Expected Utility and Climate Change

Lisa Wolring

S1271105

l.wolring@umail.leidenuniv.nl

Bruno Verbeek

Philosophy Politics Economics

2016

Contents

Abstract: This paper investigates whether expected utility theory is a proper method for decision-making given the conditions of uncertainty surrounding climate change. I explain what expected utility theory is and how Ramsey's theory of partial belief can be used to infer subjective probabilities. By the use of the Allais paradox I show that one of the axioms, namely *independence* is unlikely to be satisfied. First, it can be the case that when applying expected utility theory one does not assign utilities and probabilities in a consistent manner. Second, I explain that the example of the Allais paradox is part of a broader phenomenon in which what is considered good about an option cannot be reduced to the goodness or badness of individual outcomes of this option. I conclude that expected utility theory is an improper method for decision-making about climate change policy.

1. Climate Change and Expected Utility Theory

In this paper I will examine whether the expected utility theory is a proper or suitable method for making rational choices and developing policy proposals given the conditions of uncertainty surrounding climate change. Climate change is a complex problem. For most average people it is hard to comprehend the various possible scenarios that are discussed by climate scientists. For example, what to make of the claim that there is a probability that the temperature will rise 8°C degrees, the Antarctic will melt and the sea-level will rise 70 meters? No doubt, if this were to become reality it would be disastrous for life on earth, however, I cannot properly assess how bad this would be given that it is only probable. A general response to something this bad and big is a very short explosion of stress and panic, after which one continues with one's daily activities and forgets about it. Since I cannot act on this problem emotionally or even comprehend what this means to me, to others or to future others, maybe I can understand it in a more calculative or rational way.

One of the many problems regarding the climate change issue is that the future development of climate change is more or less uncertain. However, it is certain that climate change will cause and already is causing problems mostly for the least well off people and for various kinds of animals and plants. What is uncertain is the exact scenario that will unfold and how bad it really will be, which of course also depends upon our present actions among other things. One way to handle the problem of making a rational choice or making policy under uncertainty is to use expected utility theory. In "Valuing Policies in Response to Climate Change: Some Ethical Issues" John Broome proposes to use the expected utility theory in order to make a rational choice on how to act on climate change. By using this method one should be able to calculate the expected utility of the various actions or policies by multiplying the probability of an outcome with, what John Broome calls, the 'badness rate' of the outcome ("The most important thing" 105). The 'badness rate' is the utility of an outcome. The possible outcomes in regards to climate change are all bad, some just a little and some catastrophically bad. In this essay I will discuss expected utility theory and explain how Broome intends to use it for climate change. Broome focuses on finding a measure for the badness of the various outcomes, but he does not work out the probability part of the calculation. I will attempt this latter task by adopting Frank P. Ramsey's theory of probability as partial belief. Then I will discuss one major problem with the expected utility theory and examine how this will, among other things, be problematic for the practical application of the method to climate change. I will conclude that using the expected utility on the climate change issue is improper and unsuitable.

2. Expected Utility Theory

Expected utility theory offers a method for making rational decisions under conditions of uncertainty. Broome writes that the theory is "often presented as a theory about people's preferences- either a theory of the preferences people do have when facing uncertainty, or of the preferences they ought to have" ("Valuing Policies" 2). However he writes that he will use this theory as a "theory of value" (ibid). So will I. This means that if A has a greater utility than B it means that A is objectively better than B. By using the expected utility theory one is able to make a rational choice under uncertainty. The idea is that by calculating the expected utility of the various options open to one, one selects the rational choice or policy by opting for the one with the highest expected utility. The expected utility is calculated by combining the probability of an event with the utility of an outcome.

For example, imagine the following situation: you are sneezing very often, at least once every hour. There is a pill that will stop your sneezing. The probability that it will help is 98 percent, however there is also a probability of 2 percent that you will die after taking this pill. One can say that there is a very high chance that the problem will be solved when taking this pill, but the very small chance of dying probably weighs heavier in your calculations about what to do, since dying is valued as much more 'bad' than sneezing. Now one has to assign utility to 'stopping sneezing' or a negative value to 'sneezing'. Let us say that 'stopping sneezing' has a utility of 10 and assume that not dying has a utility of 10.000. So, when one calculates the expected utility (EU[1]) of the two options by multiplying the probability (Pr[2]) with the perceived utility (U[3]), one will conclude that one should not take the pill. Taking the pill will lead to an EU of 9.8 (0,98x10) and not taking the pill to an EU of 200 (0,02x10.000).

