SIMPLIFIED REAL OPTIONS ANALYSIS FOR CLIMATE CHANGE ADAPTATION

GRAEME GUTHRIE

START ADAPTATION PLANNING TODAY

High level summary: Simplified real options analysis for climate change adaptation

Graeme Guthrie

Introduction

This research is for planners and decision-makers who are making small to medium infrastructure investment decisions now for a future that is highly uncertain. These investments might be made by local or central government, businesses, or infrastructure management.

The research explores how decision-support methods can better accommodate the uncertainty around climate change, while being as economically efficient as possible.

It focuses on Real Options Analysis (ROA), a method that has gained some attention in adaptation investment appraisal, due to its ability to incorporate and value flexibility in investment decisions. ROA is relatively complex however, and this research examines how alternative, more familiar, approaches, could retain flexibility and economic efficiency without the requirements of ROA.

This summary document presents an overview of the project and the main findings, with more details available in the full research report. It begins with a brief overview of ROA, then a guide to alternatives, followed by the findings from this particular project.

Uncertainty and Real Options Analysis

Uncertainty about the magnitude of climate change will remain high for many decades. If climate change is less severe than communities expect, they may end up spending too much on adaptation; if it is more severe than expected, they may not adapt quickly enough. Decision-makers should respond to this uncertainty by choosing adaptation actions that minimise the overall cost of climate change to society. In principle, real options analysis (ROA) is an ideal decision-support tool because it handles the flexibility embedded in adaptation programmes, such as the ability to accelerate, delay, or rescale investment. However, ROA can be complex and resource-intensive, so it is not widely used for relatively small projects. Decision-makers need approaches that are simple enough to be useful for evaluating small- and medium-scale adaptation decisions yet retain a degree of economic rigour. This project evaluates several candidates and finds that most of the benefits of ROA can be achieved by making simple modifications to standard cost-benefit analysis (CBA).

The costs that decision-makers should try to minimise include the funds spent on adapting to climate change and the costs incurred by the community when weather-related events occur. Consider the example of a proposal to retrofit a building with new technology that will help its occupants cope with high outside temperatures. The costs occupants face depend on three types of inputs: the building's specification; the outside temperature; and economic factors such as healthcare costs and wages. Investment alters the first input, which lowers the costs faced by the building's occupants, but does not affect the other two inputs. The benefits of investment equal the avoided heat-related costs.

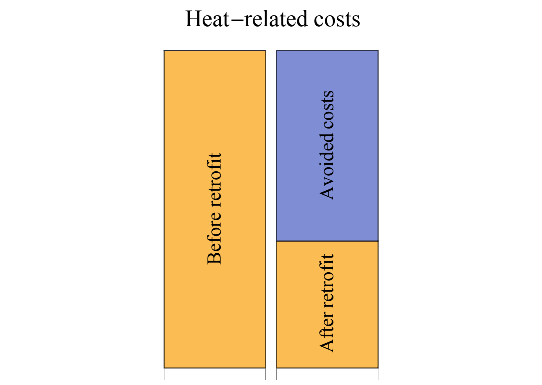

Figure 1: The benefits of improving a building's ability to cope with high outside temperatures

The two bars in Figure 1 illustrate these benefits. The left-hand bar shows heat-related costs on a particular day, assuming the building has not been upgraded. The right-hand bar decomposes this into heat-related costs on the same day, assuming the building has been upgraded, and the costs that are avoided once investment has occurred. The avoided cost (the blue portion of the bar) is what decision-makers should care about.

The heights of the bars in Figure 1 depend on the temperature and the level of the economic factors on that particular day. However, these are both uncertain when the retrofit project is evaluated. The avoided costs (the benefits of investment) are therefore uncertain as well. Decision-makers need to use the expected value of the avoided costs when they calculate the benefits of investment.

The heights of the bars will differ on different days, and climate change will also alter the distribution these have over time. Higher temperatures will become more likely and the distribution of the economic factor will also possibly change. That is, the expected value of the avoided heat-related costs (the benefits of investment) will change in response to climate change.

