RISK-OPPORTUNITY ANALYSIS
IMPLEMENTATION GUIDANCE

For policy appraisal in situations of uncertainty, diverse interests, and structural change

Working Paper draft for consultation and review. Draft as at 27 November 2024.

Authors: Simon Sharpe and Anna Murphy

If you would like to share comments on this document, please email
annamurphy@scurveeconomics.org

Contents

1. INTRODUCTION

Risk-opportunity analysis: policy appraisal for situations of uncertainty, diverse interests, and structural change

Cost-benefit analysis (CBA) is a way of organizing information to inform a decision. It provides a consistent framework for assessing the positive and negative outcomes of one or more possible courses of action, supporting judgements about whether a proposed policy does more good or more harm, and informing choices between policy options. It is particularly suited to situations of marginal change, relatively certain outcomes, and relatively narrow sets of interests.

Risk-opportunity analysis (ROA) applies the same aims and principles to more complex situations: those in which change is structural, important outcomes are uncertain, and the interests at stake are diverse. In those situations, CBA can be misleading: providing a false sense of certainty, obscuring trade-offs between different interests, and overlooking dynamics that can strongly affect the outcomes of policies.[1]

This implementation guidance focuses on the application of ROA to policy in the field of the low carbon transition. However, it can also be applied to any areas of policymaking where there are complex decisions to be made. Section 2 provides guidance on when to use ROA, and when to use CBA.

What ROA can do, and what it cannot

ROA can provide a consistent basis for considering and comparing policy options. By providing a structured process for the assessment of dynamic effects and uncertainties, it can generate insights that might not have be gained otherwise. It can also enable productive conversations among analysts and policymakers, clarifying assumptions and expectations, and converting areas of disagreement into assets for the constructive exploration of uncertainty.[2] In its outputs, it can provide a highly informative set of information to decision-makers.

ROA cannot provide definitive answers to the questions of whether a policy does more good than harm, or which policy option is best. It cannot encapsulate the value of a policy in a single number. It is intended for use in situations where to do so would be meaningless or misleading. In situations where there are definitive answers, or where the value of a policy can meaningfully be expressed in a single number, it is sufficient to use CBA.

In the complex policy problems to which ROA is intended to be applied, the need for judgement is inescapable. The quality of the analysis will depend strongly on the level of subject-matter expertise that it incorporates. The decision-maker will not know whether they have made the 'right' choice until some time after the decision, if at all.

Proportionality in analysis

The time, effort, and other resources put into policy analysis should be proportionate to the significance of the policy decision that the analysis aims to inform. In this guidance, we suggest

---

[1] HMT, 2020. Green Book Review 2020: Findings and response. pp9-10, 14-15.

[2] We can say this with confidence because although ROA itself is a relatively new concept, its component methods are not new, and have been used widely in the public and private sectors.

different 'levels' at which ROA can be implemented, from basic to advanced, so that the appropriate level can be chosen to suit the policy problem and the available resources.

2. DETERMINING WHEN TO USE RISK-OPPORTUNITY ANALYSIS

It is likely to be appropriate to use ROA when any of the following three criteria are strongly true, or when more than one of them are true:

i) Desired or expected change in the economy is structural, not marginal.
ii) Important outcomes are subject to uncertainties are that are not meaningfully quantifiable.
iii) Important outcomes exist in a diverse range of dimensions.

A judgement on each of these criteria can be made by considering the issues outlined below.

Is intended or expected change marginal or structural?

CBA has been described as a 'marginal analysis technique', which is 'generally most appropriate where the broader environment (e.g. the price of goods and services in the economy) can be assumed to be unchanged by the intervention' and working 'less well where there are potential non-marginal effects or changes in underlying relationships.'[3]

ROA is most appropriate where the opposite is true: where the intervention can be expected to change the price, quality, type, or existence of goods and services in the economy, or to change the relationships between economic variables. Either of these forms of change can be understood as 'structural change'. In these situations, it is useful not only to estimate a policy's outcomes at fixed points in time (as in CBA), but also to consider explicitly how a policy may affect processes of change within the economy.

In the low carbon transition, policies that aim to introduce clean technologies to the market, or to support their diffusion and displacement of fossil fuels, are often likely to fall within the category of structural change.

Is relevant uncertainty high or low?

In most policy decisions, there are some factors or outcomes that are uncertain. 'Uncertain' in this document means 'not meaningfully quantifiable'.[4] If a policy's outcomes can be quantified in terms of their magnitude and likelihood with a reasonable degree of confidence, then they are not uncertain. If a policy's outcomes cannot be meaningfully quantified, then they are uncertain.[5] The same applies to factors that can influence the success of a policy: these may be predictable and quantifiable, or they may be uncertain.

---

[3] HMT Green Book, 2018 edition.

[4] Note that this document does not use the Knightian distinction between 'risk', which can be quantified, and 'uncertainty', which cannot. Our usage is instead consistent with the ISO's definition of risk as 'the effect of uncertainty on objectives', with this being understood in a negative sense, whereas 'opportunity' can be understood as the equivalent in a positive sense. This is consistent with the use of the word 'risk' in policy contexts, where it usually refers to factors that are not quantifiable.

[5] Note: Randomly assigning a probability distribution to a set of outcomes, or assigning a probability distribution to a set of probability distributions, would achieve quantification, but not meaningful quantification – since any arbitrariness in analytical inputs leads to arbitrary outputs.

The important question is the extent to which uncertain factors or outcomes are relevant to policy objectives. If they are incidental, the problem may be categorized as one of relatively low uncertainty, and it will be sufficient to use CBA in the analysis. (The practice in such cases can be to use the CBA for an 'economic' analysis, and to include uncertain outcomes in a separate 'strategic' analysis). If outcomes relevant to some of the primary economic objectives of the policy are uncertain, then an economic analysis that excluded them would risk being either irrelevant or misleading. In such cases, purely quantitative analysis is likely to be inappropriate. ROA can be used instead, to provide a structured approach to considering unquantifiable (as well as quantifiable) factors and outcomes.

In the low carbon transition, there may be high uncertainty around outcomes such as the number of jobs that a policy could create, the degree to which it will incentivise private investment, or the extent to which it will accelerate innovation or build competitiveness in domestic supply chains. In some decisions, these outcomes may be central to policy objectives; in others, they may be incidental.

How important are different dimensions of policy interest?

The value of a policy's outcomes can be expressed in a single metric if the policy has significant outcomes in only one dimension, or if the relative magnitude of outcomes in different dimensions is not important to the decision-maker.

In CBA, outcomes in different dimensions are aggregated into the single metric of monetary value. There is no single method for converting outcomes in different dimensions (such as, for example, lives saved, or jobs created) into the metric of money. Many such methods have been developed, and the choice of which method to use – a choice which determines the relative weighting that the analysis gives to outcomes in different dimensions – is, unavoidably, to some degree arbitrary.[6] If only one dimension of outcomes is of interest to the decision-maker, there are no drawbacks to this approach. If the relative magnitude of outcomes in different dimensions is not important to the decision-maker, then the element of arbitrariness is likely to be acceptable. If the decision-maker has a strong interest in the relative magnitude of outcomes in different dimensions, then this approach is likely to be inappropriate, and ROA – which expresses outcomes in different dimensions separately and explicitly – may be used instead.

In the low carbon transition, policies can have outcomes in dimensions including money, jobs, emissions, air quality, energy security, international influence, and the risks of climate change. In some policy decisions, it may be only one of these dimensions that is important, with others being incidental. In other decisions, several of these dimensions may be important.

[6] Note: the same conversion method may be used consistently, and this enables comparability of analysis of outcomes in a given dimension across different policy decisions. However, the choice of conversion method remains arbitrary, and this affects the weighing of outcomes in different policy dimensions for a given policy decision.

3. COMPONENTS OF RISK-OPPORTUNITY ANALYSIS

ROA has three distinct components:

i) Dynamic assessment: assessing the effect of policies on processes of change.

ii) Multi-dimensional assessment: presenting possible outcomes in different dimensions that are relevant to the interests of the decision-maker.

iii) Uncertainty assessment: considering how policy outcomes may be affected by factors outside the control of the decision-maker.

The steps for implementing each of these components are outlined in the three following sections.

We describe these as 'components' rather than 'stages' of the analysis since they can in principle be undertaken in any order. The findings from each may influence the others, so a certain amount of iteration is often likely to be beneficial. The order in which the components are listed above can be used as a default order for the process of analysis, and we generally assume this order in the guidance that follows. At the end, the findings from the three components are combined and presented in summary form in order to inform a decision.

The relationship between the components can be understood in terms of their different areas of focus. In the dynamic assessment and multidimensional assessment, the focus is on policy outcomes, which take place in the government's *transactional environment*. The transactional environment contains actors and variables that the government (or in the general case, an organisation) can influence. The dynamic assessment considers how policy can produce structural change in this environment, and the multidimensional assessment considers the different outcomes that result from that change.