It is sometimes hard to assign an exact number to the 'badness' or 'goodness' of an outcome[4], or to the probability of an event. In the example above we can presume that the probability degrees are inferred from frequency of occurrences of death and recovery in the past. There are many situations in which we cannot rely on such frequencies, for example, because it is a new phenomenon or simply because the frequencies are not available or significant and therefore inconclusive. Still in this situation, with limited knowledge about the frequency one is able to judge certain events as reasonably more probable than others. Frank Ramsey argues that frequency cannot always generate a probability and that sometimes one needs to use subjective probabilities. Therefore Ramsey proposed a subjective probability measure. Broome also writes that in the case of climate change "there are no objective probabilities and that suitable subjective probabilities are very hard to come by" ("Valuing Policies" 8). Broome proposes the following: "I think that the only adequate way of accommodating differing views is to assess the values of all the policies several times, each time using different probabilities. This is a sort of

---

[1] "EU" = Expected Utility
[2] "Pr" = Probability
[3] "U"= Utility
[4] I will not go into detail about how to value outcomes. Broome goes into detail, read for example *Weighing Lives* and *Weighing Goods*.

sensitivity analysis" (ibid). The subjective probability measure is not further specified by Broome, therefore I will use Ramsey's method for measuring subjective probability.

Ramsey was the first to formulate probability as 'partial belief', or 'degree of belief' as it is known these days. Ramsey was a mathematician, economist and a philosopher from Cambridge. He died in 1930, just before Wittgenstein arrived in Cambridge, who he had been in contact with. Ramsey was the first to translated Wittgenstein's *Tractatus Logico Philosophicus* into English. Most people were focussed on Wittgenstein and his writings. Therefore no one payed attention to the writings of Ramsey, only when von Neumann and Morgenstern published *Theory of Games and Economic Behavior*, Ramsey's work on 'partial belief' was rediscovered.

In *Truth and probabilities* Ramsey sets out to find a way of measuring the degree of belief an agent has in a given proposition. There are situations in which we do not have enough information to easily assign a probability to a proposition or an event. I will explain Ramsey's theory of probability as partial belief and how by this method one is able to create an expected utility measure. One uses the same method to calculate the expected utility of an outcome as well as that of an option.

3. Ramsey's Theory of Partial Belief

Ramsey writes that it is commonly believed that belief cannot be measured. However, he says: "unless we are prepared to give up the whole thing as a bad job we are bound to hold that beliefs can to some extent be measured" (62). Ramsey regards the betting method as "fundamentally sound" (68). By proposing bets, Ramsey is able to order outcomes by their utility, to create a utility interval scale and to assign probabilities to events. In short, what should be done is the following: all outcomes should be ordered from least preferred to most preferred and by using the betting method their utilities must be assigned. Then the probabilities should be inferred from the utilities. Next, to establish the expected utility of an option, one should multiply the utility of an outcome with a probability of a particular event. I will now first explain how one is able to assign utilities by the use of the betting method. Then I will explain how to infer probabilities from utilities, according to Ramsey.

a. Utilities and Probabilities

The claim that one can infer probabilities from utilities seems to be a strange one. I will use an example to illustrate how one is able to do the 'trick'. First we need to utilities to all the outcomes. Suppose, for example, that Frank has the following outcomes available: bananas, pineapples and grapes. Now Frank is asked to order these fruits from best to worst. Frank prefers bananas over pineapples and grapes, and pineapples over grapes. I will use a B for bananas, a P for pineapples and a G for grapes. So, we have an order of preference of these fruits by Frank looking as follows:

$$B > P > G$$[5]

This is an ordinal scale, the items are only ordered from best to worst. This scale does not tell anything about the interval difference between B, P and G. It could be that B is just a little more preferred over P and that P is preferred much more over G. In order to calculate the interval differences between the preferences Frank has for the different fruits, we must construct an interval measure U that also represents the relative difference between the outcomes.

Suppose that the most preferred option, B is assigned a U of 1 and the least preferred option, G is assigned a U of 0. So, U (B) =1 and U (G) =0. In order to establish U (P) one should offer Frank the following choice: option 1) getting P for sure or option 2) taking the gamble between getting G or B with a probability. Now the question is: how high should the probability for getting B be, in order for Frank to be indifferent between option 1 and 2? This question asked to Frank can be written as follows:

$$P \sim [\ G\ (1-p);\ B\ (p)]$$[6]

Suppose that Frank is indifferent when he gets B with the $p = 0.6$ and G with $1-p = 0.4$ We assigned U (B) = 1 and U (G) = 0. Frank is indifferent between the two options when these have the same utility. So, in order for Frank to be indifferent between option 1 and 2, U (P) should have the same utility as U (G (1-p)) +U (B (p)):

$$U\ (P) = [U\ (G)\ 0{,}4\ ;\ U\ (B)\ 0{,}6]$$

[5] ">" means "more preferred than"
[6] "~" means "is indifferent to" and "[ ]" and in between these is the proposed bet

$$U(P) = [0 \times 0{,}4 \; ; \; 1 \times 0{,}6]$$

This gives us: U (P) = 0.6. Frank should assign 0,6 to the utility of P, because if the utility of P is 0,6, then Frank is indifferent between both options.

