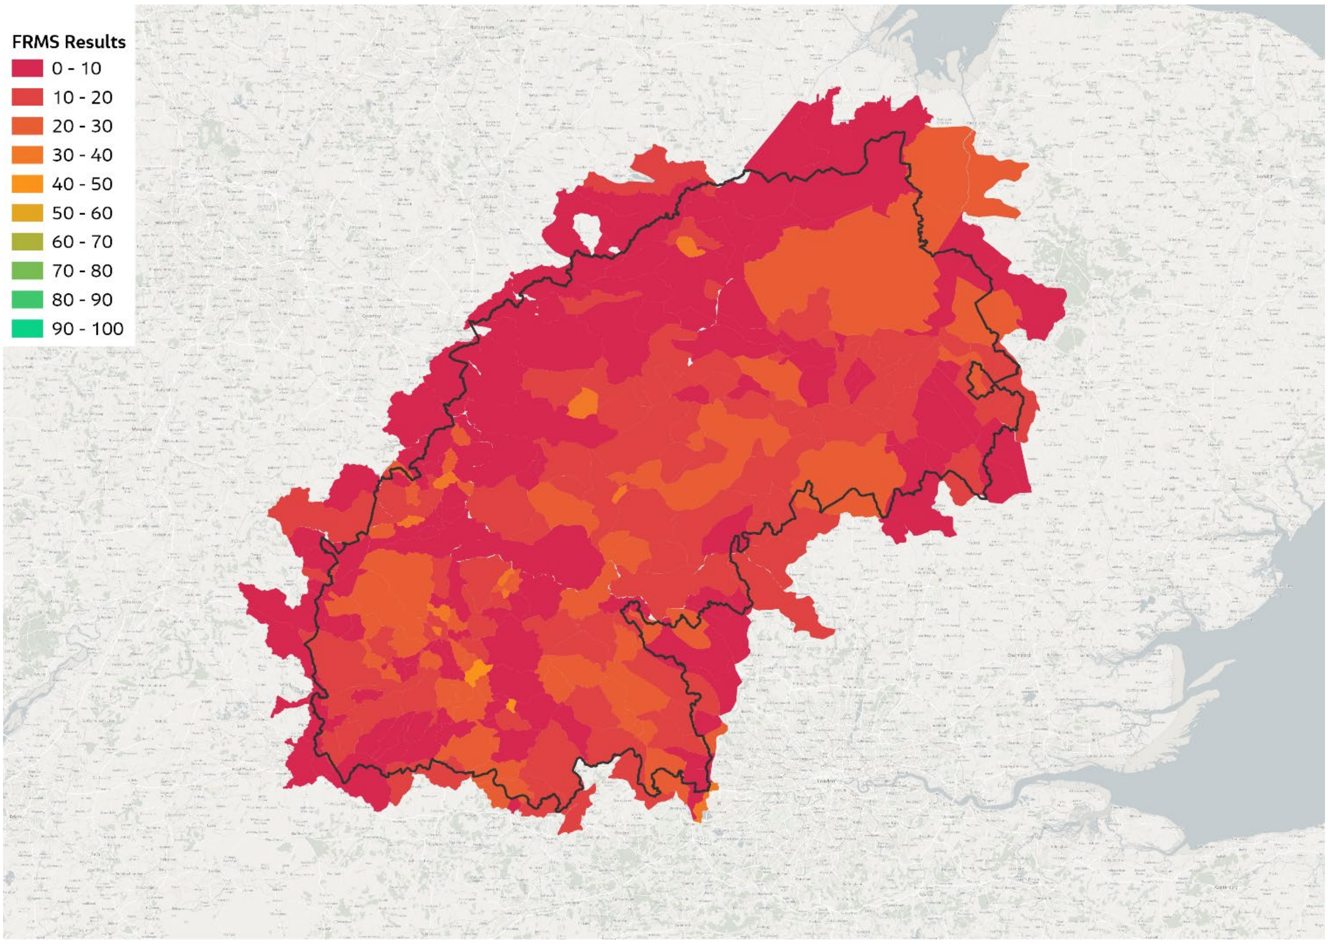

Figure 6.26. Percentage damage avoided by the optimum level of investment in surface water flood risk management interventions by flood risk management system.

Figure 6.27 and Figure 6.28 show the residual risk from the optimum level of investment, for river flooding and surface water flooding respectively. As with a number of the other map-based graphics, these show areas ranked by the level of residual risk, with 1 being the area with the highest residual risk.

