



Interfaces with Other Disciplines

Minimax decision rules for planning under uncertainty: Drawbacks and remedies

Edward Anderson^{a,*}, Stan Zachary^{b,c}^a Imperial College Business School, UK^b Heriot-Watt University, UK^c University of Edinburgh, UK

ARTICLE INFO

Article history:

Received 12 July 2022

Accepted 24 May 2023

Available online 28 May 2023

Keywords:

Decision analysis

Planning in energy systems

Minimax decision rules

Regret-based decision rules

ABSTRACT

It is common to use minimax rules to make planning decisions when there is great uncertainty about what may happen in the future. Using minimax rules avoids the need to determine probabilities for each future scenario, which is an attractive feature in many public sector settings. However there are potential problems in the application of a minimax approach. In this paper our aim is to give guidance for planners considering a minimax approach, including minimax regret which is one popular version of this. We give an analysis of the behaviour of minimax rules in the case with a finite set of possible future scenarios. Minimax rules will have sensitivity to the choice of a small number of scenarios. When regret-based rules are used there are also problems arising since the independence of irrelevant alternatives property fails, which can lead to opportunities to game the process. We analyse these phenomena considering cases where the decision variables are chosen from a convex set in R^n , as well as cases with a finite set of decision choices. We show that the drawbacks of minimax regret hold even when restrictions are placed on the problem setup, and we show how working with a structured set of scenarios can ameliorate the difficulty of having a final decision depend on the characteristics of just a handful of extreme scenarios.

© 2023 The Author(s). Published by Elsevier B.V.

This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>)

1. Introduction

We are interested in decision-making under uncertainty, and we are particularly concerned with the planning of infrastructure investment, e.g. in energy, where the value of such investments will depend on an unknown future. Many such projects will have lifetimes of decades and there are inevitably substantial uncertainties that will ultimately determine whether or not such an investment is worthwhile, and how it may compare with other possible choices of investment. Because of the long time horizons it is not easy to evaluate investments through a probabilistic analysis of how the future will evolve.

In a planning context it is usual to model future uncertainty by using a set of possible scenarios, each scenario describing what may happen in the future. The difficulty of assigning probabilities to different possible scenarios is central to our discussion. Infrastructure planning decisions often feature *deep uncertainty* (Marchau et al., 2019). These are cases where there is no agree-

ment on fundamental aspects of what may happen in the future. Often there are multiple interested parties with competing interests who cannot agree on the, necessarily subjective, probabilities for different scenarios. In a planning context the uncertainty may relate to decisions yet to be taken by governments, and will have a political flavour. In the energy area an example of this type of uncertainty in the UK relates to the future use of hydrogen as a fuel for domestic heating. One view is that a great deal of domestic heating will move to hydrogen, but another possibility is that there will be greater electrification through the use of heat pumps. What happens will depend on government actions and technology development as well as changing societal attitudes. Thus we see that if we need to plan the building of transmission lines and want to consider scenarios involving different proportions of domestic hydrogen heating being used in 2040 and beyond, then there are fundamental difficulties in assigning probabilities to these different scenarios. Carbon capture and storage represents another area where technology development and political choices will ultimately determine outcomes.

When planning takes place under uncertainty it is common for decisions to be made at different points in time, with a de-

* Corresponding author.

E-mail address: e.anderson@imperial.ac.uk (E. Anderson).

cision maker reacting to events as they occur. Moreover, the way in which decisions either restrict or open up future options will be an important consideration in practice. However, in this paper we will not consider this type of wait-and-see decision. Instead we will look at the simplest form of problem in which decisions need to be made now, and these will involve different costs and benefits depending on what happens in the future. Thus the model we use has a decision made at a single point in time. Moreover, we will not consider scenarios that depend in some way on the decisions, and thus in our model there is a well-understood finite set of possible future scenarios, that are not affected by the decisions made.

Thus we consider a formal setup which can be described using a matrix of costs associated with different decisions and scenarios. We recognise that infrastructure planning problems may be more complicated than this in practice. Nevertheless our simple problem setup is not unusual in infrastructure planning, and provides a useful starting point for any analysis even in cases where there is greater complexity. For example, our model can be extended to allow individual scenarios to evolve in different ways with given probabilities, with opportunities for further decision-making at subsequent points in time (depending on what happens). In this case the cost of a decision-scenario pair might be taken as the expected cost allowing for the further decisions to be made.

This paper considers minimax decision rules, which are frequently used in practice, partly because of risk aversion, and partly because these approaches do not need to specify probabilities on future possible scenarios. However, there are two important considerations that arise. First, while the use of minimax decision rules avoids the need to assign probabilities to scenarios, it does not address the practical issue of the choice of the scenario set itself, something which can make a big difference. Second, minimax decision rules are often used in conjunction with relative cost functions (usually defined as regret functions), rather than cost functions themselves. An argument for this comes from noting that particular scenarios may include large costs that are independent of the decision made. This would not matter when probabilities are assigned to scenarios, but does so when minimax is used. We have mentioned above a scenario that could be related to the uptake of hydrogen for domestic heating in the UK. A true minimax cost decision on the amount of electricity generation capacity built would be critically influenced by the costs associated with building hydrogen infrastructure which are unrelated to generation capacity. It is often the case that these scenario-based costs are potentially large and hard to estimate. For this reason it is attractive to consider rules which depend only on the *relative* costs in any given scenario. For example we may compare the cost if a particular project proceeds and the cost if this project does not go ahead. However, in cases where we are considering a choice between a number of alternatives there may be no obvious baseline associated with doing nothing: with regret-based rules we use the best choice in any scenario as a baseline. Nevertheless, it is well-known that regret-based rules fail to satisfy the economically desirable property of independence of irrelevant alternatives (discussed further below). This in turn leads to the possibility of gaming in which interested parties may seek to influence outcomes of the process inappropriately.

We are motivated by considering two broad types of planning decision: either the planner has to decide which of a variety of different projects should be carried out; or the planner has to determine the capacities of some infrastructure components. Thus our models use a decision vector x to be chosen from a set D , with either $D = \{0, 1\}^n$ or $D \subseteq \mathbb{R}^n$. In addition we will illustrate some observations with simple examples using a finite decision set D .

1.1. Minimax cost

Our focus is on minimax decision rules that are often used by planners, and we begin by discussing the minimax cost rule, before considering minimax regret. The minimax cost rule makes the choice that minimises the maximum cost over possible scenarios (or equivalently it maximises the minimum profit). As we make explicit in the next section, this can be considered as the limit of risk averse decisions as the degree of risk aversion increases. This is a conservative approach which is often attractive in a planning context. Rather than looking at the average performance the decision maker's attention is focused on the worst outcome. There is a substantial literature exploring the connection between minimax rules defined over a set of possible values for an unknown parameter, and optimisation carried out in a way that takes account of an appropriately defined risk measure (Anderson et al., 2020; Shapiro, 2012). Notice that the classic risk measures like VaR and CVaR are dependent on assigning probabilities to scenarios.

This approach also avoids the need to assign probabilities to different scenarios, and this is often quoted as a justification for the use of a minimax rule. In the public sector decisions may have significant financial consequences for individual firms and it is particularly important to make these in an objective way that follows a clear cut procedure – to avoid any suggestion of bias or favouritism. The minimax cost approach achieves this without the need to assign probabilities to different possible outcomes. But, as we show in this paper, the apparent objectivity is illusory, since the problem of assigning probabilities to scenarios is replaced by the problem of selecting scenarios to analyse. A particular problem with such non-probabilistic approaches is the considerable influence of more extreme scenarios (often by general consent considered relatively unlikely), where the decisions made may be very sensitive to whether or not these extreme scenarios are included in the scenario set.

Since we want to move away from a probabilistic interpretation of the scenarios, we end up looking for decisions which have good performance against each of the scenarios individually; and this is effectively a multi-objective understanding of the problem (Klamroth et al., 2017). An alternative viewpoint sees the minimax cost rule as arising from a robust choice of decision when there is a lack of information on the probabilities of different outcomes. The minimax cost rule is obtained when the planner minimises expected costs assuming the worst case assignment of the unknown probabilities for different scenarios, since this occurs when the worst scenario is given probability 1. A variation is to specify a subset, C , of the set of all possible probability distributions on scenarios, and then to look at the minimax cost over this reduced set. Gilboa & Schmeidler (1989) show that, when reframed in terms of utilities, this becomes a maximin expected utility rule that can be shown to arise axiomatically from some kinds of uncertainty aversion. We will argue that this approach, with appropriately chosen subset C is often preferable in practice to a full minimax cost approach.

1.2. Minimax regret

In this paper we will also consider minimax regret. Regret is an alternative form for the cost related to a particular scenario and decision pair: it is defined as the additional cost from the decision in comparison with the best possible decision for that scenario. The advantage of looking at regret rather than profits or costs is that our decisions are then determined by the scenarios in which the decision choice makes a significant difference. Suppose there is a single scenario in which some “disaster” occurs and all decisions do badly, with higher costs than in any other scenario. Then a minimax cost policy will choose the decision which is least bad for

this disaster scenario, even if the choice of policy makes almost no difference to the outcome in this scenario. It does not seem sensible to decide everything on the basis of a tiny improvement in a single scenario if this foregoes substantial improvement in most other scenarios, and this undesirable behaviour is avoided by using minimax regret.