The uncertainty assessment is focused on the *contextual environment*. The contextual environment is relevant to the government's interests but outside its influence; the government has to adapt and respond to changes that take place there. The uncertainty assessment considers how these factors could affect the performance of different policy options.

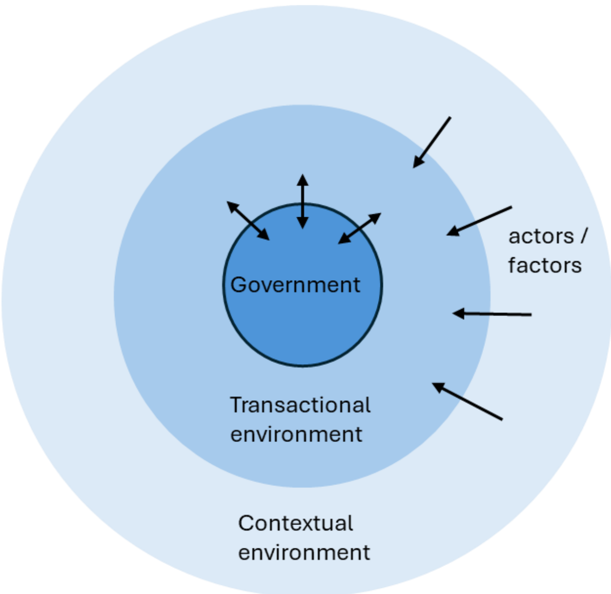

Figure 1: *transactional and contextual environments*

4. DYNAMIC ASSESSMENT

Purpose of this component of the analysis

The purpose of this component is to assess the likely effect of policy options on processes of change in the economic system of interest.

In complex systems, cause and effect are often disproportionate. Small inputs can lead to large outcomes, and vice versa. This happens because relationships between variables in a system can amplify some processes of change, and limit or prevent others. The aim here is to distinguish between policies likely to achieve a self-amplifying effect in line with their objectives, and those that could expend considerable resources while achieving little, or even be counterproductive in their effects.

Key distinctions

Marginal change and structural change

Marginal change is where the broader environment (e.g. the price of goods and services in the economy) can be assumed to be unchanged by the policy, and there are no changes in relationships between economic variables. Structural change is where the policy can be expected to change the price, quality, type, or existence of goods and services in the economy, or the relationships between economic variables. This component of ROA should be applied to problems of structural change.

Allocative efficiency and dynamic efficiency

Allocative efficiency is concerned with making the best use of a fixed set of economic resources at a fixed point in time.[7] Dynamic efficiency is concerned with creating new resources (such as technologies) or achieving structural change over the course of time. 'Decarbonisation at least cost' is an example of a dynamic efficiency objective. The aim of this component of ROA is to understand and compare the dynamic efficiency of policy options.

Introduction to the method

The core method to use for the dynamic assessment is systems mapping with causal loop diagrams. This approach is centred on the identification of feedback loops ('feedbacks' for short), which can either stabilise a system or drive it to change. There are two kinds of feedback:

- Reinforcing feedbacks (also known as positive feedbacks) amplify change in a given direction.
- Balancing feedbacks (also known as negative feedbacks) tend to limit change and stabilise the system, or a variable with it, around a constant state.

The behaviour of a system is driven by the various feedbacks within it, and the interactions between them.

For example, in the climate system, there is a positive feedback associated with the melting of sea-ice: rising temperatures melt sea ice, reducing its area; a small area of sea-ice reflects less

[7] Allocative efficiency can be defined in terms of Pareto optimality: a situation in which nobody can be made better off without someone else being made worse off.

sunlight, and more is absorbed by the sea instead; this contributes to rising temperatures. A balancing feedback is present in the Earth's emission of infrared radiation into space: the higher the Earth's temperature, the more radiation it emits, and the greater the cooling effect of this radiation. The fluctuations in the Earth's average temperature that have happened over millennia, while also remaining within boundaries that have been roughly stable over tens of millions of years, are an outcome of the combined effects of these and other feedbacks.

As can be seen from these examples, it is important to note that when reinforcing feedbacks are described as 'positive', and balancing feedbacks as 'negative', these labels do not mean 'good' and 'bad'.

Causal loop diagrams can be used first to understand the dynamics of the economic system of interest, and then to understand how a policy is likely to affect those dynamics. As well as distinguishing policies likely to be self-amplifying from those likely to be self-limiting, this method can also help to identify indirect relationships and unintended consequences of policies, and to assess whether policies implemented in combination are likely to be mutually supporting or mutually offsetting.

Steps in the method (intermediate level)

1. Determine an appropriate system boundary

All variables in the economy, society, and environment are ultimately connected in some way to each other, but policymaking cannot consider everything at once. To make progress in solving a problem, analysis must be focused as well as holistic.

The first step is to determine an appropriate system boundary for the analysis: within this boundary are variables that are important to the problem that policy is trying to address, and that can in some way be influenced by policy; outside the boundary are variables that this particular policymaking process is not trying to change. There is no objectively correct place in which to draw the boundary; it is a matter of judgement.[8]

As an example, figure 2 illustrates a possible system boundary that could be chosen by analysts working on the problem of power sector reform in the UK.

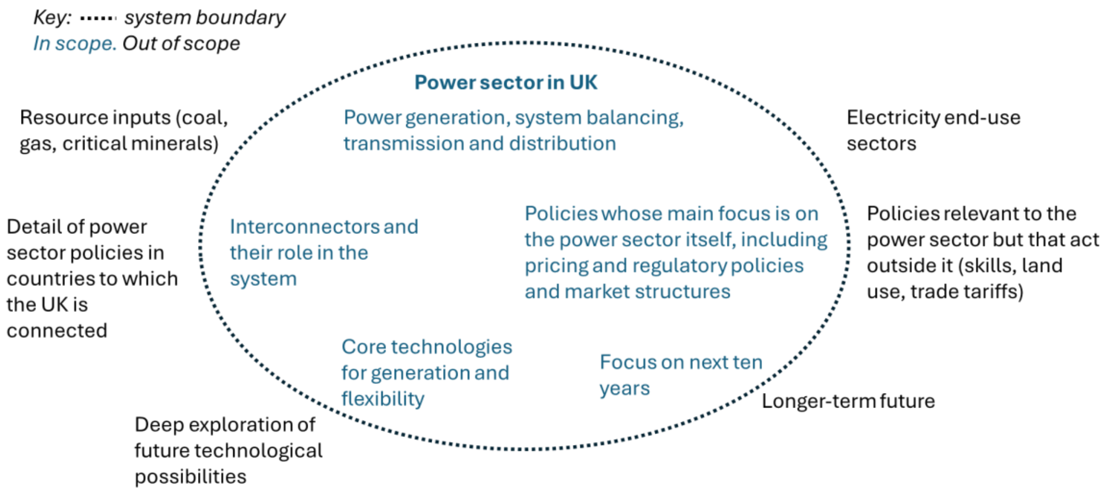

Figure 2: *example system boundary*

[8] [Placeholder: reference to REMA options assessment.]

2. Identify important variables

Identify variables that are important to the problem of interest, drawing on subject-matter knowledge. Define and name each of these so that it can be understood as something that could increase or decrease (it should be quantitatively measurable in principle, even though it may not be in practice).

For example, in a problem of power sector reform, ‘electricity price’, ‘gas power capacity’, ‘emissions’, ‘security of supply’, and ‘investment in offshore wind’ could all be variables. (Many others are possible). ‘Market design’ would not be appropriate as a variable, because it cannot be understood as increasing or decreasing. Variables should not be named as increasing or decreasing (e.g. ‘rising investment in offshore wind’) because they may move in either direction.

3. Assign directions to the relationships between pairs of variables

The direction of a relationship between a pair of variables can be either the same (an increase in A causes an increase in B; and a decrease in A causes a decrease in B) or opposite (an increase in A causes a decrease in B; and a decrease in A causes an increase in B).

For each pair of variables where one directly influences another, determine the direction of the causal relationship. To ensure rigour in the analysis, it is important that this is done by drawing on the best available information, and not arbitrarily or speculatively. The basis for determining the direction of a relationship can include:

- Logic: a relationship may be logically true without need for further explanation – e.g. increasing GDP will tend to increase GDP per capita, while increasing population will tend to reduce GDP per capita, all else being equal.
- Formulae or equations that are well-established and appropriate to the context.
- Quantitative evidence: system mapping with causal loop diagrams is a qualitative method, but its inputs should be consistent with any relevant quantitative information (provided there is confidence in the reliability of the latter).
- Other academic research findings.
- Expert knowledge: in some cases, this may be the best available information. If there is uncertainty over the level of confidence with which this information can be treated, this can be recorded for transparency, and so that the analyst can later consider whether a different input assumption would lead to different conclusions.