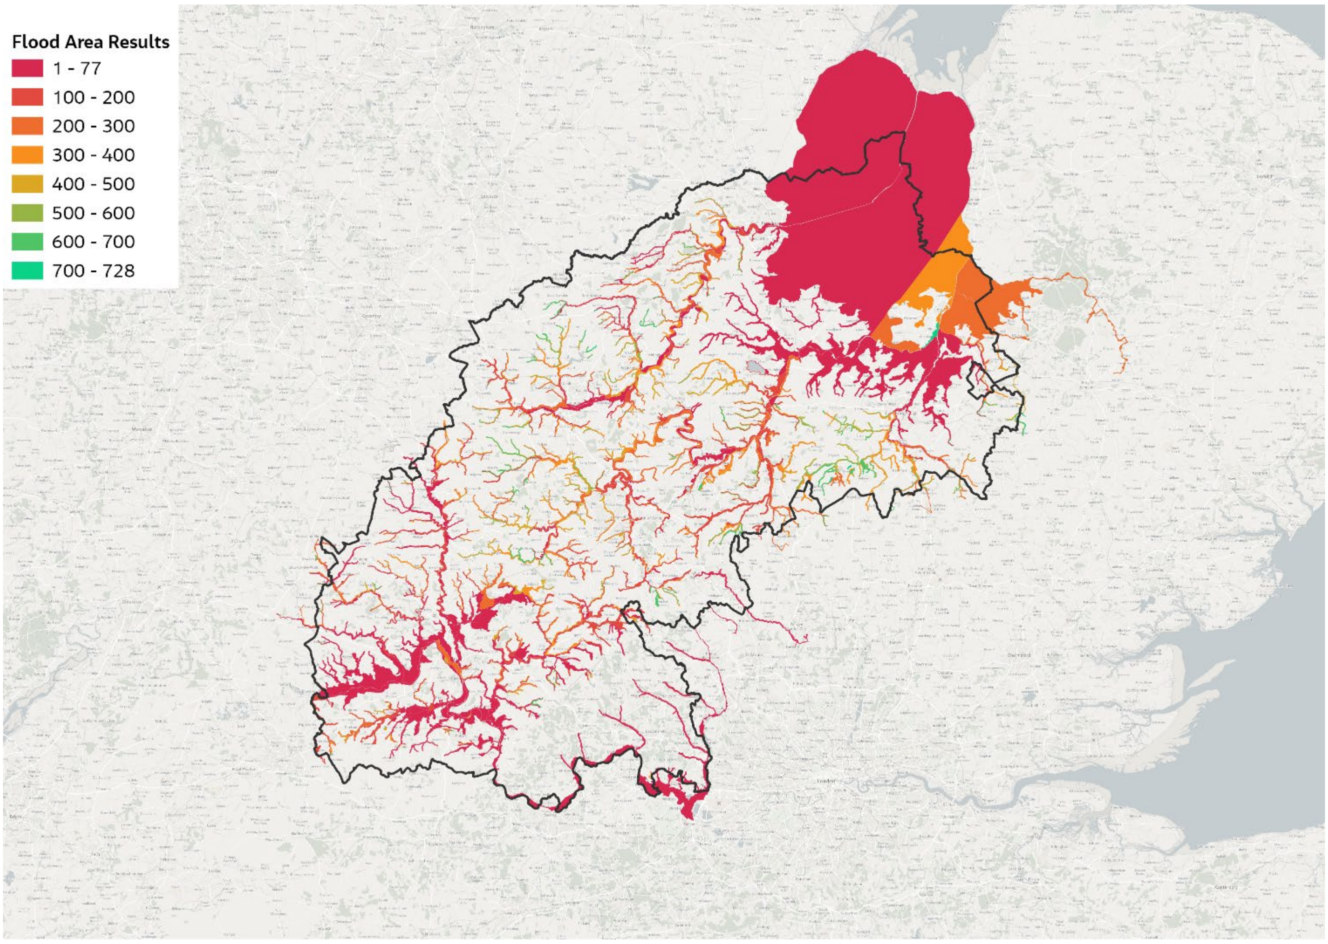

Figure 6.27. Residual risk from the optimum level of investment in river flood risk management interventions by flood area, ranked from highest (1) to lowest (728).

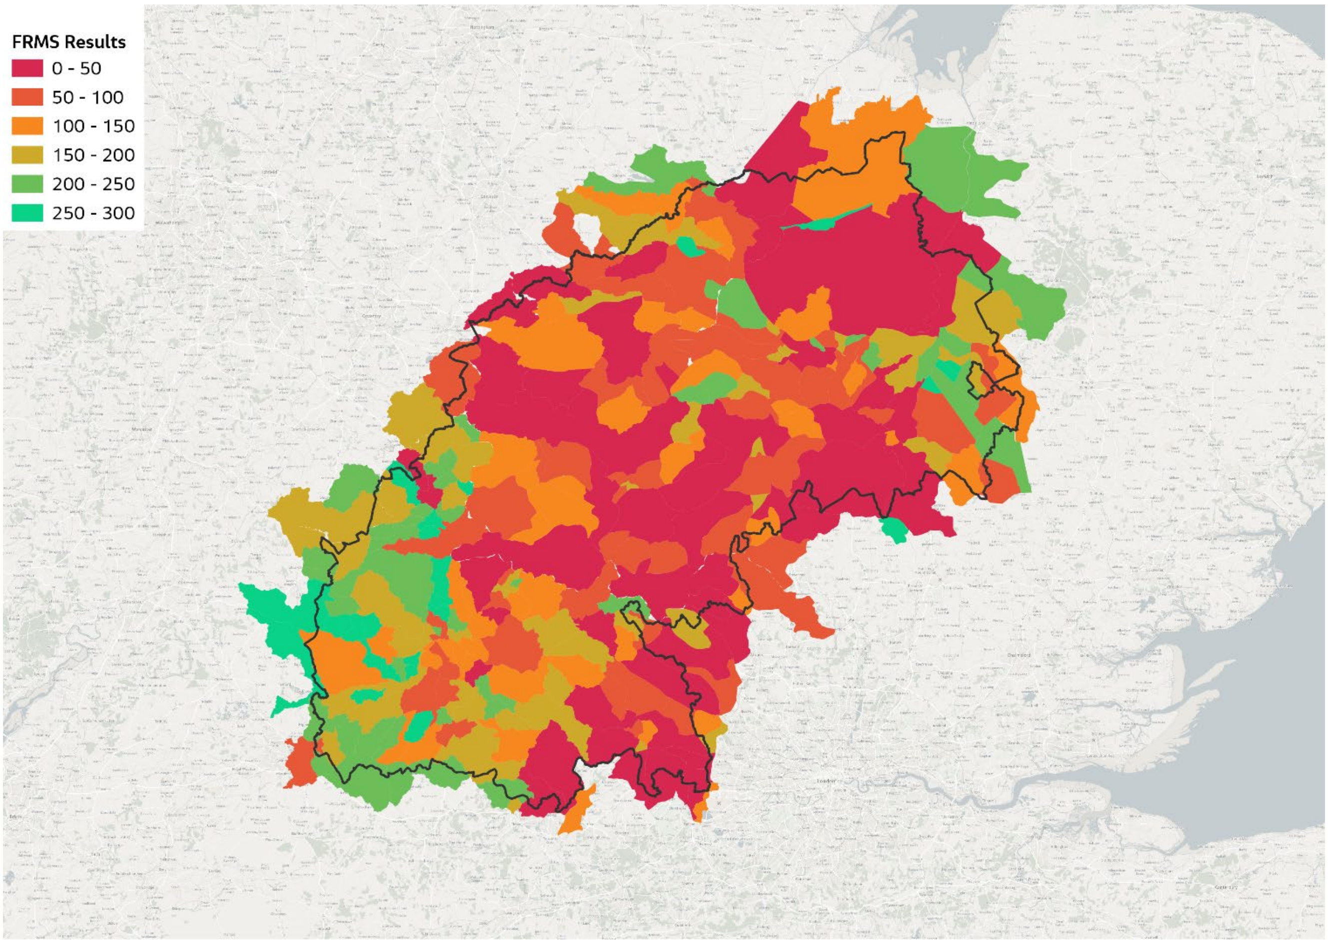

Figure 6.28. Residual risk from the optimum level of investment in surface water flood risk management interventions by flood risk management system, ranked from highest (1) to lowest (280).

6.3 Optimising investment for individual futures

In Section 6.2 above, we explored the results of the 'all futures' analysis, which applied a 'real options analysis' optimisation of investment under multiple futures. This section explores the results when investment is optimised under each individual future (and compares the results to the all futures analysis).