Minimax regret is often used in a planning context. For example, it has been proposed for decisions on facility location, waste management, and flood defences [Chang & Davila \(2007\)](#); [Daskin et al. \(1997\)](#); [Van der Pol et al. \(2017\)](#). Planning decisions which relate to climate change and emissions have characteristics that may lead planners to use minimax regret as a decision tool (e.g. [Kunreuther et al., 2013](#); [Loulou & Kanudia, 1999](#); [Sanders et al., 2016](#)). Energy applications of minimax regret have been numerous. Later we will discuss the use of this approach by National Grid ESO in the UK. Other applications include decarbonisation in the EU ([Crespo del Granado et al., 2019](#)), unit commitment ([Jiang et al., 2013](#)) and transmission planning ([Chen et al., 2014](#); [Moreira et al., 2017](#)).

The use of regret based measures involves assessment of the relative results for different decisions. There has been much discussion of the extent to which this may be a good model of the way that decisions are made in practice. Regret theory has been argued as being a natural way to explain some of the paradoxes that occur in actual decision making ([Bell, 1982](#)) where decision makers seem to have consistent preferences that cannot be easily explained under a fully rational model. The key point is that decision makers acting in an uncertain world not only consider the actual outcomes that can occur, but also what might have been possible. So the discovery that, as things have turned out, another alternative would have been preferable imparts a sense of loss or regret. Decision makers may be prepared to trade-off some amount of (expected) financial return to avoid this regret. There is debate about the extent to which these conclusions are supported by empirical evidence, and prospect theory as developed by Kahneman and Tversky is now widely used as a way of explaining how people choose between options when there is uncertainty on outcomes. But regret theory remains important and a summary of some more recent research in this area is given by [Bleichrodt & Wakker \(2015\)](#). A minimax regret policy may match the way decisions are often made in practice when there is a shortage of information and analysis available. Thus it is arguable that minimax regret aligns with our usual mental heuristics, even if these are not entirely rational.

However since Savage's original paper on minimax regret ([Savage, 1951](#)) it has been recognised that a weakness of this approach is that it fails to satisfy a property of independence of irrelevant alternatives (IIA). This property implies that adding a single action (or changing the results of an action) cannot have an impact on the choice made by the decision maker unless the new or changed action could be chosen by the decision maker. It is also the case that a minimax regret approach may fail to satisfy transitivity, so that we can have a cycle of possible decisions where each is preferred to the next if just these two are compared on the basis of minimax regret.

The failure of the minimax regret rule to satisfy these basic properties that we might expect in a good decision rule reflects more general results showing that decision makers who follow a set of reasonable axioms must act as though they are maximising expected utility. More precisely we can find reasonable assumptions under which a rational decision maker will act as if there is a utility function u and a set of subjective probabilities p_i , $i \in S$, such that the decision maker chooses the decision x that maximises their expected utility $\sum_{i \in S} p_i u(C_i(x))$ where $C_i(x)$ is the cost occurring under scenario i and decision x . The first person to provide a set of axioms under which this holds was Savage ([Savage, 1954](#)). But there is a rich literature in this area: see for example

[Anscombe & Aumann \(1963\)](#) and [Gul \(1992\)](#). These papers require a variety of different assumptions some of which are less likely to apply in the case that is our main focus, which is the planning of infrastructure investment. For examples Savage's original theorem requires the scenario set to be infinite. Many of these papers (though not the original work of Savage, or the paper by Gul) require it to be possible to define lotteries on actions (e.g. take action x half the time and action y half the time), so that the decision maker will have a preference between any pair of lotteries on actions. We may simplify the conclusions from this body of research by saying that any rule, such as minimax regret, that does not imply a probability measure on scenarios will lead to circumstances in which it fails to meet one or other axiom that one might regard as appropriate for rational decisions.

A number of authors have considered alternative sets of axioms for decisions that will imply a regret-based rule, or the minimax regret criterion. This work started with [Milnor \(1951\)](#) and includes [Hayashi \(2008\)](#) and [Stoye \(2011a\)](#). The choice of axioms vary between authors, but [Hayashi \(2008\)](#) and [Diecidue & Somasundaram \(2017\)](#) discuss sets of axioms in which only regret-based decision rules are possible. Stoye gives a set of eight axioms, which if all satisfied, imply that the decision rule must be minimax regret.

The nature of a regret-based rule implies that the choice made depends on the choice set available, even of options that are not chosen. This means that we can no longer insist on a transitivity condition, and the independence of irrelevant alternatives condition needs to be modified. Stoye's axioms that imply a minimax regret rule include two axioms related to the IIA and in both cases these are easily seen to hold for minimax regret: (axiom 4) when all choices have the same outcome in all scenarios, then introducing a new choice cannot change the decision, unless this new choice is selected; and (axiom 5) introducing a new choice with the property that in each scenario there is an existing choice that is at least as good cannot change the decision, unless this new choice is selected.

The paper [Stoye \(2011a\)](#) is complemented by another paper [Stoye \(2011b\)](#) which deals with a setup in which menu-dependent preference orders are used rather than choice sets. The results are similar and a set of axioms are derived in which minimax regret is the only possible preference order. In both papers a key assumption required for minimax regret is (symmetry) that exchanging the consequences of two events leaves the choice (or preference order) unaltered.

From the perspective of a planner the results from these axiom-based studies are revealing since they point to the areas in which problems may occur. But they are restrictive in typically assuming that some form of mixing between decisions is possible. In a planning context this is inappropriate, but instead we may consider hypothetical choices which are constructed through changes in the cost associated with different outcomes (as we do in the proof of [Proposition 4](#)).

1.3. Contribution

In this paper we will explore the implications of using minimax rules, and particularly the use of minimax regret. We highlight two key limitations that recur in different contexts. First for both minimax cost and minimax regret the number of scenarios required to determine the solution is no more than $n + 1$ in the case that there is a continuous decision problem in an n -dimensional space. This shows that the decision is critically dependent on the exact set of scenarios included in the analysis. There is a need to be very careful in the choice of scenarios, with very implausible scenarios excluded. When scenarios emerge from a consultation processes that involves many parties with different goals, as happens for example in a climate change context, then there is the possibility that a fi-

nal decision could be greatly changed as a result of an interested party pressing for a particular scenario to be included.

A second important limitation is that the minimax regret rule will often fail to satisfy the independence of irrelevant alternatives property (IIA). Here the IIA property means that the decision chosen from a set of possible decisions cannot be affected by the presence or absence of an option which is never chosen. We should note that the same terminology can be used to refer to a similar property arising in social choice theory (Ray, 1973). When a planner is choosing between projects from a number of firms and the IIA property does not hold, this makes the decision open to being gamed, for example by a firm offering a project that will not be selected but will increase the chance of some other project being chosen. We already commented on a difficulty with minimax decision rules where objectivity is compromised by the need to select exactly the scenarios to include. This issue with regret-based rules is a kind of mirror image of this with objectivity compromised by the need to select exactly the decision options that are included.

We make both a theoretical and a practical contribution in this paper. On the theoretical side we prove that if the decision variables lie within a convex set in $\mathbb{R}^n$, and a minimax rule is used, then there are only $n+1$ scenarios that determine the optimal solution. We also prove that the failure of the IIA property for regret-based rules is more complete than might be imagined. We show that any regret-based rule that fully satisfies an IIA property must be equivalent to simply minimising an expected cost for some choice of scenario probabilities. Moreover we show that even a weakened form of the IIA property fails to hold in the case where the planner chooses which of a set of projects to start. Our exploration of regret-based approaches illustrates their poor behaviour through a number of small examples.

In addition to contributing theoretically, we make some proposals that we believe are helpful in practice. The approach we recommend can be applied in cases where it is possible to identify a number of “core” scenarios, which reflect major sources of uncertainty – e.g. as to the direction of future government policy – about which it is difficult, or not appropriate, to assign probabilities. In addition, associated with each of these core scenarios may be some number of more “extreme” scenarios, reflecting deviations from the core scenario – for example, due to the severity of weather conditions – to each of which it may be reasonable either to assign a probability, conditional on the occurrence of the core scenario, or to place bounds on this conditional probability. The differentiation between core and extreme is not necessary for our approach. The key is that there are groups of linked scenarios where we can bound the relative probabilities within each group. We show how this type of structure leads to minimax problems that can be easily solved and allow many more scenarios to be involved in the determination of an optimal policy. This is exactly the structure that we find in the problem of setting overall procurement amounts in the GB electricity capacity market. We work through a case study using data from National Grid ESO to show how our recommended approach can work in practice.

In the next section we set up the decision framework and discuss the behaviour of minimax rules. We show how a minimax rule can be extended through redefining the uncertainty set in a robust interpretation of the minimax decision, and we also introduce the problem of choosing which of a set of possible projects should be undertaken. In Section 3 we consider minimax problems where the decision variables are chosen from some convex set in $\mathbb{R}^n$, and show how the total set of scenarios can be reduced to a smaller set determining the solution. In Section 4 we consider the properties of decision rules based on regret, including minimax regret. In Section 5 we look in more detail at a case where we balance capacity investment and risk. This gives an opportunity to use structural properties of the scenarios and we show how this works

out in a case study from capacity procurement for electricity generation.

2. Minimax rules

We write S for the set of possible scenarios, which we will assume is finite. We will assume a decision vector $x \in D$ and for each $i \in S$ there is a cost function $C_i(x)$ giving the costs born by the decision maker if x is chosen and then scenario i occurs. Each cost function may represent a monetary cost, or some other measure. In the case of a single scenario i , the optimal decision x would be that which minimised $C_i(x)$. But with multiple scenarios, the optimal decisions will in general differ between scenarios and so we aim to minimise some composite objective function $\tilde{f}(x)$ defined on D .