The decision-maker's goal should be to maximise the present value of these avoided costs minus the expenditure incurred. When calculating this present value, the decision-maker needs to specify the level of the discount rate. For publicly-funded projects, this will be an estimate of the social discount rate; for private projects, it will be an estimate of the opportunity cost of capital. A policy that maximises this present value will also minimize the present value of all climate-related costs to society. Provided ROA is implemented correctly, it will produce an investment policy that achieves this goal.

Decision-makers must also deal with the uncertainty around future climate change. However, decision-makers do know that they will know more in the future (they just don't know what it is that they will know!). This anticipation of uncertainty falling in the future complicates what is already a difficult problem. Decision-makers can invest now but run the risk of investing too much or too little. Alternatively, they can wait, observe how the climate has changed, and then invest the "correct" amount in adaptation. In the long run, delaying will lead to better outcomes, but society will incur the costs of delaying beneficial investment in the short run.

There are two situations in which decision-makers might look back and regret deciding to invest early. The differences between these two situations show why "distribution uncertainty" is so important when making investment decisions.

In the first case, the decision-maker's beliefs regarding the scale of climate change turn out to have been correct, but the "weather" (that is, the random draws from the climate distribution) was better than expected. For example, the weather may be unexpectedly mild after retrofitting new technology to a building. This situation is not problematic because the good luck will eventually run out and high-temperature days will occur at a frequency that aligns with the expectations at the time investment occurred. In the long run, the building's occupants will get the expected benefits from the upgrade.

In the second case, the decision-maker's beliefs turn out to have been incorrect. There may still have been fewer hot days than expected, but now it is because the change in the temperature *distribution* is less extreme than expected. The number of hot days is less than expected at the time the investment decision was made, but entirely consistent with the underlying climate distribution. This situation is bad because there is no good luck to run out. High-temperature days will continue to be infrequent and the building upgrade will not deliver the benefits expected.

The decision-maker is concerned about the second case. As decision-makers will learn more about climate change over time, delaying investment reduces "distribution uncertainty" and potentially reduces the risk of the second case above occurring. The decision-maker weighs the benefits of this risk reduction against the costs of the heat-related costs incurred in the meantime.

Many applications of ROA do not adequately handle situations in which uncertainty falls over time. The reduction in uncertainty is accompanied by a surge in volatility as the new climate information is incorporated in the expected value of future benefits. This temporary volatility surge needs to be incorporated in the models used to implement ROA.

ROA in this project

This research project extended standard ROA to cope with the situation in which uncertainty about climate change is expected to fall over time. Although the focus is on simple alternatives to ROA, we need something to measure these approaches' performance against. The benchmark used here is the socially optimal policy—that is, the policy that would result if decision-makers used "bespoke" ROA.

Consider, again, the example of retrofitting new technology to help a building's occupants cope with high outside temperatures. There are many different ways to express the optimal investment policy, but they all involve "trigger points." That is, decision-makers do not specify a fixed future date when investment should occur. The volatility in future climatic and economic conditions means that there will never be a single, fixed, optimal date when investment should occur. Instead, the best investment date depends on how climatic and economic conditions evolve. It is optimal to wait, continuously evaluate the project, and invest the first time that conditions merit it. If the climate deteriorates more quickly than expected, then it will be optimal to invest earlier than expected. If climate concerns are overblown, then it will be optimal to invest later than expected.

The most intuitive way to express the optimal trigger point is in terms of the project's benefit-cost ratio (BCR), which is the ratio of the present value of the project's benefits to its cost. That is, investment should occur as soon as the project's BCR is greater than some threshold, which is a function of the climatic conditions and the date.

Figure 3 shows how a BCR-based trigger point works. It is based on a simulated path for the economic and climatic factors which jointly determine the distribution of the future benefits from investing in an adaptation project. For each point on the path, we can calculate the BCR of investing in the project at that point in time, as well as the level that the BCR would need to exceed for investment to be socially optimal. The blue curve in Figure 3 shows that former, the yellow curve the latter. Initially, the project's BCR is too small for investment to be optimal (although it does exceed one, the threshold implied by standard CBA). The decision-maker waits until the blue curve moves above the yellow one. Only then is the BCR high enough to justify investment. For this simulation, investment occurs after waiting approximately 25 years (shown by the black dot in the graph). If we ran another simulation, we would get different paths for the economic and climatic factors, different paths for the project's BCR and the BCR threshold, and therefore a different optimal investment date. Setting a fixed investment date is not socially optimal.