The following 3 figures (Figure 6.29 to Figure 6.31) describe the optimum level of investment and corresponding net present value under each of the 27 future scenarios, but colour-coded to show the distribution of results under different variables.

Figure 6.29 shows results grouped by climate change scenario, and shows that the optimum present value investment (and net present value) is largely driven by changes in risk due to climate change, and is much less sensitive to the rate or shape of development – with results strongly clustered by the climate change scenario. Figure 6.30 shows that within those clusters, there is a clear relationship between the rate of development and the optimum level of investment – with higher rates of development leading to greater investment. Figure 6.31 shows that within those clusters, there is also a relationship between the shape of development and the optimum level of investment, with expansion scenarios consistently showing higher levels of investment, and new settlement scenarios consistently showing lower levels of investment (with hybrid – as might be expected – in between).

Note that in all 3 figures, there is an anomalous result – the net present value for one scenario appears to be an outlier (but the corresponding level of investment appears to follow the pattern that would be expected). Given the clustering of all the other results, we can be confident that this is an isolated problem, and believe it is the result of an unexpected calculation error, rather than a true outcome of the analysis, so this result is shown as a grey point.

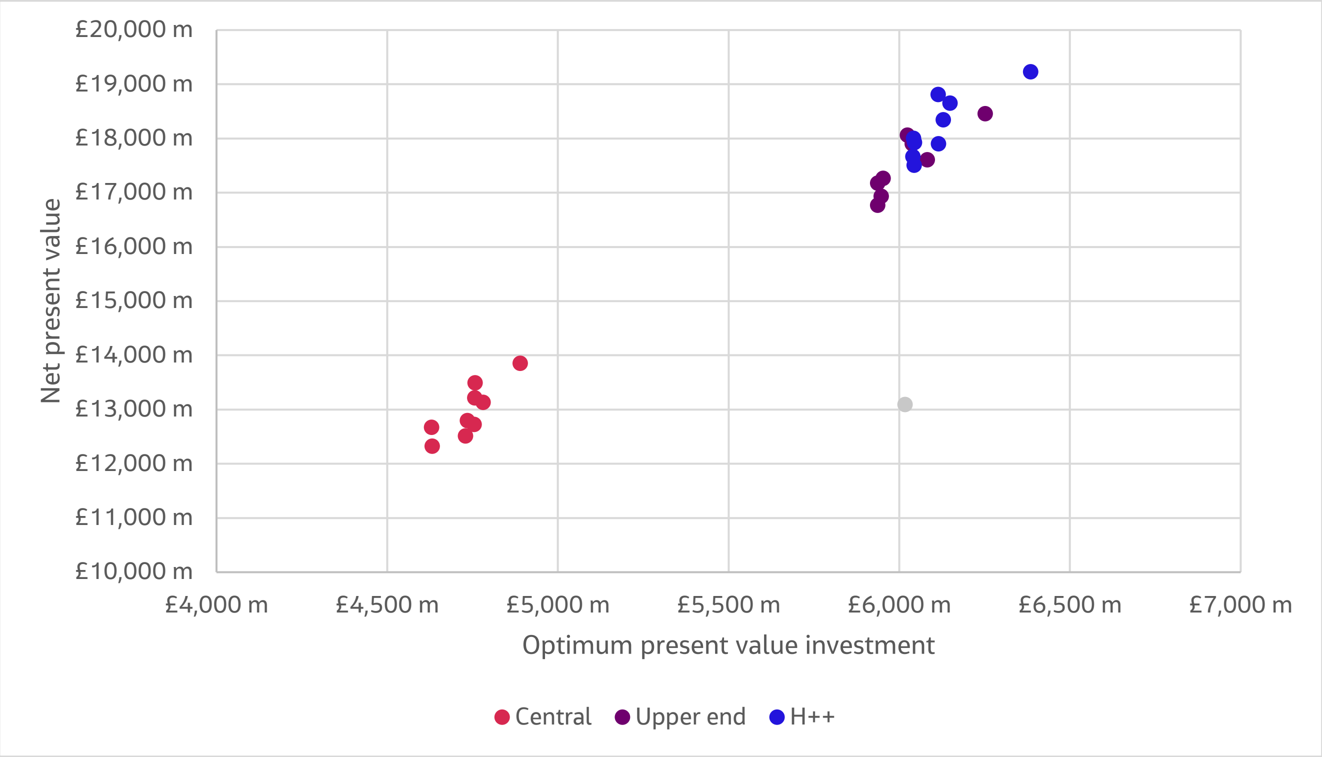

Figure 6.29. Optimum investment under each single future optimisation, grouped by climate change scenario.