It is important to consider enough variables for a chain of causation to be clear. For example, indicating that the variable ‘certainty of revenues’ influences the variable ‘cost of renewable power generation’ may be insufficient. Adding the variable ‘cost of capital’ in between the other two variables makes the chain of causation clearer.

This information should be recorded in a table, so that the inputs to the analysis can be checked later if necessary. The table can take the form illustrated below.

Driving variable	Driven variable	Direction of relationship	Basis for relationship
A	B	Same	Logically true
B	C	Opposite	Formula X
C	A	Same	Academic study Y

Note that the 'driven' variable in one relationship can be the 'driver' variable in another relationship.

4. Map the system and identify feedbacks

Construct a causal loop diagram by drawing the variables and the relationships between them. Each variable should be joined by an arrow to each other variable with which it has a direct causal relationship. Each arrow should be labelled to indicate whether the direction of the relationship is the same ('S' or '+') or opposite ('O' or '-').[9]

Wherever there is a closed loop, in which a chain of causal influence leading from a variable comes back to the same variable, this represents a feedback. Each feedback should be identified as either reinforcing or balancing, and labelled accordingly (with 'R' or 'B').

The nature of the feedback can be known from the number of opposite causal links: an even number of opposite links indicates a reinforcing feedback; an odd number of opposite links indicates a balancing feedback. The identification can be checked by walking through the set of relationships. Using the relationships in the table above as an example: an increase in A causes an increase in B; an increase in B causes a decrease in C; a decrease in C causes a decrease in A. A loop that began with A increasing has ended with A decreasing: this must be a balancing feedback. The mapping of these variables and the resulting feedback is shown in figure 3.

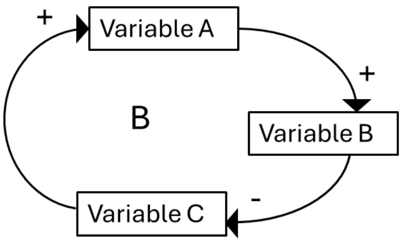

Figure 3: *example feedback*

Figure 4 below gives an example of a causal loop diagram. The diagram may have more or fewer variables, links, and feedbacks, depending on the complexity of the problem and the desired depth of analysis. The important attributes it should have are clarity in showing the causal relationships between variables, the directions of those relationships (matching the underlying knowledge), and nature of any resulting feedbacks (reinforcing or balancing).

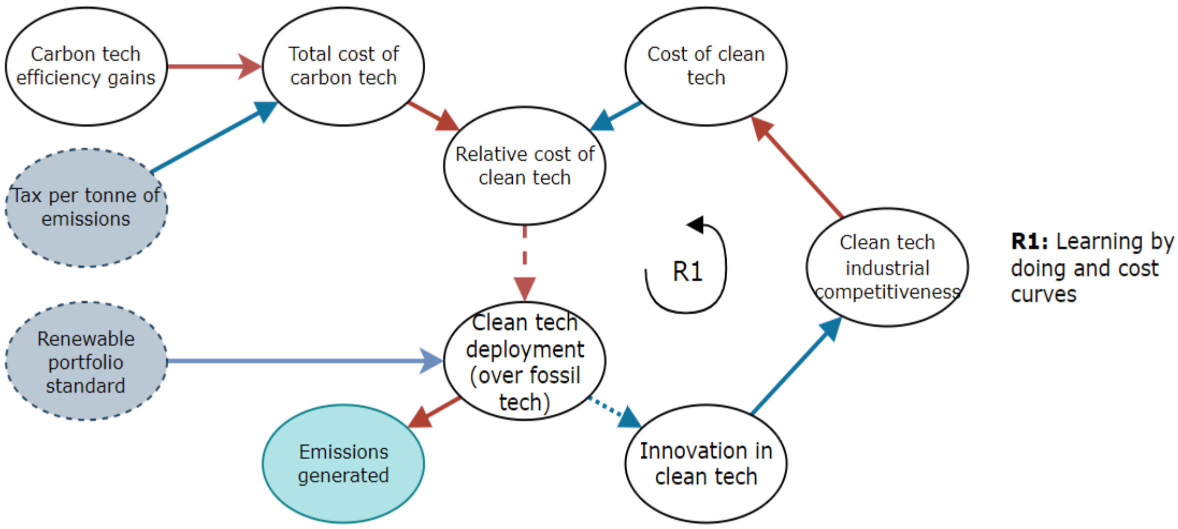

Figure 4: *example causal loop diagram. Figure key on next page.*

[9] Colour-coding of arrows can be used as an alternative to these forms of annotation.

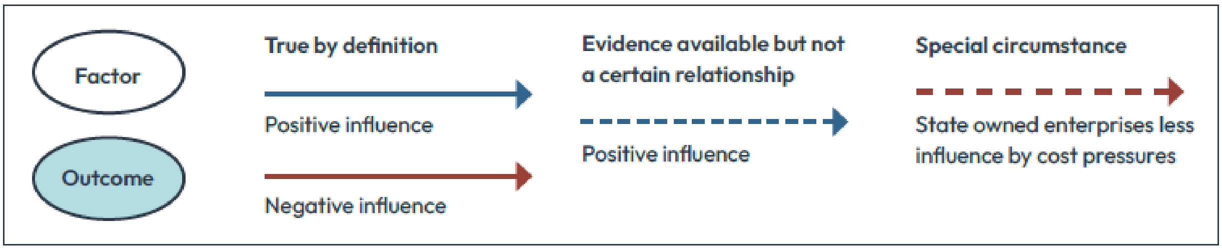

When a complex map is first drawn, the feedbacks may not be easy to see. Maps can be rearranged to make feedbacks more clearly visible, provided there is no change to the relationships between variables that they depict. (Changes to relationships between variables should only be made on the map if they reflect a change in understanding.) Feedbacks on the map can be labelled with numbers or names to distinguish them from each other.

Common errors to avoid in this step include:

- Showing incomplete or duplicate representations of causal relationships. For example, in the simple illustration given in figure 3 above, a line could be drawn from A to C, but if the real causal chain of influence is from A to B and B to C, then the line from A to C is both incomplete and duplicative, and makes the picture less clear.
- Including relationships whose direction is unjustified. Each relationship shown on the system map should be justified as described in step 3.
- Bias towards showing desired rather than likely influences of an intervention. Positive and negative, intended and unintended effects should be given equal attention.

There is also a risk of producing a system map of more complexity than can be usefully managed to inform the policy choice. The priority should be to identify the feedbacks that are most powerful in determining system behaviour and most relevant to policy interests. This is usually likely to be a relatively small number of feedbacks.

5. Determine the relevance of feedbacks to policy interests

For each feedback that has been identified as existing within the system, consider its relevance to policy interests, and determine whether it has a helpful or unhelpful effect.

For each reinforcing feedback, is it:

- Accelerating change in a desired direction (e.g. through innovation and diffusion of a low carbon technology); or
- Acting in a way that could turn a small problem into a large problem (e.g. a wage-price spiral, leading to ever higher inflation)?

For each balancing feedback, is it:

- Guiding or stabilising the system around a desired goal (like a thermostat in a home heating system); or
- Limiting change in a desired direction?

The feedbacks in the system, the identification of their effect as either helpful or unhelpful, and the reasons for that identification, can be recorded in a table. Using the two examples given above, this can be presented as follows:

Feedback	Effect	Reason
R1 (tech innovation)	Helpful	Drives clean tech cost reduction
R2 (wage-price spiral)	Unhelpful	Drives inflation

6. Consider how policy options can interact with feedbacks

For each policy option under consideration, identify the effect it would have on each of the feedbacks in the system. Add the policy to the system map by drawing the links between the policy and the variable (or variables) that it most directly influences. This should be done following the same principles as in step 3 above: there must be a good basis to justify the direction of influence.

Consider:

- Does it strengthen or weaken the helpful feedbacks?
- Does it strengthen, weaken or break the unhelpful feedbacks?
- Does it create any new feedbacks, and if so, are they helpful or unhelpful?

This step in the analysis may prompt the identification of new policy options that are worth considering. For example, if a particularly helpful feedback has been identified, it may be useful to consider policies that would most directly strengthen it. If a particularly unhelpful feedback has been identified, it is useful to consider what interventions could break any of the links in its loop.[10]

From this analysis, each policy option can be rated for its dynamic efficiency in terms of:

- Whether it is likely to be self-amplifying, self-limiting, or neutral (affecting neither kind of feedback) in its desired effect, or even counterproductive.
- Whether any self-amplifying or self-limiting effect is likely to be strong or weak.
- The degree of confidence in the above judgements.