There are many ways to define such a composite function. The classic expected utility framework uses $\tilde{f}(x) = \sum_{i \in S} p_i f_i(x)$ where the p_i are non-negative and each $f_i(x)$ is an (increasing) function of the costs $C_i(x)$. We can interpret $f_i(x)$ as the negative utility associated with costs $C_i(x)$, i.e. $f_i(x) = -u(C_i(x))$ and the p_i (after normalising to sum to one) as probabilities associated with each scenario. Then minimising $\tilde{f}(x)$ is simply maximising the expected utility for a utility function u . When u is linear this becomes simply minimising expected cost, and the risk averse case corresponds to concave u . Savage's Theorem implies that this form of function $\tilde{f}(x)$ is the only possibility if we wish to satisfy some reasonable rationality assumptions. However there is still a difficulty in assigning appropriate probabilities to the scenarios.

A second possibility is to define

$$\tilde{f}(x) = \max_{i \in S} C_i(x)$$

which leads to the minimax cost rule of choosing x to solve $\min_{x \in D} \max_{i \in S} C_i(x)$. This is equivalent to the expected utility formulation above in the case that the utility function u is sufficiently concave in costs, corresponding to extreme risk aversion. For example if the utility function is $u(z) = 1 - \exp(-kz)$ for payoff z , and we let $k \rightarrow \infty$ then for any fixed decision x , we have $f_i(x) = -1 + \exp(kC_i(x))$ and the scenario i^* with the largest cost $C_{i^*}(x)$ becomes dominant in the sum $\sum_{i \in S} p_i f_i(x)$ in any case where $p_{i^*} > 0$, no matter how small this probability is. Thus the decision taken will, in the limit, be the same as the minimax cost rule.

Later we will discuss minimax regret, but the properties we discuss for minimax cost rules will generally carry over to the regret case simply by substituting the regret function for costs. We will often write $f_i(x)$, $i \in S$, $x \in D$, for the general case where f_i is either a cost or regret function.

As an example of how a minimax rule operates, consider the following case with three scenarios and three possible decision choices:

Example 1. Suppose there are 3 possible decisions, and D is the set $\{u, v, w\}$. There are 3 scenarios and costs are as follows:

	u	v	w
Scenario A	4	0	5
Scenario B	3	5	0
Scenario C	3	2	0

It is easy to see that the minimax cost policy chooses the decision u which has a maximum cost of 4, lower than the maximum cost of 5 that can occur with each of the other options.

Shapiro & Kleywegt (2002) consider minimax decision rules and give conditions under which a solution to the minimax problem is also a solution to the problem of minimising the expected costs for an appropriately chosen set of probability weights. Their result applies in the case that D is a convex subset of $\mathbb{R}^n$ and the functions f_i are convex with S finite. Example 1 shows that their result

will not hold when there is a finite set of possible decisions and no specific structure for f_i . In this example there is no probability distribution over scenarios that generates the same choice as a minimax rule.

To establish this suppose that we have probabilities p_A , p_B and p_C for the three scenarios. Using $p_A = 1 - p_B - p_C$ the inequalities needed for u to be selected under a minimum expected cost policy are

$$4(1 - p_B - p_C) + 3p_B + 3p_C \leq 5p_B + 2p_C,$$

$$4(1 - p_B - p_C) + 3p_B + 3p_C \leq 5 - 5p_B - 5p_C.$$

These inequalities can be simplified to

$$6p_B + 3p_C \geq 4,$$

$$4p_B + 4p_C \leq 1.$$

These are incompatible with non-negative probabilities, since the second inequality implies that both p_B and p_C are less than $1/4$ which implies $6p_B + 3p_C \leq 9/4$ contradicting the first inequality. So there is no assignment of probabilities to scenarios under which the minimax regret policy of x will be chosen.

2.1. Minimax as robust optimisation

We will also generalise our discussion of minimax decision rules by considering them as robust versions of expected cost minimisers when there is a range of possible probabilities for the scenarios. We suppose that the set of allowed probability measures is restricted to a convex subset of the set of all probability measures, which we write as $\mathbb{P}_A$, and corresponds to the set C occurring in Gilboa & Schmeidler (1989). This approach is useful when a planner is able to define some core scenarios around which a wider scenario set is constructed. The core scenarios then represent a starting point around which variations are considered. If there are clear policy goals to be achieved, then the core scenarios may represent different ways in which this can be brought about. This occurs for National Grid in their use of Future Energy Scenarios which reflect different ways in which a net zero GHG emissions target may be achieved in the UK by 2050 (see National Grid, 2020). Hughes & Strachan (2010) use the term “backcasting” to refer to this way of generating scenarios. The report by Dent et al. (2020) commissioned by OFGEM in the UK also recommends an approach based around core scenarios. In the broader climate change context, the scenarios developed and used by the IPCC are the five Shared Socio-economic Pathways (SSPs), and these could be regarded as the core scenarios. However there has been criticism that these pathways are insufficiently diverse (McCollum et al., 2020). It is generally accepted that the set of scenarios used, however they are constructed, should cover the range of possible outcomes.

The introduction of a subset $\mathbb{P}_A$ allows the definition of a partial ordering amongst the probabilities of different scenarios, $p_i \geq p_j$, etc. When there are core scenarios we can then define extreme scenarios, where each extreme scenario is guaranteed less likely than the corresponding core scenario. Or we may be more explicit and define a set of scenarios around each core scenario with probabilities defined conditional on one of these scenarios occurring (the core or one of its associated extremes) – as in the case study of Section 5. Within this framework we can also fix lower bounds on the probabilities of certain scenarios. This idea, where we use a minimax setting taking account of information that restricts the possible distributions, is similar to that of Γ -minimax rules in the literature on robust Bayesian methods (Insua et al., 1995).

Given the set $\mathbb{P}_A$ of possible probability distributions over scenarios, we can formulate the robust optimisation problem:

$$\text{RO}(\mathbb{P}_A) : \min_{x \in D} \max_{P \in \mathbb{P}_A} \mathbb{E}_P(f(x)) = \min_{x \in D} \max_{(p_i: i \in S) \in \mathbb{P}_A} \sum_{i \in S} p_i f_i(x). \quad (1)$$

In the case that $\mathbb{P}_A = \{p_i, i \in S \mid p_i \geq 0, \sum p_i = 1\}$ then $\text{RO}(\mathbb{P}_A)$ becomes the standard minimax problem, since

$$\max_{(p_i: i \in S) \in \mathbb{P}_A} \sum_{i \in S} p_i f_i(x) = \max_{i \in S} f_i(x),$$

as can be seen by observing that the inequality in one direction follows from taking $p_i = 1$ for different i , and the inequality in the other direction follows since a convex combination cannot be greater than the highest value of its constituents.

An important case occurs when $\mathbb{P}_A$ is defined by a finite set of linear constraints satisfied by the probabilities of each scenario. Thus we have a matrix A of constraints with the probabilities p_i of each scenario in the set $\mathbb{P}_A = \{p_i, i \in S \mid Ap \leq 0, p_i \geq 0, \sum p_i = 1\}$. This covers the case with no constraints when A is empty, and the case with given probabilities through letting pairs of constraints such as $p_1 - 2p_2 \leq 0$, $2p_2 - p_1 \leq 0$, define the ratios between probabilities and then using the fact that $\sum p_i = 1$.

We note that the robust problem can be solved easily through dualising the inner maximisation:

$$\begin{aligned} &\text{maximise} \sum_{i \in S} p_i f_i(x) \\ &\text{subject to} \quad Ap \leq 0, \\ &\quad \sum_{i \in S} p_i = 1, p_i \geq 0. \end{aligned}$$

Then the original minimax problem can be written as the following optimisation problem: Choose $x \in D$, $w \in \mathbb{R}$ and $q \in \mathbb{R}^m$ to

$$\begin{aligned} &\text{minimise}_{x, w, q} \quad w \\ &\text{subject to} \quad w + (A^T q)_i \geq f_i(x), i \in S, \\ &\quad q \geq 0. \end{aligned}$$

There are three types of decision problem that we will investigate, depending on the set of possible decisions, D . First D may be a finite set without any particular structure. Second decisions may be required, simultaneously, on whether or not to go ahead with a number of different projects. This is a selection problem and we have $D = \{0, 1\}^n$. Third D may be a convex subset of $\mathbb{R}^n$, for example, a planner may wish to determine the right capacities for some set of transmission links to be constructed.

2.2. Deciding amongst a set of projects

An important area of application occurs when there are multiple projects that are available and the decision maker has to determine which projects to proceed with. For infrastructure investment this may involve complex interactions between projects and a combination of different criteria (Mild et al., 2015). In many cases the selection of a project portfolio is made by the decision maker on the basis of proposals put forward by other parties, giving rise to other issues when the aims of the multiple interested parties do not fully align (Armstrong & Vickers, 2010). Here we consider a simple model using a finite set of scenarios to capture the uncertainty in project value. There is a set U of n different possible projects, so we can write the decision space as $D = \{0, 1\}^n$. In scenario $i \in S$ the selection $T \subset U$ gives a cost to the decision maker of $C_i(T)$.

The minimax decision rule can be used to determine which of a number of alternative projects should be carried out. In practice projects are considered on the basis of net present values and a maximin decision rule can be applied. For consistency with the rest of this paper we will consider the negative of the NPV for a project as its cost, and then consider a minimax rule based on costs.