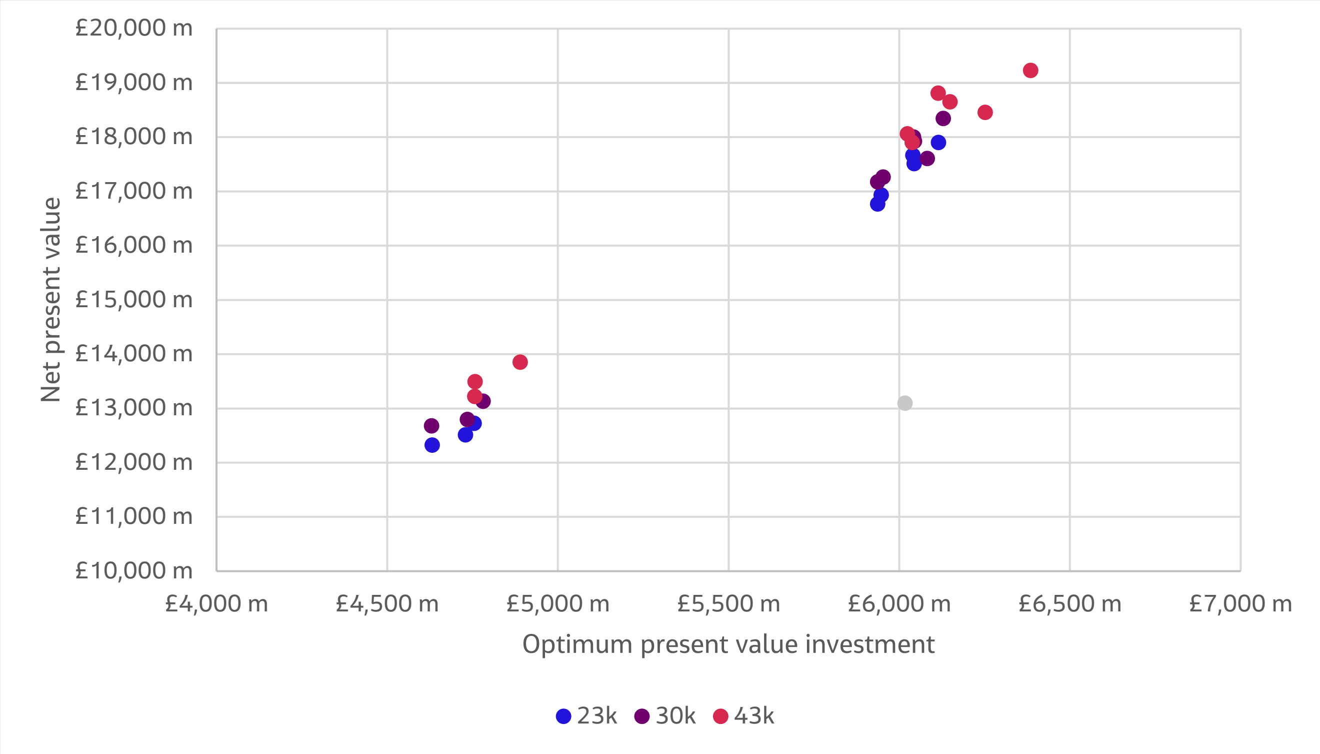

Figure 6.30. Optimum investment under each single future optimisation, grouped by development rate.

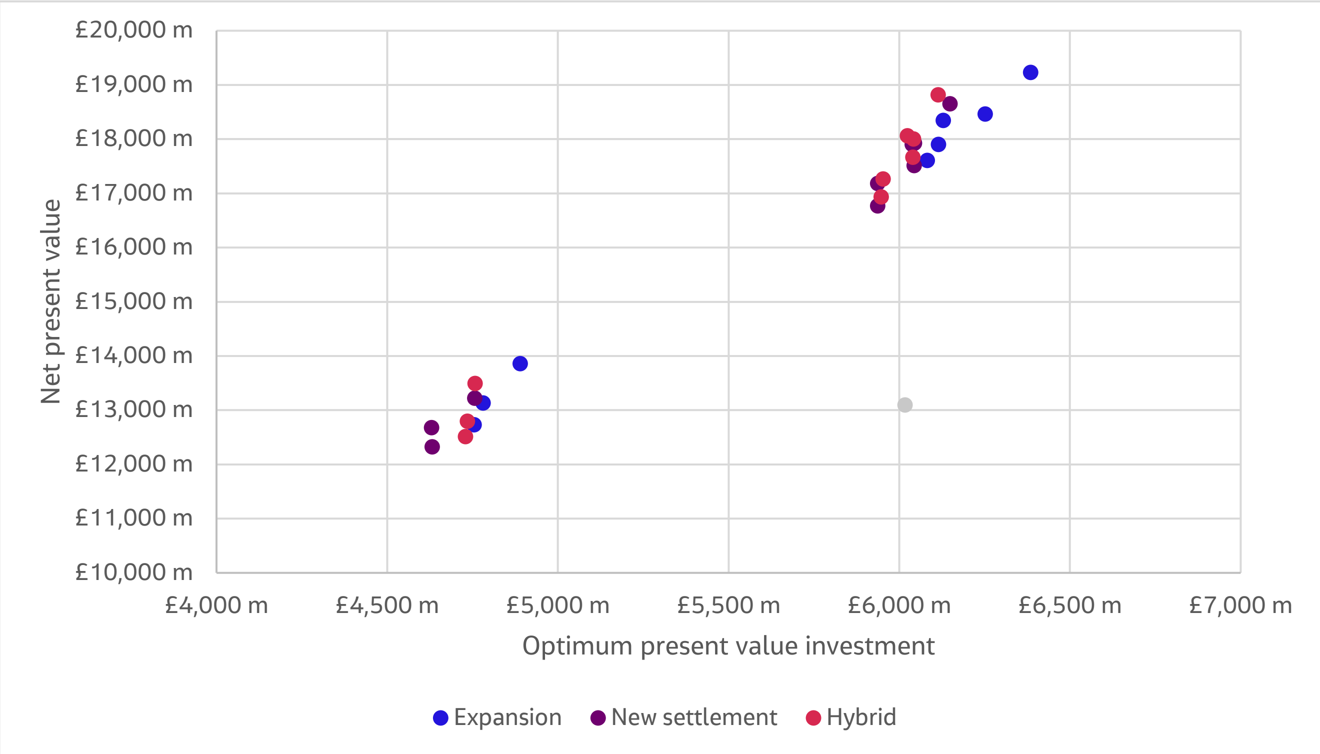

Figure 6.31. Optimum investment under each single future optimisation, grouped by development shape.

Table 6.8 shows a comparison between results from the individual future optimisation runs and the 'all futures' analysis. In general, the 'robust' investment needed to achieve a strong net present value under all futures is higher for the same future scenario than the optimum investment under a single future scenario, and the net present value is lower – although only by a small amount compared with the total present value investment. There are a handful of scenarios where the all futures analysis produces a slightly lower investment level than the corresponding single future analysis, but this is likely to be an artefact of the analysis, rather than a meaningful result.

Table 6.8. Comparison between individual future analyses and the all futures analysis.

	Investment	NPV	All futures investment	All futures NPV	Difference in investment	Difference in NPV
Central - exp_23	£4,754 m	£12,728 m	£4,757 m	£12,727 m	£3 m	-£1 m
Upper end - exp_23	£6,018 m	£13,093 m	£6,095 m	£17,088 m	£78 m	-[6]
H++ - exp_23	£6,114 m	£17,903 m	£6,181 m	£17,831 m	£66 m	-£73 m
Central - exp_30	£4,780 m	£13,134 m	£4,784 m	£13,133 m	£3 m	-£1 m
Upper end - exp_30	£6,082 m	£17,608 m	£6,159 m	£17,535 m	£78 m	-£73 m
H++ - exp_30	£6,128 m	£18,345 m	£6,195 m	£18,273 m	£67 m	-£73 m
Central - exp_43	£4,890 m	£13,855 m	-	-	-	-
Upper end - exp_43	£6,251 m	£18,460 m	-	-	-	-
H++ - exp_43	£6,385 m	£19,231 m	-	-	-	-
Central - set_23	£4,631 m	£12,325 m	£4,630 m	£12,325 m	-£1 m	£1 m
Upper end - set_23	£5,936 m	£16,770 m	£6,013 m	£16,697 m	£78 m	-£73 m
H++ - set_23	£6,043 m	£17,510 m	£6,109 m	£17,438 m	£66 m	-£73 m
Central - set_30	£4,630 m	£12,676 m	£4,629 m	£12,677 m	-£1 m	£1 m
Upper end - set_30	£5,936 m	£17,179 m	£6,014 m	£17,106 m	£78 m	-£73 m

[6] The NPV for the individual future run produced an anomalous result, so the difference has been excluded from this comparison.