The categorisation of self-amplifying, self-limiting, neutral, or counterproductive, should follow directly from the system map and the analytical steps outlined above.

The categorisation of these effects as strong or weak will involve judgement, and will draw heavily on subject-matter knowledge. It may involve consideration of the timescale on which different feedbacks operate, as well as how powerful they appear to be in driving or preventing change in the system, and how directly the policy intervention is likely to change them.

For example, for many clean technologies there is a reinforcing feedback between deployment, innovation, and cost reduction. The same deployment policy could be expected to act on this helpful feedback more strongly in a country with a large economy than in a country with a small economy. The effect of this feedback might be small over the time period of a year, but large over the timescale of a decade.

[10] It may be helpful at this point to compare the system of interest to system archetypes: patterns of dynamics that frequently occur, and typical interventions. See for example Kim, D.H., 1992. 'Systems archetypes: diagnosing systemic issues and designing high-leverage interventions' and 'Using systems archetypes to take effective action'.

The degree of confidence in these judgements may vary depending on the level of confidence in the direction of the relationships from which the system map has been constructed, and on uncertainties in the magnitude of any of the relevant causal relationships.

Tipping points

As part of step 6, the consideration of how policy options affect feedbacks, it may be useful to consider tipping points.[11] This is relevant to determining whether any self-amplifying effects of a policy are likely to be strong or weak.

A tipping point is where a small change in one variable can lead to a large change in the state and behaviour of a system. When a tipping point is crossed, the relationship between cause and effect is highly non-linear, or disproportionate.

In the low carbon transition, tipping points are likely to exist where clean technologies become preferable to fossil fuels, for a given purpose or within a given sector. This can include:

- Clean technologies becoming more profitable than fossil fuels, to industry or to investors.
- Clean technologies becoming lower cost than fossil fuels, for an equivalent use.
- Clean technologies beginning to provide a higher quality of performance than equivalent fossil fuel-using products or services.[12]

To include a consideration of tipping points in the analysis, the following questions should be addressed:

- Is there evidence that a threshold value of a variable can in principle be reached? (For example, does the scientific and economic knowledge of a clean technology support the idea that it can become lower cost than the fossil fuels against which it is competing?)[13]
- Is there good reason to believe that the system is likely to behave differently when that threshold is crossed? (This can be considered with reference to the system map that has been created in steps 1-4 above. Will this strengthen or weaken (or even create or break) important feedbacks?)
- What is the distance between the current value of the relevant variable, and its threshold value? (This will usually require quantitative analysis. For example, what is the difference in cost between a clean technology and its fossil fuel equivalent? Comparisons of the levelized cost of energy in the power sector, or equivalent comparisons in other sectors, an approach with which many government analysts are familiar, would be one way of answering this question.)
- Will the policy option under consideration cross the distance between the current and threshold values of the relevant variable, so that the tipping point is crossed?

---

[11] The Green Book (section A7.6) advises that where significant transformational change is an objective, it is important to research the likelihood, magnitude and location of tipping points.

[12] The examples of tipping points in the low carbon transition given here are not exhaustive. Tipping points may also exist in industry expectations, or in social preferences. The examples given here are those most likely to be identifiable and quantifiable with some confidence. For other examples and resources on this topic, see https://global-tipping-points.org/

[13] Note: Even if not, this does not exclude the possibility that the clean technology can be made more profitable to investors or more affordable to consumers, if policies are designed to achieve those effects.

If the last of these questions can be answered positively, then any self-amplifying effects of the policy are likely to be stronger. The policy's dynamic efficiency can be rated more highly, and with greater confidence. These findings should be included in the dynamic efficiency rating of policy options described above.

In an advanced level dynamic assessment, a wider category of 'sensitive intervention points' (SIPs) may be considered. A SIP is a point in a system where a small or moderately-sized intervention could drive disproportionately large changes. In addition to tipping points such as those described above, this category also includes critical nodes in networks, and critical points in time when there are windows of opportunity for change.[14]

Basic level dynamic assessment

A basic level dynamic assessment can be carried out by focusing on a single 'dominant' feedback loop. This can be either:

- A feedback that is currently dominating the behaviour of the system of interest, either preventing it from changing, or driving change in an undesired direction; or
- A feedback that the policy will aim to create or strengthen, which will be central to driving change in a desired direction.

If the focus is on an existing 'unhelpful' feedback (which could be reinforcing or balancing), then the analysis should show how this feedback will be changed by any policy options under consideration, as described in step 6 above. If the focus is on a 'helpful' feedback driving change in a desired direction (which must be a reinforcing feedback), then the analysis must show how this will be created or strengthened, as described in step 6. In either case, the causal relationships from which the feedback arises must be justified as described in step 3.

The function of this basic level dynamic assessment is to focus attention on the core process that most powerfully prevents or enables structural change in the desired direction.

Reinforcing feedbacks in clean technology development and diffusion (specific application of the basic level dynamic assessment to the low carbon transition)

In the context of the low carbon transition, it can be particularly useful to consider how policies can create and strengthen the reinforcing feedbacks of clean technology development and diffusion. These will often be the most powerful feedbacks aligned with the objective of driving a transition from fossil fuels to zero emission systems in a greenhouse gas-emitting sector of the economy. As such, they are a natural focus for a basic level dynamic assessment in ROA.

[14] Detailed guidance on these approaches is beyond the scope of this document. Discussion can be found in Mealy, P. et al, 2023. Sensitive intervention points: a strategic approach to climate action.

A simplified representation of the central reinforcing feedback in this process is shown in figure 5. Investment in the production of a technology tends to lead to innovation and improvement – both in the production process, and in the product or technology itself. Innovation, as well as the economies of scale that come from increasing production, tends to lower costs. Lower costs lead to higher demand for the technology. Higher demand drives increasing investment, and wider deployment.

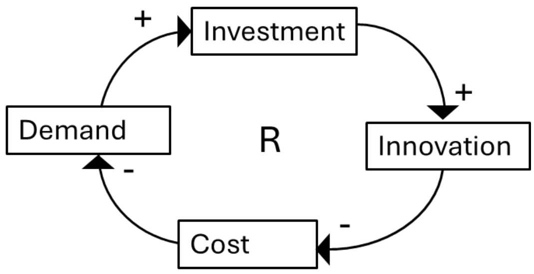

Figure 5: *reinforcing feedback in technology development and diffusion*

This core feedback can be complemented by other reinforcing feedbacks, such as:

- Complementary technologies: the more widely a technology is used, the more other technologies emerge that help to make it more useful.
- Social acceptance: the more widely a technology is adopted, the more it is noticed and considered acceptable or desirable, increasing its adoption.
- Financial risk perception: the more a technology is deployed, the less risky it may appear to investors, resulting in lower costs of capital, and more deployment.

These reinforcing feedbacks typically drive exponential growth in the deployment of a new technology in the early stages of a technology transition.

When the aim of policy is to strengthen these reinforcing feedbacks (as may be the focus of a basic or more advanced dynamic assessment in ROA), analysis should at a minimum consider the following questions:

- Is the technology consistent with the goal of zero emissions in the sector of interest? (Are there good reasons to believe it will play a significant role in a zero emission system in this sector? If not, supporting its development may divert resources from other technologies that are more important). Other consequences of supporting this technology should also be considered, in the multidimensional assessment. If there are significant negative consequences of a technology's use, for example in relation to land use, or air quality, or if there are strong limits on the scale at which it could be deployed, these may be reasons to decide that the technology is not consistent with policy goals, especially if better alternatives are available.
- Is there evidence to support the expectation that as investment and deployment of this technology increase, its costs will fall? (And if so, at what rate?) This does not happen for all technologies. Evidence can include historical trend data on technology deployment and costs, or comparative analysis across countries.[15]

[15] See Way, R. et al, 2022., 'Empirically grounded technology forecasts and the energy transition', for analysis of historical data trends for a core set clean technologies, and discussion of types of technologies for which cost reduction trends have and have not been observed. Where no data is available, it may be possible to make a judgement based on the extent to which the technology of interest shares components with other technologies for which cost reduction has been observed. (For example, offshore wind power generation shares many components with onshore wind). This is a live area of academic research.

- How exactly will any policy options under consideration strengthen this feedback? On which of its variables will they act, and how? It should not be assumed, for example, that putting a price on carbon will automatically lead to an increase in investment in the clean technology; this may instead incentivise fossil fuels to be used more efficiently. The effect of the policy should considered as described in step 6 above, and justified as described in step 3.

Note: This is a specific application of the basic level dynamic assessment to the low carbon transition. In policy problems where the aim is not to advance a technology transition, this particular set of feedbacks may not be relevant.