A simple case of this sort occurs with *independent additive costs*. Under this assumption the cost implications of different projects

are additive where the individual project costs in scenario i are $c_i(1), c_i(2), \dots, c_i(n)$. These are the extra costs associated with each project in scenario i (and are negative if that project is profitable in scenario i). We also have a total cost in scenario i of W_i if none of the projects are selected. Hence

$$C_i(T) = \sum_{k \in T} c_i(k) + W_i. \quad (2)$$

Thus each scenario has a different additive cost function over D . Note that, as we have already mentioned, the scenario-based costs W_i will have an impact on the decision made by the minimax cost rule. This is appropriate if the decision maker is very risk averse. For example suppose that the projects correspond to different flood protection measures, with scenarios representing different future possible flood patterns. Then W_i is the flood costs associated with scenario i if no flood protection projects go ahead. A highly risk averse decision maker will consider total flood costs in scenario i , even though the W_i contribution is unaltered, no matter which projects are selected.

The use of a minimax rule does not sit easily with the idea that separate projects may be considered independently of each other, as is shown by the following very simple example with independent additive costs.

Example 2. Suppose there are just two scenarios A and B and we consider whether or not to go ahead with two projects X and Y . The costs involved under different scenarios, if the project X goes ahead or not, are given in the following table:

	Do not proceed with X	Proceed with X
Scenario A	0	3
Scenario B	4	0

The numbers related to project Y are the same. Costs are additive, so we have the following costs for the 4 options.

	$\emptyset$	$\{X\}$	$\{Y\}$	$\{X, Y\}$
Scenario A	0	3	3	6
Scenario B	8	4	4	0

To put this into the framework of (2) we set $W_A = 0$, $W_B = 8$, $c_A(X) = c_A(Y) = 3$ and $c_B(X) = c_B(Y) = -4$. Then the minimax cost policy selects one of the projects X or Y for completion, but does not do both.

3. Reduced scenarios for decisions in $\mathbb{R}^n$

In the case where the decision set $D \subseteq \mathbb{R}^n$ is convex, we can establish that a small subset of the total set of scenarios will determine the minimax cost policy (or the minimax regret policy). There are several ways to formulate this property. One possibility is to ask which scenarios are essential in the sense that if they were not present then the decision maker using a minimax policy would make a different choice. When dealing with constraints Calafiore & Campi (2005) call these *support constraints*. In our formulation we can say that a scenario $s \in S$ is a *support scenario* if $\min_x \max_{i \in S \setminus \{s\}} f_i(x) < \min_x \max_{i \in S} f_i(x)$. Theorem 2 of Calafiore and Campi then constrains the maximum number of support scenarios. However, in general, keeping just the support scenarios and throwing away the others may still lead to a change in the optimal solution: it is easy to construct an example, with two similar scenarios where the optimal solution to the minimax problem remains unchanged if either of the two scenarios is dropped (so that neither is a support scenario), but where the optimal solution changes if both the scenarios are dropped.

This motivates us to make a definition in terms of subsets of the scenarios, rather than individual scenarios. For $K \subseteq S$, we write $P(K)$ for the minimax problem for the set of scenarios K , and cost

or regret functions $f_i(x)$, so

$$P(K) = \min_x \max_{i \in K} f_i(x),$$

and we write $x^*(K)$ for the set of optimal solutions to $P(K)$. We say that $K \subseteq S$ is a *support set* for $P(S)$ if $P(K)$ has the same value as $P(S)$ and $x^*(K) = x^*(S)$. Thus dropping all the other scenarios does not affect the optimal solution. In the next result we establish the existence of a support set of size $n + 1$. This result will hold provided that the functions f_i are quasiconvex, so that the level sets $A_i(z) = \{x \in D : f_i(x) \leq z\}$ are convex for each $z \in \mathbb{R}$. This is a different result from that of Calafiore and Campi.

Proposition 1. Suppose that D is a convex set in $\mathbb{R}^n$; for each $K \subseteq S$, the solution $x^*(K)$ to the minimax problem is uniquely defined; and for each $i \in S$, f_i is quasiconvex. Then there is a support set for $P(S)$ containing no more than $n + 1$ scenarios.

Proof. Suppose that $R^* = \min_{x \in D} \max_{i \in S} f_i(x)$. Let $x_0 = x^*(S)$ be the point achieving this minimum and set $S' = \{i \in S : f_i(x_0) = R^*\}$. From the definition and uniqueness of x_0 and by quasiconvexity, the level sets $A_i(R^*)$, $i \in S'$, are convex and $\bigcap_{i \in S'} A_i(R^*) = \{x_0\}$. Hence, for every choice of (direction) u , there is some $i_u \in S'$ such that $x_0 + \delta u \notin A_{i_u}(R^*)$ for all $\delta > 0$, and further there is some hyperplane through x_0 which separates $A_{i_u}(R^*)$ and the set of points of the form $x_0 + \delta u$ for $\delta > 0$. This hyperplane may be defined by its normal g_u , say. Then $g_u^\top u \geq 0$ and $g_u^\top (x - x_0) < 0$ for every $x \in A_{i_u}(R^*)$. Write $G = \{g_u : u \in \mathbb{R}^n\}$ and observe that $g_u^\top u \geq 0$ implies 0 is in the convex hull of G . By Caratheodory's theorem we can choose a set of at most $n + 1$ such hyperplanes $g_{u_1}, \dots, g_{u_k}$, $k \leq n + 1$, with 0 in the convex hull of $\{g_{u_1}, \dots, g_{u_k}\}$. This implies that, for any choice of direction u , we have $g_{u_j}^\top u \geq 0$ for at least one $j = 1, \dots, k$. Let $K = \{i_{u_1}, \dots, i_{u_k}\} \subseteq S'$ be the corresponding set of associated scenarios. Suppose that $x^*(K) \neq x_0$ and thus

$$\max_{i \in K} f_i(x^*(K)) < R^*.$$

Since $f_i(x_0) = R^*$ for all $i \in K$, $x^*(K)$ must be in the interior of the level set $A_i(R^*)$ for all $i \in K$ and thus $g_{u_j}^\top (x^*(K) - x_0) < 0$ for all $j = 1, \dots, k$ which is a contradiction, thus establishing $x^*(K) = x_0$ as required. $\square$

We can extend this result to the robust optimisation problem $RO(\mathbb{P}_A)$ in the important case that $\mathbb{P}_A$ is determined by a matrix A which is block diagonal. Thus the scenario set S may be partitioned into m disjoint components $S = S_1 \cup S_2 \cup \dots \cup S_m$, and the matrix A is such that each constraint (corresponding to some row of A) involves only scenarios within a single component of S .

We show below a version of Proposition 1 that establishes that the robust optimisation solution is determined by a set of scenarios in at most $n + 1$ of the components S_k .

Proposition 2. Suppose that D is a convex set in $\mathbb{R}^n$; f_i is strictly convex for each $i \in S$; and A is block diagonal inducing a split of S into m disconnected components $S_1, S_2, \dots, S_m$. Then there is a support set for $RO(\mathbb{P}_A)$ with scenarios in at most $n + 1$ of the components of S .

Proof. Suppose that R^* is the value of the problem $RO(\mathbb{P}_A)$, so

$$R^* = \min_{x \in D} \max_{(p_i : i \in S) \in \mathbb{P}_A} \sum_{i \in S} p_i f_i(x).$$

Let $x_0 = x^*(S)$ be the point achieving this minimum. Thus R^* is the optimal value for the inner maximisation $LP(x_0)$, which can be written as the following linear program

$$\begin{aligned} LP(x) : \quad & \text{maximise} && \sum_{i \in S} p_i f_i(x) \\ & \text{subject to} && Ap \leq 0, \\ & && \sum_{i \in S} p_i = 1, p_i \geq 0. \end{aligned}$$

We break this up into the individual components. Let

$$\begin{aligned} \text{LP}_j(x) : \quad & \text{maximise} \quad \sum_{i \in S_j} p_i f_i(x) \\ & \text{subject to} \quad A_j p \leq 0, \\ & \quad \quad \quad \sum_{i \in S_j} p_i = 1, p_i \geq 0, \end{aligned}$$

where A_j is the appropriate component of A involved in component S_j . We write $g_j(x)$ for the optimal value of $\text{LP}_j(x)$. So if $p_i^{(j)}(x)$ is an optimal solution to $\text{LP}_j(x)$ then

$$g_j(x) = \sum_{i \in S_j} p_i^{(j)}(x) f_i(x).$$

We note that g_j is strictly convex, since if we take $x_C = \lambda x_A + (1 - \lambda)x_B$ for $\lambda \in (0, 1)$ then

$$\begin{aligned} g_j(x_C) &= \sum_{i \in S_j} p_i^{(j)}(x_C) f_i(x_C) \\ &< \lambda \sum_{i \in S_j} p_i^{(j)}(x_C) f_i(x_A) + (1 - \lambda) \sum_{i \in S_j} p_i^{(j)}(x_C) f_i(x_B) \\ &\leq \lambda g_j(x_A) + (1 - \lambda) g_j(x_B) \end{aligned}$$

where the final inequality arises because $p_i^{(j)}(x_C)$ are feasible for $\text{LP}_j(x_A)$ and $\text{LP}_j(x_B)$.

Then consider the following problem with decision variables v_j , $j = 1, 2, \dots, m$:

$$\begin{aligned} \text{LP1}(x) : \quad & \text{maximise} \quad \sum_{j=1}^m v_j g_j(x) \\ & \text{subject to} \quad \sum_{j=1}^m v_j = 1 \\ & \quad \quad \quad v_j \geq 0, j = 1, 2, \dots, m. \end{aligned}$$

From a feasible solution to $\text{LP1}(x)$ we can recover a feasible solution to $\text{LP}(x)$ by taking $p_i = v_j p_i^{(j)}$ for $i \in S_j$. This is feasible since the probability ordering $Ap \leq 0$ is retained and $\sum p_i = \sum_{j=1}^m v_j (\sum_{i \in S_j} p_i^{(j)}) = 1$. Thus $\text{LP1}(x)$ is equivalent to $\text{LP}(x)$.