H++ - set_30	£6,044 m	£17,928 m	£6,110 m	£17,855 m	£66 m	-£73 m
Central - set_43	£4,756 m	£13,219 m	-	-	-	-
Upper end - set_43	£6,037 m	£17,902 m	-	-	-	-
H++ - set_43	£6,148 m	£18,652 m	-	-	-	-
Central - hyb_23	£4,729 m	£12,513 m	£4,732 m	£12,512 m	£3 m	-£1 m
Upper end - hyb_23	£5,946 m	£16,934 m	£6,024 m	£16,861 m	£78 m	-£73 m
H++ - hyb_23	£6,039 m	£17,666 m	£6,105 m	£17,594 m	£66 m	-£73 m
Central - hyb_30	£4,735 m	£12,798 m	£4,738 m	£12,797 m	£3 m	-£1 m
Upper end - hyb_30	£5,953 m	£17,265 m	£6,030 m	£17,192 m	£78 m	-£73 m
H++ - hyb_30	£6,042 m	£18,003 m	£6,108 m	£17,931 m	£66 m	-£73 m
Central - hyb_43	£4,757 m	£13,493 m	-	-	-	-
Upper end - hyb_43	£6,023 m	£18,062 m	-	-	-	-
H++ - hyb_43	£6,114 m	£18,814 m	-	-	-	-

6.4 Optimising investment for 'no development' scenarios

Table 6.9 shows the results for a real options analysis optimisation run for 'no development' scenarios, i.e. not including any new properties, for the 3 climate change scenarios. It also compares the results with the 'all futures' analysis, which provides a useful indication of the effect that future development could have on the optimum level of investment.

Table 6.9. Results for 'no development' scenarios, and comparison with all futures analysis.

		PV cost	PV baseline risk	PV benefits	PV residual risk	NPV
No development scenarios	**All climate change futures (real options analysis)**	£5,564 m	£57,432 m	£25,812 m	£31,620 m	£20,248 m
	Central	£4,615 m	£44,274 m	£16,039 m	£28,235 m	£11,424 m
	Upper end	£5,991 m	£51,749 m	£21,646 m	£30,102 m	£15,656 m
	H++	£6,087 m	£52,678 m	£22,467 m	£30,212 m	£16,380 m
All futures analysis (for comparison)		£5,634 m	£63,500 m	£28,209 m	£35,291 m	£22,575 m
Comparison		99%	90%	92%	90%	90%

7. Lessons learned

In shaping the optimisation approach, we learnt a great deal about the complexities of implementing this kind of analysis, which may be relevant to other regional or national-scale economic studies.

7.1 Interventions at different scales and timings add significant computational complexity

The core analysis in the long-term investment scenarios (LTIS) explored the level of investment to manage flood risk. In some respects, the LTIS analysis included complexities that this project has excluded – in particular annual constraints on funding and modelling the deterioration of assets. However, the core LTIS analysis was strongly focussed on **assets**, with new investment driven by assets needing replacement – therefore not explicitly optimising the timing of investment – and when LTIS introduced a portfolio of interventions in LTIS 2019, it did so by introducing a sequential analysis and with a heavily constrained (around 800 total combinations) number of scenarios that deal with interventions at different scales.

For the OxCam economic evidence project, we introduced the portfolio of responses on a 'level playing field' – meaning that we optimised investment across the portfolio, rather than for each intervention in turn. This in its own right may not have dramatically increased the computational complexity, but the scope of the project had a clear goal to optimise the **timing** of investment as well as the level of investment. This meant that rather than just considering each permutation of interventions (which could have been managed through a similar type of sequential analysis to LTIS 2019), we needed to explore a much more complex **decision tree** of possible options. This in turn meant that the number of possible permutations of levels and timings of interventions increased exponentially and meant that we were no longer able to explore every branch of that decision tree.

Future projects should explore the trade-offs between these high-level approaches, especially now we have learnt more about the computational implications of these decisions.

7.2 Different problems demand different optimisation solutions

Closely linked to the challenges outlined above, we found that the nature of the problem for OxCam was more dramatically different to the problem for LTIS than we anticipated, and different again to the needs of other (e.g. water resources planning or genetic algorithms) optimisation problems. As Section 4.1 indicates, this meant that we needed to explore a variety of approaches to identify the optimisation solution that worked best for OxCam, and that ultimately needed to develop a bespoke solution that met the project's needs. We learnt an important lesson about the iterative nature of developing this kind of analysis, and the challenge in producing those iterations when constrained by computation time.

7.3 Future projects need to adopt more scalable computing

The biggest challenge in this project was the computational complexity of the optimisation. Even with the simplifications we made, a significant level of effort to optimise the code to run efficiently and storing/reusing intermediate outputs, a full analysis still took several weeks to run. We ran the analysis across 4 high performance processing machines with data distributed across the machines, and processes run in parallel within the optimisation across the processing cores of each machine. But future projects could take a cloud computing approach to solving this problem – and design it to be massively scalable from the outset. For example, the slowest part of the analysis was the present value impacts (and therefore benefits) calculation, which was run in parallel for each flood area across available processing cores – but the number of cores was a long way short of the number of flood areas, meaning that the processes effectively formed a 'queue' and were processed in turn across the available cores. Using cloud computing, these processes could be run completely in parallel.

7.4 Real options analysis provides a bridge between optimisation and adaptation

The results described in Section 6 show that it is possible to develop an optimisation approach which is compliant with Treasury economic analysis guidance while still representing the need to adapt future decisions to an evolving understanding of risk – and the lynchpin to that connection is the real options analysis. We have also shown that we can apply the principles of a real options analysis to a problem space much larger than most 'standard' options appraisals.

What the results perhaps do not do – at least not without further interpretation – is help us to piece together a true adaptive plan, or specific triggers for making different decisions. But this is, as much as anything, a by-product of the nature of the analysis, and of the study area. For example, Thames Estuary 2100 identified patterns in the flood risk management policy to adopt in different parts of the estuary. For the OxCam project, we start with a larger area with much greater differences in the level of protection offered to manage flood risk (and in the level of risk) – making it much harder to form a coherent regional scale narrative around an adaptive plan. For example, the trigger point for a given intervention at a given location could be quite different to any other locations. What we can say, however, is the pattern of investment needed regionally – what investment is cost beneficial in all futures, and for investment which is dependent on future risk, what the spectrum of timing (and level) of investment is under different futures. This is presented in the summary report.