Advanced level dynamic assessment

Participatory systems mapping

If there is a high level of complexity in the economic system of interest, and a high degree of uncertainty around the causal relationships that exist and the feedbacks that are important, then it may be difficult or insufficient for an individual analyst to create and document causal loop diagrams and conduct the dynamic assessment following the steps outlined above. In such cases, the method can be enhanced by using participatory systems mapping, in which a group of analysts and experts in the system of interest construct the system map together through a series of interactive workshops.

This process can broadly follow steps 1-6 above, but with each of these steps being conducted in a participatory manner. The following elements should also be considered:

- Choice of participants: Aim to involve stakeholders with knowledge of all important parts of the system of interest. Small but diverse groups tend to work well.
- Facilitator: Using an experienced facilitator can help ensure a high-quality process.
- Verification: A complex systems map is likely to benefit from several rounds of review and revision by the group of participants.
- Analytical focus: In steps 5 and 6, it may be useful to investigate subsystems – sections of the systems map that are of particular interest. This can include: a) looking at the range of influences on a particular outcome variable of interest; b) looking at the range of variables affected by a particular policy option; and c) investigating interactions between different policy options that could be implemented together.[16]

Dynamic modelling

If there is a need to quantitatively assess the dynamic effects of policy options, and sufficient time and resources, simulation models can be used for this purpose. Simulation models represent processes of change over the course of time, and so can be used to test how policy options may affect these processes. Unlike optimisation models, which aim to show what is *desirable* – for example, a least-cost technology mix within a given sector, simulation models aim to show what is *likely* to happen with or without policy interventions.

[16] Detailed guidance on participatory systems mapping can be found in Barbrook-Johnson, P. and Penn, A.S., 2022. Systems mapping: how to build and use causal models of systems. Available to download.

There are two main structural forms of simulation model: system dynamics models, and agent-based models. System dynamics models may be appropriate in situations where the relationships between variables are reasonably well-known, and able to be confidently quantified. The feedbacks are built into the structure of a system dynamics model, and the model's outputs represent their combined effects. Agent-based models may be appropriate in situations where there is greater uncertainty about the structure of the system, but a reasonable degree of knowledge about how agents (such as businesses, consumers, or investors) make decisions. In an agent-based model, feedbacks emerge from agents' interactions with each other and with their environment. The purpose, core assumptions, and insights to be gained from these two types of model, alongside optimisation models, are compared in the table below.

Model category	**Optimisation**	**Simulation**	**Simulation**
Model structure	Technology cost optimisation	System dynamics	Agent-based
Model use	Identify technology mix to aim for	Test effects of policies	Test effects of policies
Core assumptions	Technology costs and constraints	Relationships between variables	Decision-making of agents
What it shows, as a contribution to dynamic analysis		How interacting feedbacks affect system outcomes	How interacting agents affect system behaviour

<u>Note of caution</u>: It should not be assumed that quantitative analysis is inherently superior to qualitative analysis. Models are simplifications of reality; their outputs depend on assumptions about structure, relationships, and variables. When using a model in policy appraisal, it is as important to understand what assumptions the model contains, and on what evidence they are based, as it is to understand the model's outputs. If the assumptions are not well-founded, or if the model omits important variables, its outputs may be misleading.

When time, resources, and technical capability allow, the ideal approach is for a model to be constructed by the analysts who wish to use it in the process of policy appraisal. This will ensure the most thorough understanding of the model's assumptions, capabilities and limitations. When this is not possible, and support with modelling is sought from academics or consultants, analysts should ask the following questions to determine how appropriate a model is likely to be in contributing to the analysis of a given policy problem:

- *Model structure:* Why is this model structure appropriate for this policy problem? How does it correspond to what is known about system behaviour, what is not known, and what we wish to find out?
- *Model resolution:* In what level of detail are the relevant market structures, technologies, policy options, and other variables and relationships of interest represented in the model? Which aspects of the system are not represented, and how will this affect the model's validity and usefulness?
- *Model validity:* What are the core assumptions that drive the model's behaviour? What is the theoretical or empirical basis for these assumptions?

It is also useful to consider whether a model has an appropriate degree of complexity. While the omission of important variables could detract from the model's ability to represent system behaviour accurately, equally there are risks from including too many variables, about each of which some assumptions must be made. The more assumptions that are made in model design, the more difficult it is to be sure of their combined effect. Models that are relatively simple in structure can, for this reason, be understood with a greater degree of confidence.

An overarching question to ask is: how will the use of this model contribute insights that add to those already gained from a qualitative dynamic assessment? If the answer to this question is not clear, then a different model (or no model) may be needed.

When quantitative dynamic modelling is used, its findings should be used to supplement the dynamic efficiency rating of policy options described in step 6 above. As well as comparing individual policies, the modelling can be used to compare combinations of policies, distinguishing those that are mutually reinforcing from those that are mutually offsetting.

5. MULTI-DIMENSIONAL ASSESSMENT

Purpose of this component of the analysis

The purpose of this component of the analysis is to present information on the outcomes of policy options in a way that enables the decision-maker to consider trade-offs between outcomes in different dimensions.

Key distinctions

Analytical judgement and political judgement

The purpose of analysis in policy appraisal is to inform a decision by making the pros and cons of different options as clear as possible. Analysis aims to be descriptive and objective, setting out the facts of a situation and the likely consequences of policy choices. Analytical judgement is exercised with the aim of doing this well. Political judgement is normative and subjective; it is concerned with what the government believes is good and wants to achieve. In this component of ROA, a clear separation should be maintained, so that the analysis does not encroach on matters of political judgement.

Implicit assumption and explicit consideration

The relative importance of outcomes in different dimensions has been assumed implicitly if it is no longer visible in the final information presented to the decision-maker. In this component of ROA, such implicit assumptions should be avoided, so that the relative importance of outcomes in different dimensions can instead be considered explicitly, by the decision-maker.

Introduction to the method

The approach in this component of the analysis is to use the best available information to assess the likely and possible outcomes of policy options in all dimensions that are relevant to policy interests, and to preserve the integrity of that information by keeping it in the metric most natural to each dimension.

The best available information may be found in different forms, including academic research papers, modelling studies, stakeholder consultations, the government's own knowledge and experience, and precedent from similar policy decisions in the UK and in other countries.

While some policy outcomes, such as a policy's immediate costs, may concern the present and therefore can in principle be precisely knowable, most policy outcomes – including all opportunities and risks – concern the future, and so are usually not precisely knowable. For that reason, this component of the analysis includes a consideration of uncertainty. The focus here is on the uncertainty around the outcomes of policy, which take place in the transactional environment, as shown in figure 1 above. (The effect of uncertainty in the contextual environment – factors outside the government's influence – is considered in the uncertainty assessment, described in the next section).

Since unquantifiable outcomes (opportunities and risks) may be as important to policy objectives as quantifiable outcomes (costs and benefits), these should be assessed in a structured way and on an equal basis; the analysis should not be limited to the latter.

Steps in the method

1. Identify a set of outcome dimensions of significant importance to policy.

The identification of outcome dimensions of significant importance to policy should be done in consultation with Ministers, at the beginning of the process.

Important dimensions include those in which the policy is intended to achieve some desired outcomes, and those in which the policy is likely to have significant outcomes which may be either positive or negative.

In this step, analysts should consider the policy objectives and interests of other government departments as well as their own.

2. For each outcome dimension, determine the part of the probability distribution that is of greatest relevance to policy objectives.

For each outcome dimension, the analyst should determine whether policy objectives are primarily concerned with: a) the expected or most likely outcome[17]; b) the worst case outcome (limiting its likelihood); or c) the best case outcome (maximising its chances)? Alternatively, it may be decided that all of these are of equal interest.

For example, in policymaking on the transition to clean electricity, the expected outcome may be most relevant in the outcome dimension of electricity prices (what prices will consumers typically pay?); the worst case may be most relevant for the outcome in the dimension of security of supply (how likely is a shortfall in supply, or another kind of system failure?); and the best case may be most important for outcomes in the dimension of innovation (if one of a portfolio of technologies being tested is successful, what new capabilities will this create?).

There are several possible ways of making this determination:

[17] Note: Depending on the shape of the probability distribution, the expected value and most likely value are not necessarily the same.

a) With reference to established policy objectives, which may clearly indicate a focus on creating opportunities, limiting risks, or improving expected outcomes.
b) Through consultation with stakeholders. This question can also be included in the consultation with Ministers referred to in step 1.
c) With reference to precedent, in the handling of similar policy problems.
d) With regard to legal obligations. For example, some health and safety laws require the probability of extremely bad outcomes (such as accidental death of a person at work, or the collapse of a bridge) to be kept below very low threshold levels, clearly indicating a focus on the worst-case end of the probability distribution.
e) Through mathematical analysis of the shape of the probability distribution. If the data suggests a normal distribution, expected values can be meaningfully calculated, and these may be an appropriate focus for the analysis. If the data suggests a heavy-tailed distribution (one where the probability of extreme outcomes does not decrease exponentially), there is a strong risk that an expected value would be misleading (or impossible to calculate); and in this case, a focus on the tail of the distribution (chances of success, or risk of failure) is likely to be more appropriate.[18]

Whichever of these methods is used, the reasons for the choice of which part of the probability distribution is most relevant to policy objectives in each outcome dimension should be clearly and transparently documented. This choice is important, as it will influence the way that options are assessed (see next step), and the way they are finally presented to inform a decision (see Section 7).