Now observe that $\text{LP1}(x)$ has a solution which is simply $\max\{g_j(x) : j = 1, 2, \dots, m\}$ and, since strict convexity implies a unique solution to the minimax problem, we have converted the problem into one for which [Proposition 1](#) applies (with $g_j(x)$ being quasiconvex, since it is convex). Thus there is a subset $L \subset \{1, 2, \dots, m\}$ with at most $n + 1$ elements and

$$R^* = \min_x \max_{j \in L} g_j(x).$$

Since $g_j(x) = \sum_{i \in S_j} p_i^{(j)}(x) f_i(x)$ where $p_i^{(j)}(x)$ is an optimal solution to $\text{LP}_j(x)$, this establishes the desired result. $\square$

In the case that there are a limited number of core scenarios each associated with a set of extreme scenarios, this result implies that only $n + 1$ of the core scenarios will be used in determining the minimax decision. But the restrictions on probabilities within each component of S will imply a relatively large number of individual scenarios contributing to the final result. We believe this can represent an attractive compromise in practice.

We note that the property of there being a reduced set of $n + 1$ scenarios, that holds when D is convex, does not apply in the discrete project setting, even with independent additive costs. In the appendix we demonstrate this with an example where there are three projects, but 5 scenarios are all involved in the determination of the optimal project choice.

4. Regret-based decision rules

In this section we turn to regret-based decision making. As we discussed in the Introduction it is arguably desirable that decision making should be based on relative costs under each scenario. We consider the extent to which minimax regret is open to being gamed. While it is well-known minimax regret does not satisfy the IIA property, we go further and show that regret-based decision making only avoids this problem if it essentially probabilistic.

As mentioned in the Introduction using a simple minimax cost may give rise to decisions that seem inappropriate. The result of a minimax decision rule can change if for a particular scenario every decision has some cost added. In this sense the decision can depend on a cost associated with a scenario that is not impacted by the decision made. This is a particular difficulty in a planning context where there are often costs that do not depend on the decision made, but do depend on the scenario that eventuates. Moreover if there is a single scenario which has a high, but almost constant, cost across decisions then the outcomes in this single scenario will determine the choice made.

For each scenario $i \in S$, we define $R_i(x)$, the regret function defined on D , as

$$R_i(x) = C_i(x) - \inf_{z \in D} C_i(z). \quad (3)$$

Thus $R_i(x)$ represents the ‘regret’ felt by the decision maker if she chooses x and scenario i occurs. Since $\inf_{x \in D} R_i(x) = 0$, $i \in S$, if in retrospect the best choice of x was made for the scenario that occurs then there is no regret. Assuming that scenario i occurs with probability p_i , the optimisation problem facing the decision maker who wants to minimise expected regret is to find $\min_x \sum_{i \in S} p_i R_i(x)$, and this is equivalent to finding $\min_x \sum_{i \in S} p_i C_i(x)$ which is the expected cost. The minimax regret rule sets

$$\bar{f}(x) = \max_{i \in S} R_i(x)$$

so that the decision maker solves $\min_{x \in D} \max_{i \in S} R_i(x)$. This will produce a different result to the minimax cost rule.

We mentioned above the possibility of the minimax cost policy doing badly on some problem instances where one scenario gives high costs for all decisions, but with only a small difference between them. To support the idea that minimax regret may do better on average we compare the behaviour of the two decision rules on a random problem instance. We consider the simplest possible case with just two scenarios and three possible decisions x, y, z . Thus a problem instance is determined by the six cost values $C_i(x)$, $C_i(y)$ and $C_i(z)$, $i = 1, 2$. In a symmetric version of this problem we can suppose that each scenario is equally likely. Thus the expected cost for a given problem instance is the average cost over the two scenarios for the decision chosen, which can then be compared between the two decision rules. When the six cost values are all drawn randomly and independently from $[0, 1]$ then the minimax cost rule gives an expected cost of $12/35 = 0.3429$. The minimax regret rule does better – though it is hard to give an exact analysis, we can show by simulation that the expected cost is 0.3286.

As we have mentioned the independence of irrelevant alternatives property fails to hold for the minimax regret problem. Thus removing an option which is not chosen can have an impact on the optimal choice. In fact, different choices for the characteristics of a new option that is added can lead to many different existing choices becoming optimal under the minimax regret rule. As an example of this we show that if a decision has the minimum cost under any of the scenarios, then it can become the minimax regret choice if a new option is added with correctly chosen costs.

Proposition 3. Suppose that there is some $y \in D$ and $k \in S$, where $C_k(y) < C_k(x)$ for $x \neq y$. Then we may define a new possible decision

z and a set of values $C_i(z)$, $i \in S$ such that y is chosen for the minimax regret problem $\min_{x \in D \cup \{z\}} \max_{i \in S} R_i(x)$.

Proof. We write $M = \max_{j \in S} \max_{x \in D} C_j(x)$ for the global maximum over the existing sets D and S , and $L = \max_{j \in S} \min_{x \in D} C_j(x)$ for the largest of the minimums over D (which may be negative). Consider $z \notin D$ and set $C_i(z) = M$ for $i \in S \setminus \{k\}$ and $C_k(z) = C_k(y) - M + L$. Then, since $M > L$, the decision z replaces y as the preferred choice in scenario k with $R_k(y) = C_k(y) - C_k(z) = M - L$ and $R_k(x) = C_k(x) - C_k(z) > C_k(y) - C_k(z) = M - L$ for $x \neq y, z$. Thus $\max_{i \in S} R_i(x) > M - L$ for $x \neq y, z$. For $i \in S \setminus \{k\}$ we have $\min_x C_i(x) \leq L$, and so we can deduce that $R_i(z) = M - \min_x C_i(x) > M - L$ for $i \in S \setminus \{k\}$, so $\max_{i \in S} R_i(z) > M - L$. Also $R_j(y) = C_j(y) - \min_x C_j(x) < M - L$ for $j \neq k$, so $\max_{i \in S} R_i(y) = M - L$. This establishes the result. $\square$

The failure of the IIA property for minimax regret is not due to the minimax operation. In fact the use of any decision function based on regrets will fail the IIA property, unless the decision problem reduces to simply minimising expected cost under some choice of probabilities over scenarios. Here we take the IIA property to mean that if decision x is preferred to decision x' , then this preference order is not changed by any change in the costs associated with a third decision x'' .

Proposition 4. Suppose that, for a given decision problem, the decision space D consists of at least three points, and the decision is made by minimising a continuous non-decreasing function of the set of regrets $f^*(R_i(x) : i \in S)$ over $x \in D$ and in addition this decision rule satisfies the IIA property, then the decision process is equivalent to minimising the expected costs for some set of probabilities p_i , $i \in S$.

Proof. It will be convenient to assume a total of s different scenarios and regard the set of regrets as a vector in $\mathbb{R}^s$, so $f^* : \mathbb{R}^s \rightarrow \mathbb{R}$. Moreover, since regret functions are non-negative, the function f^* is naturally defined on $\mathbb{R}_+^s$. The IIA condition implies that, for any given scenario (assume without loss of generality it is scenario 1) and for any two possible decisions x and x' in D , we have that

$$\begin{aligned} f^*(R_1(x), \dots, R_s(x)) &> f^*(R_1(x'), \dots, R_s(x')) \\ \Rightarrow f^*(R_1(x) + a, \dots, R_s(x) + a_s) &> f^*(R_1(x') + a, \dots, R_s(x') + a_s) \end{aligned} \quad (4)$$

for all $a \geq -\min(R_1(x), R_1(x'))$. The reason for this is that the values $R_1(x)$ and $R_1(x')$ of the regret R_1 may be changed by the same amount a by suitable adjustment of the cost $C_1(x'')$ associated with one or more of the remaining decisions x'' ; the IIA property implies that this should not change the direction of the inequality on the left hand side of (4).

Iteration of (4) implies that

$$\begin{aligned} f^*(R_1(x), \dots, R_s(x)) &> f^*(R_1(x'), \dots, R_s(x')) \\ \Rightarrow f^*(R_1(x) + a_1, \dots, R_s(x) + a_s) &> f^*(R_1(x') + a_1, \dots, R_s(x') + a_s), \end{aligned} \quad (5)$$

provided that all of $R_1(x) + a_1, \dots, R_s(x) + a_s$ and $R_1(x') + a_1, \dots, R_s(x') + a_s$ remain non-negative. Moreover we can easily establish the same implication with inequalities reversed. And the combination of both implications shows that

$$\begin{aligned} f^*(R_1(x), \dots, R_s(x)) &= f^*(R_1(x'), \dots, R_s(x')) \\ \Rightarrow f^*(R_1(x) + a_1, \dots, R_s(x) + a_s) &= f^*(R_1(x') + a_1, \dots, R_s(x') + a_s). \end{aligned} \quad (6)$$

Now suppose that $f^*(\tilde{R}) = f^*(\tilde{R} + \tilde{a})$ for arbitrary vectors $\tilde{R}$ and $\tilde{a}$, with $\tilde{R} \geq 0$ and $\tilde{R} + \tilde{a} \geq 0$. We will show that $f^*(\tilde{R}) = f^*(\tilde{R} + q\tilde{a})$ for any $q \in \mathbb{R}$. Since f^* is continuous the set of points having the same value is closed, and hence it will be enough to show that $f^*(\tilde{R}) = f^*(\tilde{R} + q\tilde{a})$ for any rational number q with $\tilde{R} + q\tilde{a} \geq 0$.