7.5 Outputs are clearly shaped by the nature of inputs

It was necessary to define a series of interventions – and 'levels' for each intervention – and to define a series of decision points, and a range of specific future scenarios. This does not diminish the usefulness of the outputs, but the results show signs that the outputs are being influenced by the nature of the inputs. For example, the development scenarios have assumptions baked into them about the level of development on the floodplain – this is likely to be a direct contributing factor in the level of correlation between development and future risk, and why it appears that development is less influential than climate change. Further work could:

- Explore different development scenarios. For example, development scenarios based on evidence of local plans, or similar modelled development scenarios – but ones with stronger or weaker planning control, with different levels of development in areas at risk.
- Explore a wider range of 'sizes' of intervention – both in terms of extremes (i.e. allowing for a 'very small' and 'very large'), and in creating sizes 'between' the existing sizes.
- Explore the possibility of defining interventions as a continuous spectrum, rather than discrete 'sizes'. This could be achieved simply through interpolation of the modelled sizes of intervention.
- Explore more targeted storage interventions, rather than the catchment-wide benefit offered by those developed for OxCam. This could be achieved through 1D modelling of the water level reducing effects of different storage options, in place of the simplified method adopted in this project.

8. References

Clarke, J., Thacker, S., Pant, R., & Zorn, C. (2018). *Long Term Investment Scenarios: Infrastructure Impacts Analysis Technical Report.* Environment Agency.

Clarke, J., Wicks, J., Steel, M., Smith, C., & Neve, P. (2018). *Long term investment scenarios 2019: Effectiveness of the full range of responses (technical report).* Swindon: Environment Agency.

Department for Environment Food and Rural Affairs. (2020, May). *Evidence Review of the Concept of Flood Resilience: Final Report FD2716*. Retrieved from http://randd.defra.gov.uk/Default.aspx?Menu=Menu&Module=More&Location=None&Completed=0&ProjectID=20293

Environment Agency. (2010). *Flood and coastal erosion risk management appraisal guidance.* Bristol: Environment Agency.

Environment Agency. (2014). *Flood and coastal erosion risk management. Long-term investment scenarios (LTIS) 2014.* Retrieved from https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/381939/FCRM_Long_term_investment_scenarios.pdf

Environment Agency. (2018, March). *Accounting for adaptive capacity in FCERM options appraisal. User guide - SC110001/R1*. Retrieved from https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/690257/Accounting_for_adaptive_capacity_in_FCERM_options_appraisal_-_user_guide.pdf

Environment Agency. (2019). *Long-term investment scenarios (LTIS) 2019*. Retrieved October 1, 2020, from https://www.gov.uk/government/publications/flood-and-coastal-risk-management-in-england-long-term-investment/long-term-investment-scenarios-ltis-2019#different-ways-to-manage-risk

Halcrow Group Ltd. (2014). *Long Term Investment Strategy (LTIS) Improvements - Part 1 Technical Documentation: FACET Modelling Documentation.* Swindon: Environment Agency.

HM Treasury. (2018). *The Green Book: Central government guidance on appraisal and evaluation.* Retrieved from https://www.gov.uk/government/publications/the-green-book-appraisal-and-evaluation-in-central-governent

Lawrence, J., Bell, R., Blackett, P., Stephens, S., & Allan, S. (2018). National guidance for adapting to coastal hazards and sea-level rise: Anticipating change, when and how to change pathway. *Environmental Science and Policy, 82*, 100-107.

Lawrence, J., Haasnoot, M., McKim, L., Atapattu, D., Campbell, G., & Stroombergen, A. (2019). Dynamic Adaptive Policy Pathways (DAPP): From Theory to Practice. In V. Marchau, W. Walker, P. Bloemen, & S. Popper, *Decision Making under Deep Uncertainty: From Theory to Practice* (pp. 187-199). Cham: Springer. Retrieved from https://link.springer.com/chapter/10.1007/978-3-030-05252-2_9

Lumbroso, D., & Ramsbottom, D. (2018). Flood risk management in the UK, Putting climate change adaptation into practice in the Thames Estuary. In *Resilience: The science of adaptation to climate change* (pp. 79-87). Elsevier.

Mallawaarachchi, V. (2017, 07 08). *Introduction to Genetic Algorithms — Including Example Code*. Retrieved from Towards Data Science: https://towardsdatascience.com/introduction-to-genetic-algorithms-including-example-code-e396e98d8bf3

Neve, P., Smith, C., & Steel, M. (2016). *Long-term investment scenarios 2014 - a chapter-by-chapter explanation.* Environment Agency.

Prutsch, A., Grothmann, T., Schauser, I., Otto, S., & McCallum, S. (2010, November 8). *Guiding principles for adaptation to climate change in Europe ETC/ACC Technical Paper 2010/6.* Retrieved from EEA Topic Centre on Climate Change and Air Quality: https://climate-adapt.eea.europa.eu/metadata/guidances/guiding-principles-for-adaptation-to-climate-change-in-europe

Stanford University. (n.d.). *Introduction to Mathematical Optimization.* Retrieved from http://web.stanford.edu/group/sisl/k12/optimization/MO-unit1-pdfs/1.1optimization.pdf

Ur Rehman, A., Awuah-Offei, K., Baker, D., & Bristow, D. (2019). Emergency evacuation guidance system for underground miners. *Society of Mining, Metallurgical and Exploration Engineers.* Denver.

Appendix A. Stretching the approach

A.1 Introduction

The approach described in the main report recognises the constraints of time and budget, and the balance that needs to be struck between economic optimisation and adaptation. Further work beyond the current study could consider the questions in the following sections, to put a greater emphasis on an adaptive approach. Indeed, the work within the current scope could be seen as initially fulfilling the steps in Figure C.1 from Portfolio Analysis through to Adaptation Pathways.