Figure 6 presents a stylised illustration of the relationship between probability and impact (magnitude) for a given outcome of an imaginary policy, with the policy interest in different parts of the probability distribution represented by different actors.[19]

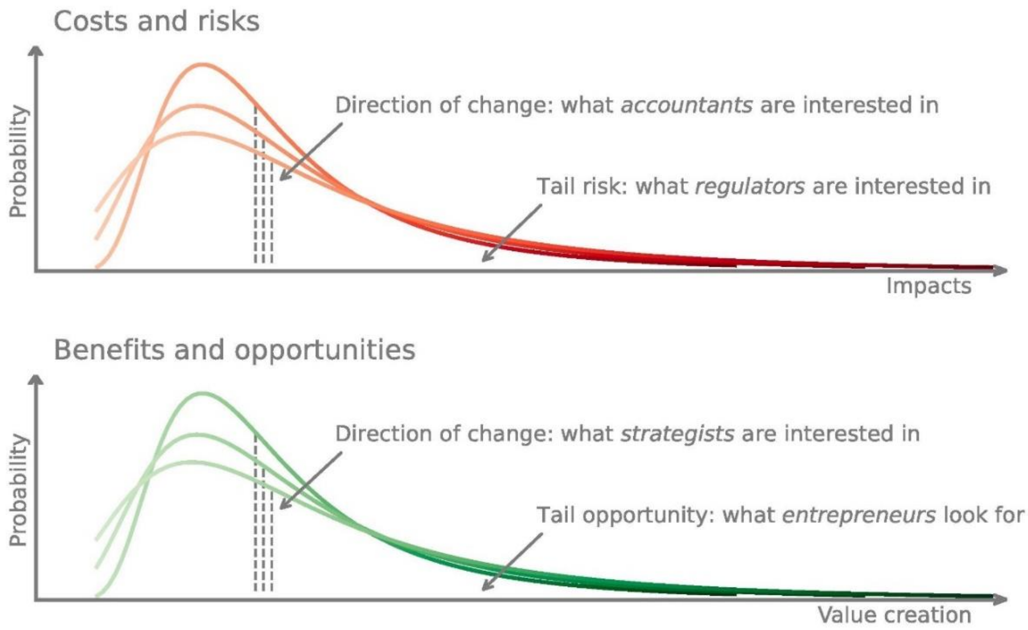

[18] Pollitt, H., and Mercure, J-F., forthcoming.
[19] Figure source: Mercure, J-F., et al, 2021. Risk-opportunity analysis for transformative policy design and appraisal. Note: the use of actors to represent interests in different parts of the probability distribution is illustrative; in reality, their interests may differ.

Figure 6: *probability, impact, and policy interests. Figure note: In each of two dimensions, one representing costs and risks, the other representing benefits and opportunities, the 'most likely' magnitude of outcome is found at the point along the x-axis corresponding to where the coloured line is highest on the y-axis. Larger or smaller magnitudes of outcome can be found either side of this point, with the probability decreasing as more extreme values are reached. Policy can change the shape of the distribution (coloured line), making the most likely outcome larger or smaller in magnitude, or making extreme outcomes more or less likely.*

3. For each outcome dimension, determine the timeframe that is of greatest relevance to policy objectives.

A point (or points) or period in time of interest to policy must be identified for each dimension, so that any quantitative values can be calculated accordingly. These time points or time periods should be chosen so as to be maximally relevant to policy objectives, and consistent with any available information. (They may be adjusted after step 4 discovers what information is available).

The time focus does not have to be the same for each dimension. There must be consistent use of a time point or period *within* each dimension, so that different policy options are compared on an equal basis.

The reasons for the choice of time points or periods should be clearly documented.

4. Assess expected (or worst case / best case) outcomes of policy options in each dimension by drawing on the best available information, in whatever form it exists.

This step of the process is, substantially, the same as the gathering and assessment of information that is done in CBA before the conversion of expected policy outcomes into monetary value. The aim is to understand the potential outcomes of policy options as thoroughly as possible. As noted above, the best available information could come from the government's own knowledge and experience, stakeholder consultations, academic studies, or other sources.

The differences from CBA in this step are as follows:

- The information on policy outcomes in each dimension should be kept in its original metric (the one most natural to its dimension).
- In each dimension, the information should be assessed with a focus on the part of the probability distribution identified as being of most relevance to policy objectives (as described in step 2).
- Qualitative information should be given equal weight to quantitative information, in the assessment. If policy outcomes in a given dimension are uncertain (not meaningfully quantifiable), the best available information should be used to describe the expected, worst case or best case outcomes in qualitative terms.

This assessment should be undertaken with reference to the time points or periods determined in step 3.

5. Combine this information with findings from the dynamic assessment.

Information gathered in step 4 should be combined with the findings of the dynamic assessment to determine how the effects of policy options in each of the relevant outcome dimensions are likely to change over time.

The outcome of a policy option in each dimension should be assigned a direction (increase or decrease) and a description of its main dynamic effect. 'Reinforcing' means the policy is likely to create or strengthen a reinforcing feedback such that change in this direction is likely to be self-amplifying. 'Balancing' means the policy is likely to create or strengthen balancing feedbacks so that change in this direction is likely to be self-limiting. 'Neutral' means neither dynamic applies: the effect of the policy is likely to be stable over time.

The table below illustrates how this combined information should be presented.

	Dimension 1 e.g. costs	**Dimension 2 e.g. solar PV deployment**	**Dimension 3 e.g. price of electricity**	**Dimension 4 e.g. security of supply**	**Dimension 5 e.g. innovation**
Point of interest	*Expected value*	*Expected value*	*Expected value*	*Tail risk (of system stress event)*	*Opportunity (of tech improvement)*
Time focus	*Present and next X years (discounted)*	*Over next X years (cumulative)*	*In X years' time*	*In any given year*	*Over next Y years*
Policy A	+£Xm Neutral	+X GW Reinforcing	-£X/MWh Reinforcing	+X% probability Reinforcing	Strong chance Reinforcing
Policy B	+£Ym Neutral	+Y GW Balancing	-£Y/MWh Balancing	-Y% probability Balancing	Unlikely Neutral
Policy C	+£Zm Neutral	+Z GW Neutral	+£Z/MWh Neutral	-Z% probability Neutral	Medium likelihood Reinforcing

Notes on the content of this table:

- The outcomes in each dimension are shown in the metric most natural to that dimension. In this example, costs are shown in monetary terms (GBP), solar photovoltaics deployment is shown in GW of capacity, the electricity price is given in £/MWh, the risk of a system stress event is measured in terms of its probability, and innovation is described in qualitative terms because the likelihood of some useful technological progress being made cannot be meaningfully quantified. Note that each of these outcome dimensions is qualitatively different from the others.

- Each cell in the table contains information on: i) the expected, worst case, or best case outcome at the specified time point or period. This includes a '+' or '-' sign to indicate the direction of the change (increase or decrease). (In the innovation column, since change in only one direction is possible ('improvement'), this is indicated in the column header); and ii) the main dynamic effect of the policy in relation to that outcome (reinforcing, balancing, or neutral, as described above).

Notes on the interpretation of this table:

- Presenting the information in this way makes the trade-offs between outcomes in different dimensions clearly visible. Policy A is good for solar deployment, innovation, and cheap electricity, but is bad for security of supply. Policy B is less good (dynamically) for solar deployment, cuts the price of electricity but with less favourable dynamics than Policy A, and is expected to achieve nothing in terms of innovation, but improves security of supply. Policy C performs moderately well in most dimensions, but increases the price of electricity. The three policies also each have a different cost.

- There is no objectively correct answer to the question of which of the three policies A, B, or C is best. That depends on the relative importance of each dimension to the decision-maker. This is a political choice, not an analytical choice.

Distributional impact assessments

Distributional impact assessments consider how the costs and benefits of a policy affect different groups within society. The Green Book advises that distributional analysis is important where there may be significant redistributive effects between different groups within UK society, resulting from a proposed policy, and that the level of detail and complexity devoted to this analysis should be proportionate to the likely impact on those affected.

When distributional impact assessment is used as a complement to CBA, it is appropriate for this to take a quantified and monetised form. An assessment of a policy's effect of a policy on the incomes of different social groups may use an equivalised income approach to incorporate some other effects of the policy within the analysis.