Consider an arbitrary rational $q = m_1/m_2$ with both m_1 and m_2 being integers. Then we claim that $f^*(\tilde{R}) = f^*(\tilde{R} + (1/m_2)\tilde{a})$. To establish this we observe that if $f^*(\tilde{R}) > f^*(\tilde{R} + (1/m_2)\tilde{a})$ then from (5) $f^*(\tilde{R} + (1/m_2)\tilde{a}) > f^*(\tilde{R} + (2/m_2)\tilde{a})$ and we can continue in this way m_2 times to show $f^*(\tilde{R}) > f^*(\tilde{R} + \tilde{a})$ which is a contradiction. In exactly the same way we get a contradiction if $f^*(\tilde{R}) < f^*(\tilde{R} + (1/m_2)\tilde{a})$. Thus we have established that $f^*(\tilde{R}) = f^*(\tilde{R} + (1/m_2)\tilde{a})$. But now we may use (6) repeatedly to show $f^*(\tilde{R} + (1/m_2)\tilde{a}) = f^*(\tilde{R} + (2/m_2)\tilde{a})$, $f^*(\tilde{R} + (2/m_2)\tilde{a}) = f^*(\tilde{R} + (3/m_2)\tilde{a})$ and so on (provided that we do not reach a vector with a negative component). After m_1 repetitions we have the equality we want: $f^*(\tilde{R}) = f^*(\tilde{R} + (m_1/m_2)\tilde{a})$.

Thus the set $L = \{x : f^*(\tilde{R}) = f^*(\tilde{R} + x)\}$ is the intersection of a linear subspace with $\mathbb{R}_+^s$. Continuity of f^* implies that the set $\tilde{R} + L$ divides $\mathbb{R}^s$ into two half spaces where $f^*(\tilde{R}) < f^*(x)$ and $f^*(\tilde{R}) > f^*(x)$ and so L has dimension $s - 1$. We have L defined by its normal vector p with $L = \{x : p^\top x = 0\}$.

Now consider any other vector of regrets $\tilde{R}'$ with $f^*(\tilde{R}) \neq f^*(\tilde{R}')$. Then we may use (6) with $\tilde{a} = \tilde{R}' - \tilde{R}$ to show that

$$f^*(\tilde{R}') = f^*(\tilde{R}' + x) \text{ for any } x \in L.$$

Thus f^* is a function with sets of equal values all being translations of the linear subspace L . From this we deduce that we can write $f^*(x)$ as a function of the scalar $p^\top x$.

We may scale so that $\sum p_i = 1$, and then f^* non-decreasing shows $p_i \geq 0$, so these can be interpreted as probabilities. We have established that $f^*(\tilde{R}) > f^*(\tilde{R}')$ if and only if $\sum p_i R_i > \sum p_i R'_i$ and the choice between two decisions is made by minimising the expected regret under probabilities p_i . As we observed earlier this is equivalent to minimising expected costs, since

$$\sum_{i \in S} p_i R_i(x) = \sum_{i \in S} p_i C_i(x) - \sum_{i \in S} p_i \inf_{z \in D} C_i(z).$$

$\square$

The failure of minimax regret to satisfy the IIA property also allows the possibility of non-transitivity in decision-making in the sense made clear in Example 4 given in the appendix.

4.1. Minimax regret for project choice

We now return to the question of choosing between projects using a minimax regret decision rule. We use the same notation as before, so that U is the complete set of projects. Thus the regret function for a selection T_0 is $R_i(T_0) = C_i(T_0) - \inf_{T \subset U} C_i(T)$ and the minimax regret choice is $\min_{T \subset U} \max_{i \in S} R_i(T)$.

An example of the use of minimax regret as a basis for project decisions of this sort occurs with the UK National Grid process for Network Options Assessment (National Grid ESO, Network Options Assessment Methodology, 2020). Each year a process is carried out to determine which network reinforcement projects should be started. A minimax regret analysis is carried out as part of this decision process.

In this case with $D = \{0, 1\}^n$ the IIA property states that changes in the costs or availability of a particular selection cannot alter the preference between two other selections. For example adding a restriction that two particular projects interfere with each other and cannot both be carried out together will eliminate all the selections including both these projects. The IIA property implies that adding this restriction will not affect the optimal solution unless the optimal solution has both these projects selected. However, there is a weaker project-based IIA property that is also of interest. Suppose that a decision rule Γ is defined as the mapping from a set of available projects to the set of projects selected, and we write $S_\Gamma(A) \subseteq A$ for the set of projects selected under Γ given a set A of available projects. Then Γ has the project-based IIA property if $n \notin S_\Gamma(1, 2, \dots, n)$ implies that $S_\Gamma(1, 2, \dots, n) = S_\Gamma(1, 2, \dots, n - 1)$. In

other words leaving out of the consideration set a project which is not chosen makes no difference to the choice. The project-based IIA property follows from the usual IIA property simply by leaving out all the project selections involving the missing project one at a time. Since this is weaker than the IIA property, we may wonder whether the project-based IIA property will hold for the minimax cost rule with independent and additive costs. The next example shows that it does not.

Example 3. There are 3 projects (X, Y, Z) and 3 scenarios. The cost values are as follows

	X	Y	Z
Scenario A	-1	1	0
Scenario B	-1	-2	3
Scenario C	1	-2	-2

In the case that $W_A = W_B = W_C = 0$ we get cost values

	$\emptyset$	{X}	{Y}	{Z}	{X, Y}	{X, Z}	{Y, Z}	{X, Y, Z}
A	0	-1	1	0	0	-1	1	0
B	0	-1	-2	3	-3	2	1	0
C	0	1	-2	-2	-1	-1	-4	-3

and regret values

	$\emptyset$	{X}	{Y}	{Z}	{X, Y}	{X, Z}	{Y, Z}	{X, Y, Z}
A	1	0	2	1	1	0	2	1
B	3	2	1	6	0	5	4	3
C	4	5	2	2	3	3	0	1

Hence the minimax regret policy chooses just project Y (with a worst regret of 2).

Now suppose that project Z is not available. Then we obtain the following regret values:

	$\emptyset$	{X}	{Y}	{X, Y}
A	1	0	2	1
B	3	2	1	0
C	2	3	0	1

The minimax regret policy chooses both projects X and Y (with worst regret of 1). Thus the presence or absence of project Z will determine whether project X is selected, even though Z is not itself selected.

The absence of a project-based IIA property is important because it means that making decisions amongst a set of projects, when different projects are proposed by different parties, is liable to gaming. This occurs if a firm or organisation can put forward a project knowing that it will not be chosen, but with the intention of pushing the decision maker into selecting a second project in which the firm has an interest.

5. Balancing capacity investment and risk

In a planning context it is often the case that there is a set of investments that need to be determined and there is uncertainty with respect to the demand that will occur. Investment gives greater capacity to meet high demands on the system. In these circumstances, rather than an individual scenario corresponding to a single set of demands, it may be better to let each scenario represent a distribution of these demands, but with a specific mean. The costs associated with the risk of demand exceeding capacity are then related to the safety margin between the set of capacities built and the average demand levels.

In this case each scenario i is associated with a vector of (mean) demands a_i . The decision to build a vector of capacities x then has investment cost $k^\top x$ for some vector k of costs, together with the costs associated with the risk of demand exceeding capacity in scenario i , which are often simply a function $h(x - a_i)$ of the difference between mean demand and capacity.

A new element in our discussion here is that there is a structure on the scenarios arising from the position of the mean demand vectors a_i in $\mathbb{R}^n$. Our next result shows that the extreme scenarios (for the convex hull of the vectors a_i , $i \in S$) are the ones that determine the minimax solution.

Proposition 5. Suppose that the costs in each scenario $i \in S$ are given by

$$f_i(x) = h(x - a_i) + k^\top x, \quad (7)$$

for some quasiconvex function h . Suppose further that the scenario $j \in S$ is such that a_j is a convex combination of the a_i corresponding to the remaining scenarios, so that, for some $S_j \subset S$ with $j \notin S_j$, we have $a_j = \sum_{i \in S_j} \lambda_i a_i$ where each $\lambda_i \geq 0$ and $\sum_{i \in S_j} \lambda_i = 1$. Then the scenario j can be removed from S without any effect on the solution of either the minimax cost problem or the minimax regret problem.

Proof. We begin by showing that the regret functions are also of the form (7). Let x^* minimise $h(x) + k^\top x$. Then x^* also minimises $h(x) + k^\top(x + a_i)$ and so $x^* + a_i$ minimises $f_i(x)$. Thus the regret in each scenario i is given by

$$R_i(x) = h(x - a_i) + k^\top x - h(x^*) - k^\top(x^* + a_i)$$

so that $R_i(x) = h(x - a_i) + k^\top(x - a_i)$ plus a constant. Thus by redefining the function h we obtain that $R_i(x)$ is also of the form (7) without any additional linear term (i.e. it is of the form $f_i(x)$ with $k = 0$).

Thus it is enough to show the result for the minimax cost problem. But now note that from $a_j = \sum_{i \in S_j} \lambda_i a_i$ for $j \notin S_j \subset S$, we may deduce

$$\begin{aligned} f_j(x) &= h\left(\sum_{i \in S_j} \lambda_i (x - a_i)\right) + k^\top x \\ &\leq \max_{i \in S_j} h(x - a_i) + k^\top x \\ &= \max_{i \in S_j} f_i(x) \end{aligned}$$

where the inequality follows from the quasiconvexity of h . Thus, removing scenario j leaves the minimax problem unchanged as required. $\square$

Since, by Caratheodory's theorem, there is a set $K \subseteq S$ of size at most $n + 1$ such that every a_j , $j \in S$, may be written as a convex combination of the a_i , $i \in K$, it follows that (analogously to Proposition 1), under the conditions of Proposition 5, both the minimax cost and minimax regret decisions are determined solely by those scenarios in the set K .