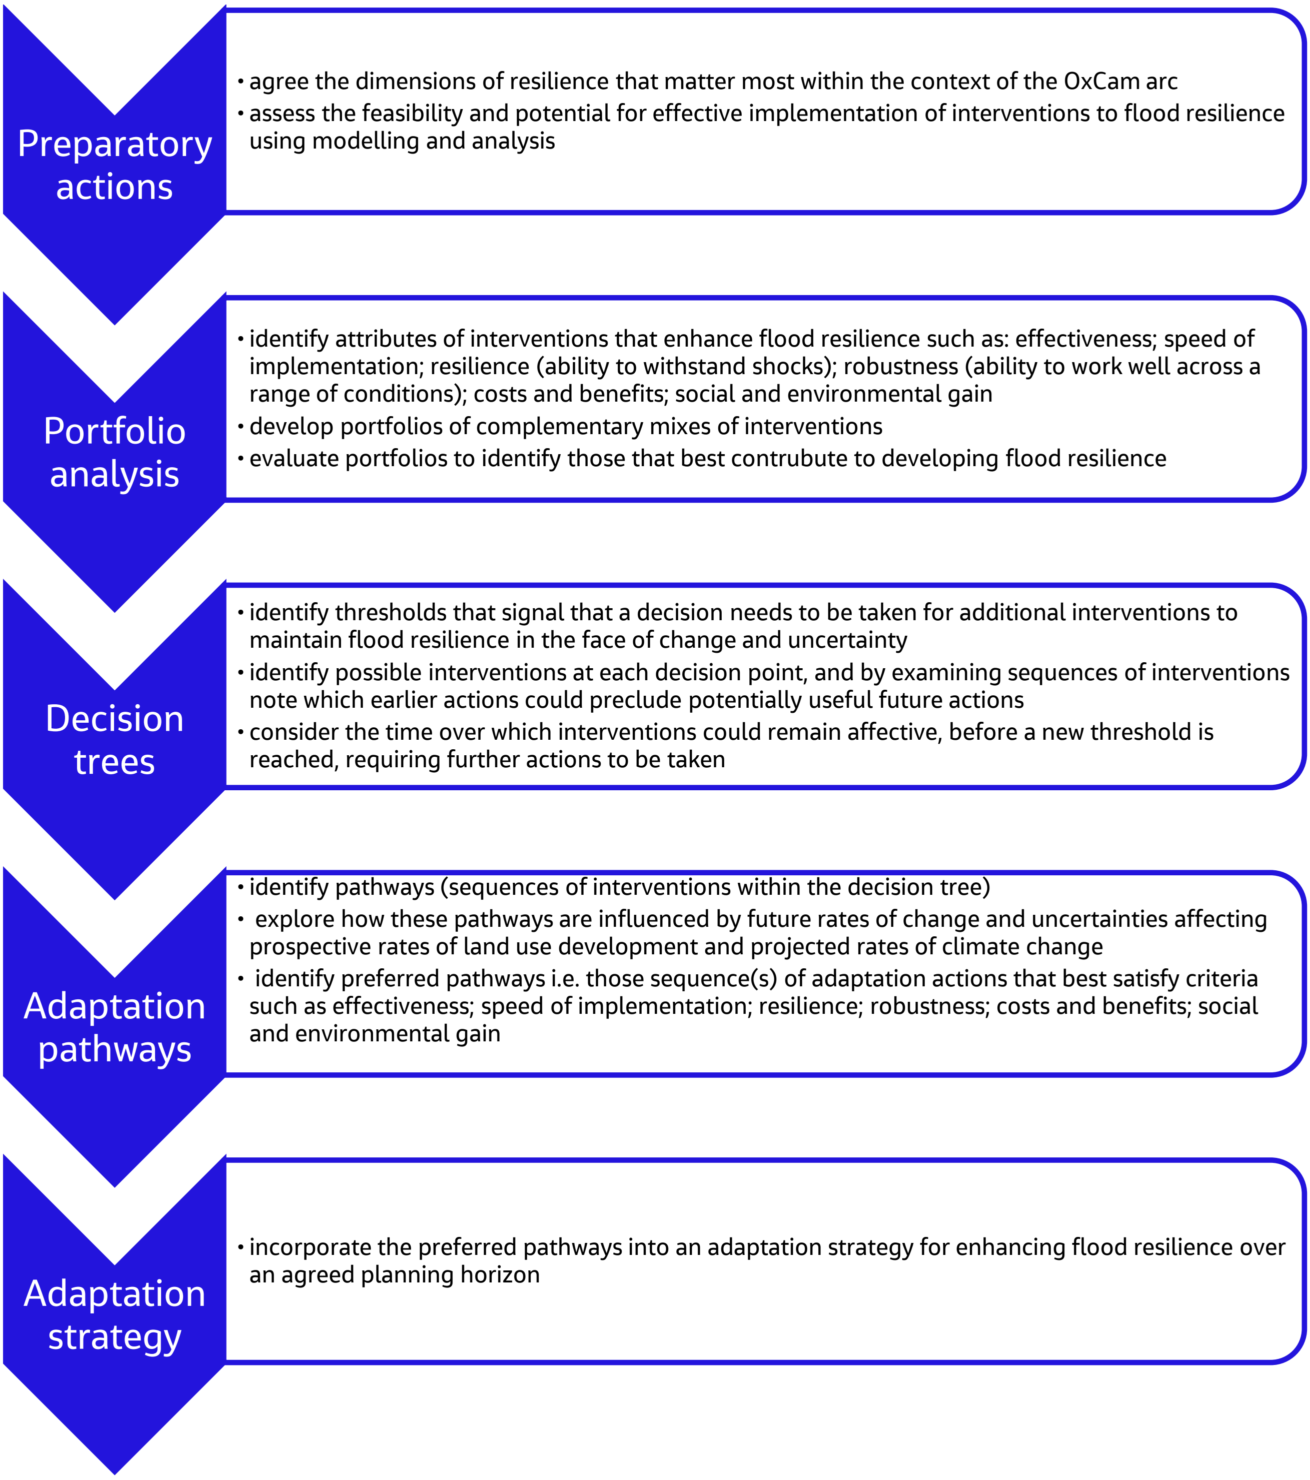

Figure C.1 Initial actions within an adaptive planning process for developing flood resilience in the OxCam Arc.

Importantly, however, is the recognition that the steps in the current scope are the first iteration in an adaptive planning process. The outcomes of this study inform a further set of questions and analysis whereby the identified pathways can be refined and embedded more in the local context. One of many illustrations of the iterative process of adaptive planning is illustrated in Figure C.2. Implementing an adaptive planning process, as envisaged by Defra (2020), involves:

- Establishing what level of resilience is already in place in order to determine where interventions to build additional resilience could best be placed.
- Determining what types of interventions could be effective and developing portfolios of interventions that can be shown to reduce the risk and impacts of flooding in ways that improve the resilience capacities of places, communities and the environment.
- Sequential and adaptive implementation of potentially effective interventions informed by monitoring and evaluating the effectiveness of measures implemented within an iterative and adaptive planning process of continuous learning and improvement.