When distributional impact assessment is used as a complement to ROA, it may be used in relation to any (one or more) of the outcome dimensions of interest. For example, low carbon transition policies may have distributional effects in dimensions including costs (differently affecting income groups), jobs (differently affecting regions of the country, as well as sectors), and air quality (differently affecting age groups). The principle of whether, where, and in what level of detail it should be applied remains the same: it depends on the likely significance of a policy's redistributive effects between different groups within UK society, in any given outcome dimension.

If distributional impact assessment is to be used with ROA in relation to more than one outcome dimension, it should be done separately for each. This maintains consistency with the ROA approach of presenting outcomes of policy in each dimension in their original metrics, ensuring a maximum of transparency and allowing trade-offs between outcomes in different dimensions to be considered explicitly. The distributional impact assessment may be quantitative or qualitative, depending on the nature of the best available information.

The findings of a distributional impact assessment used with ROA should be considered together with the findings of the dynamic assessment, to identify whether the dynamic effects of the policy are likely to cause any distributional effects of the policy to increase or decrease over time. Beyond identifying the immediate impact of the policy on different social groups, this may be used to consider the potential cumulative impact of the policy over time on communities or regions.

Basic and advanced level multi-dimensional assessments

In the multi-dimensional assessment there is no difference in method between a basic, intermediate, or advanced level of analysis. The differences are only of breadth and depth: the number of dimensions that are considered, and the depth of analysis undertaken within each one. As in the other components of ROA, the extent of the analytical work should be proportionate to the significance of the policy decision that it aims to inform.

6. UNCERTAINTY ASSESMENT

Purpose of this component of the analysis

The purpose of this component of the analysis is to consider how policy outcomes may be affected by factors outside the control of the decision-maker.

As well as developing a better understanding of the pros and cons of the policy options originally being considered, this may lead to their adaptation, or to the identification of new options.

Key distinctions

Probability and uncertainty

If the probability of an event can be quantified, then for the purposes of this guidance, it is not uncertain. Uncertain factors are those whose probability cannot be confidently quantified.

Transactional environment and contextual environment

(See above.) This part of the analysis is concerned with the contextual environment: where there are factors that affect the performance of policies and the achievement of policy objectives, but over which the government has no control and little if any influence.

Introduction to the method

The core method to use for the uncertainty assessment is scenario analysis. This is a structured approach to imagining and considering different possible futures, in ways that are relevant to policy objectives. It involves identifying relevant uncertainties; developing a range of possible future scenarios that could materialise; and considering how alternative policy options[20] would perform in those scenarios – and how they could be improved.

Scenario analysis is primarily a qualitative method, because it is concerned with factors that are important and often unquantifiable, but it can incorporate quantitative information when this is available.

Steps in the method (intermediate level)

1. Determine the scope of the exercise

[20] Scenarios can also be used to test strategies. A strategy is usually supported by many policies, but there is no unambiguous division between one and the other. Here we refer to policies, or policy options, rather than strategies, for consistency with the rest of this guidance.

If the default order of analysis in ROA is being followed and this component is being carried out last, then the scope of the exercise will already be clear. If this component is being carried out first, then the first step is to determine the scope: this means clarifying what policy or policy options are being considered, for what purpose, and in what context. As in step 1 of the multi-dimensional assessment, Ministers should be consulted on policy objectives, if these are not already well-established.

2. Identify critical uncertainties

Identify uncertain factors that could affect policy outcomes. This process of identification should at a minimum include consultation with analysts and policy teams working in the relevant area. It should ideally include consultation of wider stakeholders or experts, and review of relevant literature.

Plot these factors on a grid of uncertainty against impact, as shown in figure 7, such that their relative positions reflect their relative uncertainty and their relative impact on policy objectives. This positioning is a matter of judgement, and ideally should not reflect the judgement of a single analyst. While the extent of participation in the analysis must follow the principle of proportionality, at the intermediate level this step (and the following steps, 3-6) should include at least several analysts and policy officials with a strong understanding of the problem at hand.

Note that in this grid, the y-axis is drawn to represent increasing uncertainty, not increasing certainty. This reflects the fact that in scenario analysis, uncertainty is a resource: if highly uncertain factors are not considered, the process will be less useful. Uncertainty can reflect disagreement about what is likely among members of the analytical team or among external experts; or it can reflect a shared sense of uncertainty.

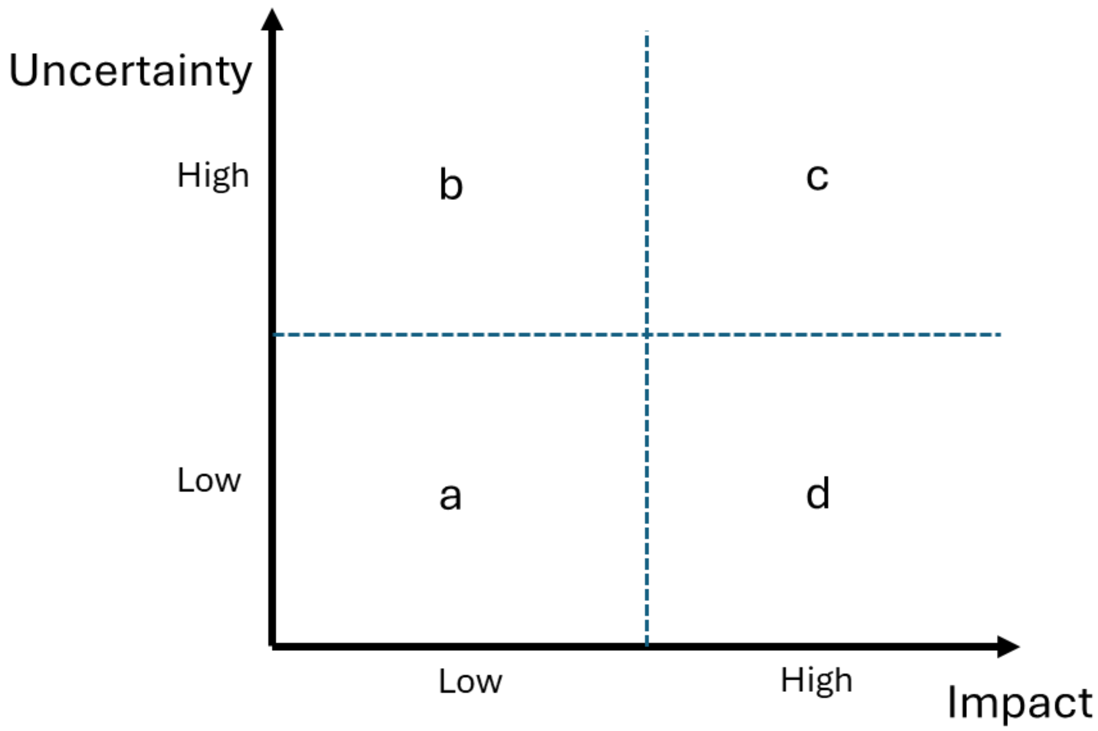

Figure 7: *uncertainty and impact grid*

Factors identified as having relatively high uncertainty and high impact of policy objectives (quadrant c) should be taken as the focus for the remaining steps in the scenario analysis.

Factors identified as having relatively low impact on policy objectives (quadrants a and b) can be deprioritised. Those having relatively low uncertainty but high impact on policy objectives can be considered through incorporation into the multi-dimensional analysis – with an iteration of that component, if necessary.

3. Cluster the critical uncertainties and select scenarios

The critical uncertainties (factors) identified in step 2 should be grouped into clusters representing two or three main categories of uncertainty. If factors are closely related such that change in one would have a predictable effect on change in the other, these factors should be included in the same cluster.

Each category of uncertainty should be considered as an axis, such that it can be realised as either 'high' or 'low', 'strong' or 'weak', 'fast' or 'slow', etc, in different possible futures. In this way, the two or three main categories of uncertainty correspond to four or eight possible scenarios (each scenario containing each category of uncertainty in one of its two possible states).

From this set of possible scenarios, choose the three or four scenarios that are the most relevant and challenging to policy objectives, as well as being plausible. Challenging can mean either that they make policy objectives difficult to achieve, or that they differ strongly from the dominant assumptions about the future conditions in which the policy will be implemented.

Note: the aim is not to select a scenario that is the most likely to be realised.

4. Develop scenario narratives

Develop each of the chosen scenarios into a narrative about the path from the present to the future. These narratives should include known pre-determined factors, as well as the uncertain factors that are the main focus.

Quantitative analysis can be included if it is available; this should be driven by the narrative (which can deal with greater uncertainty), and should not itself determine the narrative.

5. Build a scenario impact matrix

Consider how each policy option would perform in each scenario. As in previous steps, this step should involve consultation with analyst and policy colleagues, and ideally with external experts and stakeholders. At the same time, consider how each policy could be adapted to increase its chances of success in each scenario, and what the cost or difficulty of that adaptation might be.