5.1. Application to electricity capacity procurement in Great Britain

Each year National Grid ESO (the GB electricity system operator) produces an Electricity Capacity Report (ECR), the purpose of which is to recommend a generation *capacity-to-secure*, via a “capacity auction”, in order to secure adequate GB electricity supplies in a given future winter. The analysis of the ECR is usually based on consideration of a small number of core scenarios for the future winter under study. These core scenarios are mostly based on the Future Energy Scenarios which National Grid ESO produce annually for a number of purposes. They represent varying assumptions about the evolution of the electricity system and the evolution of demand. Additionally, for each of a number of these core scenarios, there is defined a number of variant scenarios – referred to as “sensitivities” – each corresponding to the “parent” core scenario, but varying assumptions about weather conditions and their consequences in the winter under study. The set D of possible decisions is essentially the set of capacities which might be recommended to be secured, and may be identified with $\mathbb{R}_+$. For each i




Fig. 1. Regret functions $R_i(x)$ for each of the five core and two most extreme variant scenarios considered by the 2015 ECR. The dotted line corresponds to the minimax regret capacity-to-secure of 47.79 GW.

in the scenario set S , there is a cost function C_i on D given by

$$C_i(x) = \text{VOLL} \times \text{EEU}_i(x) + \text{CONE} \times x, \quad (8)$$

where x is the possible generation capacity to be secured (in MW) and $\text{EEU}_i(x)$ is the corresponding *expected energy unserved* (in MWh) over the given winter. The constant VOLL is the *value of lost load* (£/MWh) while the constant CONE is the *cost of new entry* (procurement cost per unit of new generation capacity for the winter under study, in £/MW). The functions $\text{EEU}_i(x)$ are estimated from data and are all approximately of the form $\exp(-\lambda(x - a_i))$ for some λ independent of i , representing the exponential decay of risk as a function of the level of generation secured.

The ECR reports use minimax regret analysis – on the set S of all scenarios considered (both core and variants) – to determine a recommended generation *capacity-to-secure*. Since the cost functions (8) are all approximately of the form (7), it follows from Proposition 5 that the recommended *capacity-to-secure* is then the same as that which would be determined by the use of only the two most extreme scenarios, i.e. those with the least and greatest values of the parameters a_i .

While it is not considered appropriate to attempt to assign probabilities to the core scenarios – since they correspond to, for example, different evolutions of public policy – it might reasonably be possible to assign a conditional probability to each of the variant scenarios, given the occurrence of the parent core scenario. While the National Grid ESO analysis does not do this, it is nevertheless careful to exclude variant scenarios whose overall unconditional probabilities of occurrence are considered to be very small, as these would typically have extreme values of the parameters a_i and thus have an undue influence on the outcome of the minimax regret analysis. The outcome of the National Grid ESO analysis is therefore very sensitive to decisions about which variant scenarios should be included.

It is natural to consider here the robust optimisation problem $\text{RO}(\mathbb{P}_A)$ introduced in Section 2. It is convenient to consider the data of the 2015 ECR (National Grid EMR Electricity Capacity Report, 2015), as best illustrating the points we wish to make. This report was concerned with the provision of adequate electricity capacity for the winter 2019–20. The report considered a set S of five core scenarios for the above winter. In increasing order of the parameters a_i , these were labelled DECC (a scenario proposed

by the Department of Energy and Climate Change), Gone Green, Slow Progression, Consumer Power and No Progression. For each of the two central scenarios, Slow Progression and Consumer Power, there were defined a further seven variant scenarios, or sensitivities, corresponding to lower or higher levels of conventional generating plant availability (than that considered in the corresponding core scenario), lower or higher levels of demand, and a weather-related lower level of wind generation.

In the 2015 ECR the constants VOLL and CONE were taken to be £17,000/MWh and £49,000/MW respectively. Fig. 1 plots the regret functions $R_i(x) = C_i(x) - \inf_{z \in D} C_i(z)$ corresponding to the five core scenarios and the two most extreme of the 14 variant scenarios (i.e. those with the least and greatest values of the parameters a_i). The latter two scenarios essentially determine the outcome of the minimax regret analysis, which is to obtain a *capacity-to-secure* of 47.79 GW.

We can see that there is considerable arbitrariness in this process. Notice that there are no variant scenarios introduced for three of the core scenarios. These are more extreme than the other core scenarios and so the inclusion of variants would impact the analysis. The factors involved in generating variant scenarios, such as high or low levels of demand, will apply to each of the core scenarios equally. Thus excluding these scenarios within the planning process is hard to justify.

To demonstrate the robust optimisation approach we will define the seven “missing” variant scenarios associated with each of the remaining three core scenarios. Thus the scenario set S may be partitioned as $S = \bigcup_{i=1}^5 S_i$, where each S_i contains eight scenarios – the parent core scenario, labelled $i1$ and a further seven variant scenarios, or sensitivities, labelled ij for $j = 2, \dots, 8$. For each core scenario $i1$ for which it is necessary to define the additional variant scenarios ij , the corresponding cost functions $C_{ij}(x)$ given by (8) are obtained by taking, for each $x \in D$, the ratio $\text{EEU}_{ij}(x)/\text{EEU}_{i1}(x)$ to be the geometric mean of the two corresponding ratios for those two core scenarios for which the variant scenarios are already defined and for which we have data. (These two ratios are in general very similar, and so the data for the additional variant scenarios – although not directly available – is likely to be sufficiently accurate for the present illustrative purposes.)

We refrain from assigning probabilities to the sets S_i , but for each i , we place constraints on the conditional probabilities p_{ij} of




Fig. 2. The functions $f_i(x)$ given by (10), each associated with a parent scenario and its seven variant sensitivities. The dashed and dotted lines correspond respectively to the robust optimisation and minimax regret capacities-to-secure of 47.65 and 47.79 GW.

the scenarios within the set S_i , given the event that one of the scenarios within the set S_i occurs. With sufficient weather and other data, we might provide good estimates of these conditional probabilities; however, lacking this, we content ourselves with the requirement that, for some $a < 1$, for any of the variant scenarios ij , $ij \in S_i$, its conditional probability p_{ij} satisfies the constraints

$$ap_{i1} \leq p_{ij} \leq p_{i1}, \quad (9)$$

where $\sum_{ij \in S_i} p_{ij} = 1$ and where p_{i1} is the conditional probability of the parent scenario in the set S_i (again given that some scenario within the set S_i occurs). Hence, for each $i = 1, \dots, 5$ and each positive capacity-to-secure x , define

$$f_i(x) = \max_{(p_{ij}: \sum_{ij \in S_i} p_{ij} = 1)} \sum_{ij \in S_i} p_{ij} R_{ij}(x), \quad (10)$$

where $R_{ij}(x)$ is the regret function associated with the scenario ij and where the above conditional probabilities p_{ij} satisfy also the additional constraints (9). It is then easily seen that the robust optimisation problem $RO(\mathbb{P}_A)$ given by (1) becomes

$$\min_{x \in D} \max_{i=1, \dots, 5} f_i(x), \quad (11)$$

i.e. the standard minimax problem applied to the functions $f_i(x)$.

We take the constant a in (9) to be given by $a = 0.1$. Fig. 2 plots the functions $f_i(x)$, $i = 1, \dots, 5$, each associated with a parent scenario and its seven variants. The value $x \in D$, i.e. the capacity-to-secure, which solves the robust optimisation problem given by (11) is then $x = 47.65$ GW, which is just a little less than the 47.79 GW obtained from the minimax regret analysis of the 2015 ECR. The results are similar, and we believe that the original recommended capacity is essentially correct. National Grid ESO have carried out this exercise being careful to exclude more extreme scenarios that might unduly have affected the results of the analysis. Our robust optimisation approach generates the required result using a principled and consistent approach, without the need to make arbitrary decisions on which extreme solutions to include.

6. Conclusions

An important property of the minimax decision-making criterion (whether applied to regret functions or to cost functions), is that the outcome of the analysis is typically determined by a small number of scenarios. This is an issue that needs to be considered by planners making use of these robust approaches, since sensitivity to the scenarios chosen reduces the apparent advantage (from a public planning perspective) of there being no requirement to assign probabilities to scenarios. We explore this issue and show how it can be helpful to reformulate the minimax rule through making use of an uncertainty set which is a convex subset of the set of all probability measures. This allows the use of a partial ordering between the probabilities for “core” and “extreme” scenarios that in turn will widen the scenario set considered. This leads to effective decision rules that do not make any assumptions about the probabilities of core scenarios capturing the deep uncertainty regarding these, but at the same time make effective use of information that is available on the relative likelihood of different deviations around these core scenarios. We have demonstrated how this can work on a case study from the GB electricity capacity market.

Our work has further focused on minimax regret analysis. It is well known that this fails to satisfy usually accepted conditions of economic rationality and in particular the IIA property. We give a variety of examples to show how widespread this behaviour is, both with continuous and discrete decision sets, whenever regret values are the basis for decisions. We show that with a continuous decision space the IIA property can only be retained when the regret values for different scenarios are combined through taking an expectation over a set of probabilities, which then reduces the problem to one of minimising expected costs.

One context in which minimax regret can be used as a planning tool occurs when projects are put forward for potential funding. The planner must then determine the selection of projects that will proceed. If the projects are proposed by a number of different firms, then the absence of even a weakened form of the IIA property leaves minimax regret open to being gamed. Our analysis shows the importance of being careful in these circumstances.