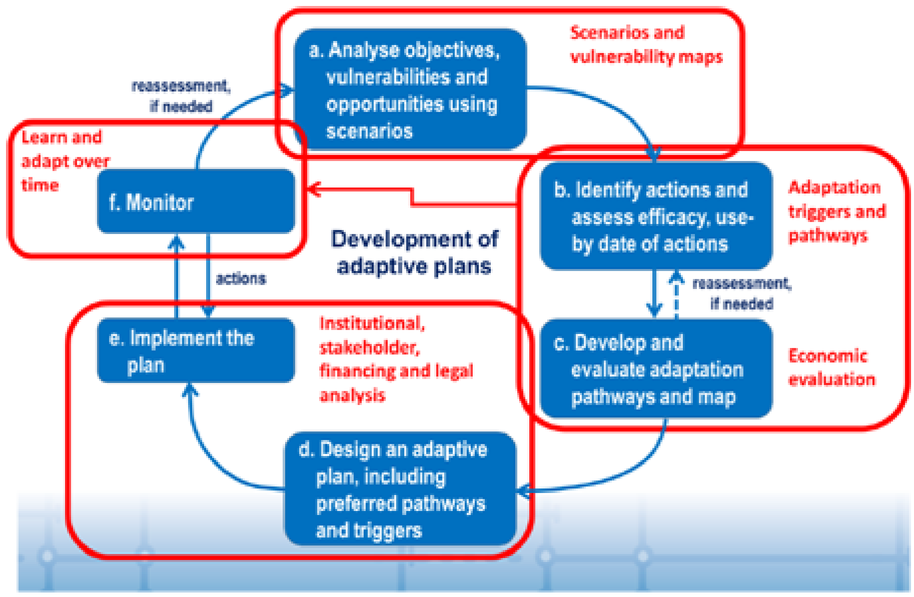

Source: Adapted from Haasnoot et al (2013)

Figure C.2. The iterative process of adaptive planning – taken from New Zealand Government Guidance for Local Government on Coastal Hazards and Climate Change (2017).

A.2 Sustainability

One of the first principles of an adaptive approach is ensuring that what we choose at any decision point is sustainable, does not add to climate change and does not cause harm to other parts of society, the economy or the natural environment. With the hugely complex web of interrelated activities and impacts in the OxCam arc, ensuring any development adheres to this principle is significantly challenging. However, the following could be considered:

- **What types of development, which are located in areas requiring flood management, are sustainable?** The National Planning Policy Framework requires land in Flood Zones 2 and 3 to only be considered if no suitable sites are available in Zone 1, and that the development in these higher flood risk areas can be made safe. The current scope is looking at providing this management through off-site (e.g. raised embankments or flood storage) or on-site (i.e. property flood resilience measures). However, compliance with NPPF requires development to be constructed appropriately (e.g. raised, flood resistant), and to also provide safe access, to not increase flood risk elsewhere and to manage residual risks. Although these issues would ultimately have to be considered at the future scale of local implementation, the additional costs of building appropriately (e.g. raised floors) could be included, and modelling could be refined to (i) report flood depths (and hazard if velocities are available) in likely areas of access and consider costs of making access safe (e.g. raised road network), (ii) represent likely displacement of flood waters and consider the cost of neutralising these, and (iii) consider the costs of managing residual risks through uplifts on allowances for warning, evacuation, essential services etc.

- **Will the flood risk management be *affordable* in the future?** NPV assessments use standard discounting which minimises costs and benefits occurring further into the future. In an anticipated development lifetime of 100 years, costs and benefits beyond approximately 50 years may not heavily influence the outcome of the assessment. Conversely, impacts of climate change will most likely increase, so that the most significant impacts are anticipated towards the end of the lifetime of the development. Therefore, any larger investments in flood management which may be required further into the future may seem to be low cost in today's NPV assessment and not heavily influence any decision on investment, whereas consideration of undiscounted (cash) costs may suggest that large future investments required to manage flooding may be considerable. The optimisation undertaken on the same modelled scenarios could consider the cash costs of future investment as an additional factor in deciding whether the level of investment can be sustained.
- **How could you broaden the economic assessment to consider more than economic sustainability?** The assumptions in the current economic assessment include some allowances for positive and negative impacts on society and the environment. However, many of these factors are growing in importance and could be included in greater depth in a broader appraisal. Carbon and natural capital are two examples which are already finding a basis in legislation and guidance through achieving 'net zero' and 'net gain', are included to some degree in standard appraisals, and which are important to fully consider to get a fuller picture of achieving sustainable development. These aspects of flood risk management are informed by their own studies, and typically included in overall appraisals either in the economics through monetisation (which is recognised as an emerging science) or separately through some form of multi-criteria analysis. The best approach to balancing the current emphasis on the economy, with considerations of the environment and society, would enable a fuller picture of sustainable development.

A.3 Flexibility

Flexibility in adaptation goes beyond recognising that different interventions could be selected at decision points in the future, to ensuring that our early actions maximise our ability to implement any of those different actions. In other words, any pathway must be considered in its entirety to ensure that whatever actions may be required at the end of the pathway are not closed off because of actions taken – or importantly – not taken, earlier in the pathway. This suggests the following could be considered:

- **What is the full cost of maintaining flexibility?** Flood storage, for example, or natural flood management may be legitimate outcomes towards the end of the pathway in terms of giving the highest NPV, but earlier decisions to raise embankments and not secure land for future storage and NFM may mean the later actions cannot be implemented. In its fullest sense, the future cost of implementing storage should therefore consider the earlier costs of securing the land required which, because of the effects of discounting and loosing the opportunity to develop the land, would likely increase the overall costs of flood storage. This is challenging in a modelling and assessment environment where the exact location of future storage is not considered, but doing so would strengthen the link between this high level assessment and the more detailed localised assessments being undertaken in catchment.
- **How could you embed flexibility into the pathways?** Further, the current scope includes calculating adaptive measures (e.g. flexibility, robustness) for the highest performing NPV pathways. However, a stretch assessment could consider how the outcomes of these measures could be fed back into the optimisation, so that the pathways are refined to maximise flexibility.

A.4 Effectiveness

Adaptation measures – in this case flood risk management – must effectively manage flood risk in efficient and equitable ways. Following the lead from government, the current scope is considering what interventions are optimal in terms of maximising NPV, rather than aiming to achieve any tolerable standard of flood protection. Trying to define what is a tolerable standard of protection, or what an appropriate level of resilience is, is a challenge which is beyond this study and should be done at as part of setting high level policy for local areas in close consultation with the communities. However, exploring the additional costs (or reduction in NPV) required to provide a greater level of resilience may inform this debate and could be explored with the modelling and costing data already being produced. Different optimisation targets could be set which explore the question: **What more would it cost to provide a greater level of resilience?** This would essentially expand the current assessment to consider incremental cost-benefits.