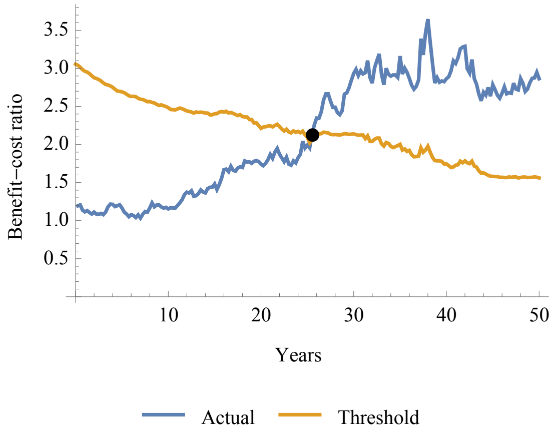

Figure 3: Optimal decision-making using trigger points

There are two important points to note. First, the optimal BCR threshold is much larger than one. That is, the present value of the project's benefits needs to be much larger than the project's cost in order for investment to be socially optimal. This reflects the value of the flexibility embedded in the project. The option to delay and learn more about climate change before investing is very valuable. If it is infeasible to carry out bespoke ROA, then it is important that our evaluation approach achieves similar results. The flexibility embedded in this project is too valuable to ignore.

The second important point is that the optimal BCR threshold (the yellow curve in Figure 3) has complex behaviour. In particular, it varies over time as the decision-maker's expectations about future climate change evolve in response to new information. It will be difficult to calculate in general, and too expensive to calculate for many small-medium scale adaptation projects. This brings us back to the central issue of this research project: How can we calculate the level of these trigger points given the analytical resources we can justify allocating to the project?

Alternatives to ROA

The aim of this research is to identify project evaluation techniques that are easy to implement and result in decisions that are close to the ones that emerge from ROA. Two conditions are needed for a project-evaluation technique to be easy to implement:

1. All of the required inputs should be readily available or easily calculated.
2. Converting these inputs into recommended actions should be straightforward.

Bespoke ROA sometimes fails the first test, but it almost always fails the second test.

Four simple rules emerged from the research described here as candidates to approximate ROA.

1. Standard CBA applied at fixed regular dates. The project is evaluated at fixed, regular intervals (perhaps every five or ten years) until the present value of the project's benefits is greater than the cost of completing the project, when investment occurs.

2. Constant BCR threshold. The project is evaluated at fixed, regular intervals until the ratio of the present value of the project's benefits to its cost is greater than some constant threshold. When using CBA, the decision-maker invests if the project's BCR is greater than one. This modification increases the threshold for the BCR to a constant that is calculated using a formula that depends on the characteristics of relevant economic factors.

These characteristics relate to how the economic benefits of the project fluctuate over time. For example, if identical weather events occurred on two different dates, one year apart, the economic impacts of the events would almost certainly be different. For example, when considering projects to reduce flooding risk, repair costs and per-capita disruption costs will change, as potentially will the number of people affected. The threshold we use for the BCR depends on the mean and standard deviation of the annual change in the economic impact of identical weather events. We usually have an estimate of the mean, and sometimes of the standard deviation, from when we carry out CBA.

Figure 4 shows suggested BCR thresholds for different combinations of the mean and standard deviation of the economic factor's annual growth rate. These thresholds are significantly greater than one. That is, a project's benefits need to exceed its cost by a considerable margin in order for investment to occur. This is consistent with the option to delay investment and learn more about the scale of climate change having considerable value.