Despite these drawbacks from using minimax regret instead of minimax cost, we expect that it will continue to be used in many public planning contexts. The alternative of minimax cost implies

an estimate of the broader societal costs that vary between one scenario and another, and this is likely to be impossible in practice.

Acknowledgements

The authors would like to thank the Isaac Newton Institute for Mathematical Sciences for support during the Mathematics of Energy Systems programme (<https://www.newton.ac.uk/event/mes/>), when early work on this paper was undertaken.

Appendix A. Further examples

Example 4 (Failure of transitivity). We take the decision space D to consist of three possibilities $D = \{x, y, z\}$ and we again consider a scenario space $S = \{A, B, C\}$. The costs of the decisions under each scenario are given by

	x	y	z
Scenario A	4	0	2
Scenario B	4	6	0
Scenario C	0	0	5

It is now readily verified that, under the minimax regret criterion, if the decision z is not available then the decision y is preferred to the decision x , while if the decision x is not available then the decision z is preferred to the decision y , and, finally, if the decision y is not available then the decision x is preferred to the decision z .

Example 5 (Project choice and more than $n + 1$ active scenarios). We consider choosing between n projects so $D = \{0, 1\}^n$ and we give an example where more than $n + 1$ scenarios are active, in the sense that removing any one of them changes the minimax regret solution. We take 3 projects (X , Y and Z) and 5 scenarios (A , B , C , D and E) with the following set of costs

costs	project X	project Y	project Z
scenario A	6	−2	−4
scenario B	2	4	4
scenario C	4	−8	−1
scenario D	−6	6	0
scenario E	−2	−7	1

The costs for different choices of sets of projects are:

	$\emptyset$	$\{X\}$	$\{Y\}$	$\{Z\}$	$\{X, Y\}$	$\{X, Z\}$	$\{Y, Z\}$	$\{X, Y, Z\}$
scenario A	0	6	−2	−4	4	2	−6	0
scenario B	0	2	4	4	6	6	8	10
scenario C	0	4	−8	−1	−4	3	−9	−5
scenario D	0	−6	6	0	0	−6	6	0
scenario E	0	−2	−7	1	−9	−1	−6	−8

We can convert this into a set of regret figures:

	$\emptyset$	$\{X\}$	$\{Y\}$	$\{Z\}$	$\{X, Y\}$	$\{X, Z\}$	$\{Y, Z\}$	$\{X, Y, Z\}$
scenario A	6	12	4	2	10	8	0	6
scenario B	0	2	4	4	6	6	8	10
scenario C	9	13	1	8	5	12	0	4
scenario D	6	0	12	6	6	0	12	6
scenario E	9	7	2	10	0	8	3	1
max regret	9	13	12	10	10	12	12	10

The minimax regret policy with all scenarios present is $\emptyset$, i.e. to accept none of the projects. But remove scenario A and we choose $\{X, Y\}$; remove scenario B and we choose $\{X, Y, Z\}$; remove scenario C and we choose $\{X, Z\}$; remove scenario D and we choose $\{Y\}$; and remove scenario E and we choose $\{Z\}$. Thus dropping any scenario will cause a change in the decision made, so that all 5 scenarios contribute to the decision.

References

Anderson, E., Xu, H., & Zhang, D. (2020). Varying confidence levels for CVaR risk measures and minimax limits. *Mathematical Programming*, 180(1), 327–370.

- Anscombe, F. J., & Aumann, R. J. (1963). A definition of subjective probability. *Annals of Mathematical Statistics*, 34(1), 199–205.
- Armstrong, M., & Vickers, J. (2010). A model of delegated project choice. *Econometrica*, 78(1), 213–244.
- Bell, D. E. (1982). Regret in decision making under uncertainty. *Operations Research*, 30(5), 961–981.
- Bleichrodt, H., & Wakker, P. P. (2015). Regret theory: A bold alternative to the alternatives. *The Economic Journal*, 125(583), 493–532.
- Calafiore, G., & Campi, M. C. (2005). Uncertain convex programs: Randomized solutions and confidence levels. *Mathematical Programming*, 102(1), 25–46.
- Chang, N. B., & Davila, E. (2007). Minimax regret optimization analysis for a regional solid waste management system. *Waste Management*, 27(6), 820–832.
- Chen, B., Wang, J., Wang, L., He, Y., & Wang, Z. (2014). Robust optimization for transmission expansion planning: Minimax cost vs. minimax regret. *IEEE Transactions on Power Systems*, 29(6), 3069–3077.
- Daskin, M. S., Hesse, S. M., & Revelle, C. S. (1997). α -reliable p -minimax regret: A new model for strategic facility location modeling. *Location Science*, 5(4), 227–246.
- Dent, C., French, S., & Zachary, S. (2020). Decision making for future energy systems: Incorporating rapid change and future uncertainties. Report for OFGEM available at <http://www.ofgem.gov.uk/publications-and-updates/decision-making-future-energy-systems>.
- Diecidue, E., & Somasundaram, J. (2017). Regret theory: A new foundation. *Journal of Economic Theory*, 172(1), 88–119.
- Gilboa, I., & Schmeidler, D. (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2), 141–153.
- Crespo del Granado, P. C., Skar, C., Doukas, H., & Trachanas, G. P. (2019). Investments in the EU power system: A stress test analysis on the effectiveness of decarbonisation policies. In H. Doukas, A. Flamos, & J. Lieu (Eds.), *Understanding risks and uncertainties in energy and climate policy* (pp. 97–122). Springer. <https://doi.org/10.1007/978-3-030-03152-7>.
- Gul, F. (1992). Savage's theorem with a finite number of states. *Journal of Economic Theory*, 57(1), 99–110.
- Hayashi, T. (2008). Regret aversion and opportunity dependence. *Journal of Economic Theory*, 139(1), 242–268.
- Hughes, N., & Strachan, N. (2010). Methodological review of UK and international low carbon scenarios. *Energy Policy*, 38(10), 6056–6065.
- Insua, D. R., Ruggeri, F., & Vidakovic, B. (1995). Some results on posterior regret γ -minimax estimation. *Statistics and Decisions*, 13, 315–331.
- Jiang, R., Wang, J., Zhang, M., & Guan, Y. (2013). Two-stage minimax regret robust unit commitment. *IEEE Transactions on Power Systems*, 28(3), 2271–2282.
- Klamroth, K., Köbis, E., Schöbel, A., & Tammer, C. (2017). A unified approach to uncertain optimization. *European Journal of Operational Research*, 260(2), 403–420.
- Kunreuther, H., Heal, G., Allen, M., Edenhofer, O., Field, C. B., & Yohe, G. (2013). Risk management and climate change. *Nature Climate Change*, 3(5), 447–450.
- Loulou, R., & Kanudia, A. (1999). Minimax regret strategies for greenhouse gas abatement: Methodology and application. *Operations Research Letters*, 25(5), 219–230.
- Marchau, V. A., Walker, W. E., Bloemen, P. J., & Popper, S. W. (2019). *Decision making under deep uncertainty: From theory to practice*. Springer Nature.
- McCollum, D. L., Gambhir, A., Rogelj, J., & Wilson, C. (2020). Energy modellers should explore extremes more systematically in scenarios. *Nature Energy*, 5(2), 104–107.
- Mild, P., Liesiö, J., & Salo, A. (2015). Selecting infrastructure maintenance projects with robust portfolio modeling. *Decision Support Systems*, 77, 21–30.
- Milnor, J. (1951). Games against nature. *RAND Project Air Force Report*. Santa Monica CA.
- Moreira, A., Strbac, G., Moreno, R., Street, A., & Konstantelos, I. (2017). A five-level MILP model for flexible transmission network planning under uncertainty: A min-max regret approach. *IEEE Transactions on Power Systems*, 33(1), 486–501.
- National Grid (2020). Future energy scenarios. <https://www.nationalgrideso.com/future-energy/future-energy-scenarios>.
- National Grid ESO, Network Options Assessment Methodology (2020). Report available at <https://www.nationalgrideso.com/document/174231/download>.
- National Grid EMR Electricity Capacity Report (2015). Report available at <https://www.emrdeliverybody.com/CM/Capacity.aspx>.
- Ray, P. (1973). Independence of irrelevant alternatives. *Econometrica*, 41(5), 987–991.
- Sanders, A., Hart, A., Ravishanker, M., Brunert, J., Strbac, G., Aunedi, M., & Pudjianto, D. (2016). *An analysis of electricity system flexibility for Great Britain*. London, UK: Carbon Trust / Imperial College.
- Savage, L. J. (1951). The theory of statistical decision. *Journal of the American Statistical Association*, 46(253), 55–67.
- Savage, L. J. (1954). *The foundations of statistics*. New York, NY, USA: John Wiley and Sons, Inc.
- Shapiro, A. (2012). Minimax and risk averse multistage stochastic programming. *European Journal of Operational Research*, 219(3), 719–726.
- Shapiro, A., & Kleywegt, A. (2002). Minimax analysis of stochastic problems. *Optimization Methods and Software*, 17(3), 523–542.
- Stoye, J. (2011a). Axioms for minimax regret choice correspondences. *Journal of Economic Theory*, 146(6), 2226–2251.
- Stoye, J. (2011b). Statistical decisions under ambiguity. *Theory and Decision*, 70(2), 129–148.
- Van der Pol, T. D., Gabbert, S., Weikard, H. P., van Ierland, E. C., & Hendrix, E. M. T. (2017). A minimax regret analysis of flood risk management strategies under climate change uncertainty and emerging information. *Environmental and Resource Economics*, 68(4), 1087–1109.