Use a scenario impact matrix, as shown below, to summarise the findings.

	Scenario 1	Scenario 2	Scenario 3	Scenario 4
Policy A	effect of scenario 1 on performance of policy A, and potential adaptation	effect and adaptation	effect and adaptation	effect and adaptation
Policy B	effect and adaptation	effect and adaptation	effect and adaptation	effect and adaptation
Policy C	effect and adaptation	effect and adaptation	effect and adaptation	effect and adaptation

Note: this step in the process is usually found to be most productive when participants consider the performance of the set of policy options in each scenario, one scenario at a time (i.e. working down each column before moving to the next column, not working along each row before moving to the next row.)

Detailed notes should also be kept, to record the issues considered and to justify the summary content included in the table.[21] These should refer to supporting evidence where it is available, and acknowledge uncertainty wherever it exists.

6. Draw conclusions

Rate each policy's performance under each of the scenarios – for example, as good, poor, or moderate. These ratings should take into account the potential to adapt the policy to improve its performance in a given scenario.

For each policy, use its performance across the set of scenarios to determine whether it is:

a) Robust: it performs well in all scenarios; or
b) Resilient: it is relatively easy to adapt as scenario uncertainties become resolved; or
c) Contingent: it performs well only if one particular scenario is realised (representing a strategic bet on this future).

Which of these three characteristics is preferable depends on the nature of the problem and the objectives of the decision-maker. (As with other components of ROA, where there is no objectively 'correct' answer, the aim is to make the choices and trade-offs as clear and transparent as possible to the decision-maker.)

Note: if at this stage, the scenario analysis has led to the identification of policy options that were not previously considered but appear likely to perform well, these options may be introduced into the other components of the analysis through an iteration of the dynamic assessment and multi-dimensional assessment.

A note on participation

Scenarios are especially valuable for those who participate in the process of creating them. The process is designed to challenge assumptions, and this is more difficult to achieve with respect

[21] The scenarios and the detailed notes accompanying them may prove useful beyond informing the immediate policy decision that is the intended focus of the analysis. The process of developing scenarios can help analysts become better able afterwards to identify indicators of change that signify a direction of travel consistent with one scenario or another. This can support earlier identification of new trends, and the adaptation of policies or strategies with less difficultly or lower costs and risks.

to non-participants. Since Ministers make policy decisions but cannot be involved in the underlying analysis (for reasons of time), the best alternative to Ministerial participation must be found. Three elements should be considered:

- If possible, Ministers should be consulted at the beginning of the analytical process, confirming policy objectives and outcome dimensions of significant interest (as in step 1 of the multi-dimensional assessment), and at the same time identifying any uncertain factors that Ministers consider important. The latter can then be incorporated into steps 1 and 2 of the uncertainty assessment.
- The senior official(s) responsible for making the policy recommendations to Ministers at the end of the analytical process should be involved or consulted in each of the steps of the uncertainty assessment described above.
- Analysts should aim to demonstrate the rigour with which the process has been followed. This includes demonstrating: a) the quality of input information to the scenario analysis; b) the quality (including expertise and diversity) of participation; and c) the extent to which the scenarios considered challenged the dominant narrative and assumptions.

Basic level uncertainty assessment

A basic level uncertainty assessment can be carried out by considering and documenting the answers to three questions, with respect to each policy option:

a. What could happen, outside the government's control, that would most strongly prevent this policy from achieving its objectives?
b. If that were to happen, could the policy be adapted to meet its objectives, and if so, how, and at what cost?
c. What could happen, outside the government's control, that would most strongly increase the extent to which this policy achieves its objectives?

Question (a) should be addressed for each policy option independently. This means that the single, critically important, uncertain factors identified for each policy may or may not be the same. As a subsequent step, the effect of each of the uncertain factors that has been identified should be considered with respect to each of the policy options.[22] The findings should then be summarised in a table equivalent to the scenario impact matrix shown in step 5 above, with single uncertain factors (in their 'worst case' state) taking the place of scenarios as column headings.

If the 'best case' outcomes are particularly relevant to policy objectives (as determined in step 2 of the multi-dimensional assessment) then this process can be repeated for question (c).

[22] Note: This assumes a very small number of policy options are being considered (and the number of uncertain factors is either the same or smaller). If a large number of policy options are being considered, the combination of policy options and uncertain factors may be too large for this approach to be practical. In this case, the complexity of the problem is likely to justify an intermediate-level uncertainty assessment, including the prioritisation and clustering of uncertainties described in steps 2 and 3 above.

Conclusions as to the performance under uncertainty of the policy options should then be drawn in the same way as described in step 6, above.

Advanced level uncertainty assessment

An advanced level uncertainty assessment can follow the same steps 1-6 as described above, but with greater participation of experts and stakeholders, and more extensive consultations among participants. An experienced facilitator of scenario exercises should be used in this case. Steps 2-6 may be conducted in facilitated workshops, and these may be supplemented with: i) interviews with individual participants ahead of workshops, to gather perspectives and insights, and ii) a process of verification after workshops in which scenarios' plausibility, internal coherence, and factual foundations are checked and confirmed with participants (after step 4), and the conclusions of the process are reconfirmed with participants at the end of step 6.

As in the dynamic assessment, it is possible to use quantitative modelling as an input to the uncertainty assessment, to complement the primarily qualitative process described above. A Robust Decision Making approach to modelling can be used – following a consultative, participatory approach to define objectives, uncertainties, and policy or strategy options – to generate a large number of possible scenarios, estimate the performance of policies (or strategies) within those scenarios, and evaluate trade-offs between policy or strategy options.[23] An advantage of this approach is that computational power can be used to explore a larger number of scenarios than would otherwise be possible, and statistical analysis can be used to systematically explore the robustness of policies under the influence of different combinations of factors. A risk with this approach is that important uncertainties may remain outside the model: necessarily, a quantitative model is a simplification, and is more limited than a qualitative narrative in the possibilities that it can describe. Given this, when quantitative modelling is used in the uncertainty assessment it should be as a complement to the core method of scenario analysis described above, and not as a replacement.

7. COMBINING THE COMPONENTS

As has been noted previously, work on any one component of ROA may lead to the identification of new policy options to be considered, adaptations of the policy options being considered, or new factors to be included in other components of the analysis. For that reason, a certain degree of iteration between the components is likely to be helpful before entering into this final stage of the analysis.

The purpose of this final stage is not to add new information, but to combine the most important findings from each of the three components of ROA into an overview that is easily accessible to the decision-maker.

---

[23] Workman, M. et al., 2021. Climate policy decision making in contexts of deep uncertainty – from optimisation to robustness.

This overview should be supported by detailed notes and tables presenting findings from each of the three analytical components, as described above.

ROA grid

The overview of findings is presented in an ROA grid: a table constructed by adding key findings from the uncertainty assessment to the table combining information from the dynamic and multi-dimensional assessments. This should be presented as shown below.

	Dimension 1 e.g. costs	**Dimension 2 e.g. solar deployment**	**3**	**4**	**5**	**Performance in relation to critical uncertainties**
Point of interest	*Expected value*	*Expected value*	*Exp*	*Risk*	*Opp*	
Time focus	*Present and next X years (discounted)*	*Over next X years (cumulative)*	*In X years' time*	*In any given year*	*Over next Y years*	
Policy A	+£Xm Neutral	+X GW Reinforcing	-£X/MWh	+X%	Strong	*Contingent* Only performs well if X happens
Policy B	+£Ym Neutral	+Y GW Balancing	-£Y/MWh	-Y%	Unlikely	*Resilient* Can be adapted without difficulty to what happens with X, Y or Z
Policy C	+£Zm Neutral	+Z GW Neutral	+£Z/MWh	-Z%	Medium	*Robust* Performs well whatever happens with X, Y or Z

Notes: Information has been removed from columns 3, 4, and 5 to make the rest of the content more easily visible.

The content of all columns except the final column should be the same as in the table constructed in step 5 of the multi-dimensional assessment (after any changes to this table that may have resulted from subsequent iteration).

The final column, 'performance in relation to critical uncertainties', should include:

- A single-word description of the policy's performance under uncertainty (robust, resilient, or contingent, as determined in step 6 of the uncertainty assessment).
- A brief mention of any critical uncertainties that are strongly relevant to the policy's performance. For example, if policy B is resilient to some of the critical uncertainties but not to others, this should be stated. The same applies to policy C's robustness.
- Brief descriptive wording that communicates the key findings from the uncertainty assessment. The wording shown above is illustrative, not prescriptive.

As a whole, the table should be no larger than can fit on a single side of A4 paper, in landscape format. Including detail that makes it any larger than this is likely to make the main points less clear to the decision-maker. Detail can be included in the supporting information.