		Average annual growth in economic factor			
		-1%	0%	1%	2%
Annual volatility in economic factor	6%	1.12	1.21	1.41	1.80
	8%	1.18	1.29	1.49	1.89
	10%	1.25	1.37	1.59	2.00
	12%	1.32	1.46	1.69	2.12
	14%	1.40	1.55	1.80	2.25

Figure 4: Suggested thresholds for a project's BCR

3. Constant IRR threshold. The project is evaluated at fixed, regular intervals until its internal rate of return (IRR, the discount rate that sets the present value of its benefits equal to its cost) is greater than some constant threshold. This threshold is calculated using a formula that depends on the characteristics of relevant economic factors.

When using CBA, the decision maker effectively invests if the project's IRR is greater than the discount rate (the social discount rate for publicly-funded projects, the opportunity cost of capital for private projects). This modification uses a threshold for the IRR that is larger than the social discount rate.

Figure 5 shows suggested IRR thresholds for the same range of parameter settings as Figure 4. The key insight this time is that the thresholds are significantly greater than the social discount rate (5% in this example). Once more, a project's benefits need to exceed its cost by a considerable margin in order for investment to occur. And, once again, this is consistent with the option to delay investment and learn more about the scale of climate change having considerable value.

		Average annual growth in economic factor			
		-1%	0%	1%	2%
Annual volatility in economic factor	6%	5.71%	6.04%	6.62%	7.40%
	8%	6.09%	6.43%	6.97%	7.68%
	10%	6.50%	6.85%	7.35%	8.00%
	12%	6.94%	7.29%	7.76%	8.36%
	14%	7.41%	7.76%	8.20%	8.76%

Figure 5: Suggested thresholds for a project's IRR

4. Fixed-delay comparison. A decision-maker using this rule invests if and only if the net present value (NPV) of investing immediately is greater than the NPV of waiting and investing after a fixed delay. For example, the decision-maker compares the NPV of investing now with the NPV of investing in two years' time and only invests now if the former NPV is larger than the latter one. The decision-maker should repeat this process at fixed regular dates until either the investment test is satisfied or the project is abandoned altogether.

Figure 6 assesses these four approaches against the two criteria above. That is, the required inputs should be readily available or easily calculated, and converting these inputs into recommended actions should be straightforward. The fixed-delay-comparison approach requires no additional inputs and even the approaches using BCR and IRR thresholds only require one additional input—which should be readily available from the data used to carry out standard CBA. None of the approaches require calculations that are more difficult than the ones used to carry out standard CBA. In short, all the approaches considered are viable alternatives to standard CBA.

Approach	Additional information needed	Complexity relative to CBA
CBA at fixed dates	None	None
Constant BCR threshold	Volatility of application-specific economic factor	Minor: BCR needs to exceed a fixed threshold, calculated using a simple formula
Constant IRR threshold	Volatility of application-specific economic factor	Minor: Replace social discount rate with a higher discount rate, calculated using a simple formula

Fixed-delay comparison	None	At each evaluation date, two implementations of CBA rather than one

Figure 6: Simple alternatives to bespoke ROA

These four rules are therefore simple enough to implement to be viable candidates for alternatives to bespoke ROA in situations where the cost of carrying out bespoke ROA cannot be justified. The key remaining issue is whether they deliver outcomes that are close to optimal. That is, how much welfare will be lost if decision-makers use one of these rules instead of bespoke ROA?

Evaluating these alternatives to ROA

This research project investigated the performance of these four rules using a theoretical framework that incorporates the economic and climate risk analysts face. This framework was used to estimate the maximum attainable welfare that representative adaptation projects can generate—that is, how much welfare is possible if decisions are made using bespoke ROA? Then we suppose the decision-maker uses one of the simple alternative rules to choose an investment date and ask a related question: how large is overall welfare if the decision-maker uses each of the four simple rules above? This allows us to estimate the amount of potential welfare that is lost by using each of the simple alternatives to ROA. If a particular simple decision-making rule performs well in the theoretical model, then we can have some degree of confidence that it will also perform well in the real world.

This process was carried out for three different types of adaptation projects. The projects were selected to be representative of the wide variety of available adaptation projects that decision-makers will confront in the future.

1. Simple timing options. For this class of projects, the decision-maker's most important decision is when to invest. Other considerations, such as the precise details of what to invest in, are less important. For example, the decision-maker might have the flexibility to decide when to retrofit a building to enhance its performance during high temperatures.

2. Ongoing investment. Some adaptation projects feature ongoing investment programmes involving relatively large numbers of relatively small investments. The rate of investment in these programmes is usually able to be scaled up or down reasonably easily and the project starts to generate benefits even before the overall programme is complete. For example, the decision-maker might be responsible for upgrading an urban stormwater system.

3. Scale-and-timing flexibility. For projects in this family, the decision-maker must choose when to invest and what to invest in. For example, the adaptation project might involve increasing the height of an existing sea wall, with the decision-maker having to choose the new height of the wall.

Each of the project-evaluation approaches considered here can be expressed in terms of a threshold for the BCR. That is, they imply that the decision-maker should invest in an adaptation project the first time the project's BCR is greater than this threshold. One way to evaluate the performance of the ROA alternatives is therefore to compare these implied BCR thresholds with the threshold corresponding to the socially optimal investment rule. This approach is illustrated in Figure 7, which superimposes the implied BCR thresholds for each rule on the optimal threshold. This particular case corresponds to an example from the first class of projects above. The results for the other project classes are similar.

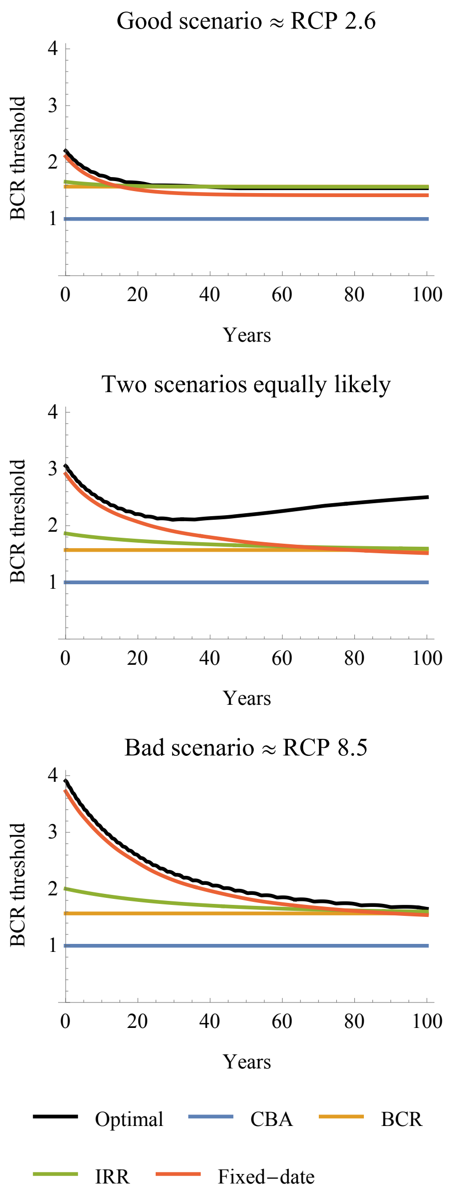

Figure 7: Comparing investment policies implied by simple evaluation approaches with the socially optimal policy

The example in Figure 7 assumes there are two possible climate scenarios (corresponding approximately to RCP 2.6 and RCP 8.5). At any point in time, the state of knowledge about climate change is summarised by the probability that the decision-maker attaches to each of the scenarios. As the decision-maker learns more about climate change, these probabilities change. There is therefore a different BCR threshold for each probability that the bad scenario holds. The graphs in Figure 7 show these thresholds for three different probabilities. The good scenario is certain to hold in the top graph, the bad scenario is certain to hold in the bottom graph, and the decision-maker believes the two scenarios are equally likely in the middle graph. In each case, the level of the BCR threshold changes as we look further into the future. (The dates on the horizontal axis of each graph are relative to 2023.)

In all cases, the fixed-delay-comparison approach implies the BCR threshold that is closest to the optimal threshold. Indeed, for the next few decades, it is remarkably close. Of the others, using the suggested constant threshold for the IRR will result in investment timing that is closest to the optimal policy. Overall, given that the various alternative approaches are all straightforward to implement, the fixed-delay comparison appears to be the best candidate for a simple alternative to bespoke ROA.

Ultimately, we are interested in overall welfare (the present value of benefits minus costs), not the timing of investment. After all, it is possible that a rule delivers timing that is close to optimal, but that the gap is still large enough to cause large welfare losses. Conversely, even rules that appear to cause large timing differences may still result in only small welfare losses. For this reason, this research project has also estimated the loss of welfare that would occur if each alternative decision-making approach was used instead of bespoke ROA.

This process involves two steps. First, calculate the level of welfare that results if the decision-maker adopts an optimal investment policy—that is, use ROA to solve for the optimal investment policy. Second, using the same framework, suppose the decision-maker uses one of the alternative rules instead. Third, compare the present values of benefits minus costs in the two cases.

When this approach is applied to the building-retrofit example, using standard CBA at five yearly intervals results in 69% of the maximum attainable welfare being achieved on average. This average is calculated over 180 different parameter settings, so standard

CBA performs reasonably well. When the decision-maker uses the constant BCR threshold, 96% of maximum attainable welfare is achieved on average across the same 180 situations. Using the constant IRR threshold achieves 97% of maximum attainable welfare and the approach using a fixed-delay comparison achieves 98% of maximum attainable welfare. Similar performance is achieved when these rules are applied to the other classes of adaptation projects ("ongoing investment" and "scale-and-timing flexibility").

There are four main results for practitioners to take away from this research:

- In all the cases considered in this research project, the value of the flexibility embedded in the adaptation project is economically significant. It is important that decision-makers attempt to exploit this investment flexibility in ways that benefit society as a whole.
- If standard CBA is used, then projects should be evaluated at fixed intervals. Decision-makers should resist the temptation to reevaluate a delayed project ahead of schedule in response to news that the project's benefits have increased.
- Constant BCR and IRR thresholds are better than standard CBA. A constant IRR threshold is preferred, because it implies a non-constant BCR threshold with qualitative properties that are similar to the socially optimal threshold.
- If even better performance is required, then decision-makers can use the fixed-delay-comparison approach. This simple rule delivers small welfare losses, even though the calculation is not much more complicated than standard CBA.

These insights are based on extensive sensitivity analysis that shows the results hold over a wide range of parameter settings for three broad classes of adaptation projects. However, the examples analysed are only snapshots of the wide spectrum of possible adaptation projects. **It is not appropriate to extend the results here beyond the broad classes of adaptation projects considered in this research**. For example, the fixed-delay-comparison approach works well here, but it may not work as well in other situations. In particular, we cannot be sure it will work as well with adaptation projects that have more complicated option structures. **There will always be a place for bespoke ROA, such as when analysing large or complex adaptation projects.**

Finally, although the focus of this research project was on simple alternatives to ROA for use in evaluating small-medium adaptation projects, the underlying framework could be used for larger projects. It was necessary to extend standard ROA in order to cope with the situation in which uncertainty about climate change is expected to fall over time. This new framework could be the basis of bespoke ROA for larger adaptation projects. This would involve retaining the parts of the current model that specify the uncertainty surrounding future economic and climatic conditions, but replacing the current stylised approach to estimating the annual flow of investment benefits with practitioner-generated benefit-flow functions. With such a change, the framework used here would be suitable for relatively large or complicated adaptation projects, for which simple rules may not be sufficient.

Additional resources

Guthrie, G. (2009). *Real Options in Theory and Practice*. Oxford University Press, New York, NY.

Guthrie, G. (2019). Real options analysis of climate-change adaptation: Investment flexibility and extreme weather events. *Climatic Change*, 156(1–2):231– 253.

Guthrie, G. (2023a). Optimal adaptation to uncertain climate change. *Journal of Economic Dynamics and Control*, 151:104621.

Guthrie, G. (2023b). Simplified real options analysis for climate change adaptation. https://deepsouthchallenge.co.nz/resource/report-simplified-real-options-analysis/