In this manner it is possible to establish utilities for all fruits that are utility-wise in between B and G, representing their preference ordering. For example one could add an apple (A). Let's suppose that Frank prefers A over P and B over A. This means that U (A) is in between 0.6 and 1. The question should be asked to Frank with what probability *q* he will be indifferent between A for sure or the gamble between B and G. This looks as follows:

$$A \sim [\, G \; 1 - q; \; B \; q]$$

If Frank will be indifferent between A for sure and the gamble between P and B if q=0,7 then U (A) = 0,7. Because: 0 x 0,3 + 1 x 0,7 = 0,7

If Frank finds an outcome that is better than our best option B, he should do the following. Let us name this outcome, which is better than our best option, AA. Frank should find any outcome, for example D, on the utility scale such that U (B) = ½ x U (AA) + ½ x U (D). So, if U (B) is 1 and U (D) is 0,9 then U (AA) = 1,1, since 1 = ½ U (AA) + 0,45. The same procedure applies to determine the utility of an outcome to which G is preferred. Once all the relevant outcomes are assigned a utility, Frank should be able to calculate the expected utility of the options available provided he knows the relevant probabilities.

Suppose that Frank wants to make a rational choice under uncertainty. He has two options to choose from. The choice is between randomly picking from fruit basket one or from fruit basket two. Imagine fruit basket number one with the following content: 5 apples, 2 bananas, 2 pineapples and 1 grape. Imagine fruit basket number two with the following content: 1 apple, no bananas, 6 pineapples and 3 grapes. Because Frank knows the content of the fruit baskets, he is able to calculate how probable it is to get any fruit from basket one or basket two. Furthermore there are events in the world, namely for example event Ev_2 in which someone came into the room just at the moment you pick a fruit and because of that you moved your hand a little to the right. Events have a probability of occurring. There are also outcomes, for example the outcome ‘apple’ or the outcome ‘banana’. These outcomes have the utilities as we assigned by using the betting method above, getting an apple for example has a utility of 0,7. So, in our calculation there are three elements; options, events and outcomes.

Suppose Frank finds that in event$_1$ (Ev$_1$[7]) he will pick an apple in both cases. In Ev$_2$ he will pick an apple from basket one and a banana from basket two. In Ev$_3$ he will pick a banana from basket one and a grape from basket two. In Ev$_4$ he will pick a grape in both cases. The question Frank wants to answer is whether it is better, meaning having a higher expected utility, to take one random piece of fruit from basket one or from basket two.

Take a look at Table 1:

	Ev $_1$	Ev $_2$	Ev $_3$	Ev $_4$	Ev $_5$
Pick from basket one	U(A)	U(A)	U(P)	U(B)	U(G)
Pick from basket two	U(A)	U(P)	U(P)	U(G)	U(G)

Table 1

Suppose Frank has acquired the following information:

Basket 1: 5 A, 2 P, 2 B, 1G

Basket 2: 1 A, 6 P, 0 B, 3 G

U (B) =1 U (A) = 0,7 U (P) = 0,6 U (G) = 0

Pr (Ev$_1$) = 0,1 Pr (Ev$_2$) = 0,4 Pr (Ev$_3$) = 0,2 Pr (Ev$_4$) = 0,2 Pr (Ev$_5$) = 0,1

B > A > P > G

Frank is now able to fill in table 1:

	Pr (Ev 1)0,1	Pr (Ev $_2$) 0,4	Pr (Ev $_3$) 0,2	Pr (Ev $_4$) 0,2	Pr (Ev $_5$) 0,1
Pick from basket one	U(A) 0,7	U(A) 0,7	U(P) 0,6	U(B) 1	U(G) 0
Pick from basket two	U(A) 0,7	U(P) 0,6	U(P) 0,6	U(G) 0	U(G) 0

Table 2

To be able to calculate the expected utility for basket one and two, one should multiply the probability of the event with the utility of the outcome and add these numbers for fruit basket one and for fruit basket two.

Fruit basket one: 0,1 x 0,7 + 0,4 x 0,7 + 0,2 x 0,6 + 0,2 x 1 + 0,1 x 0 = 0,67.

Fruit basket two: 0,1 x 0,7 + 0,4 x 0,6 + 0,2 x 0,6 + 0,2 x 0 + 0,1 x 0 = 0,43

It is clear that randomly picking from basket one has a higher expected utility than randomly picking from basket two. So, Frank should pick a fruit from basket one. For Frank, the

[7] "Ev" stands for event, which is a possible scenario in the world, for example the sun shines or it rains. Events have a probability of occurring.

expected utility of picking from fruit basket one is higher, given his judgement of the value of fruits. However there are situations in which there is no such information.

b. Probabilities and Utilities

By using Ramsey's method of partial belief, one is able to infer probabilities from utilities. Partial belief is the degree to which one beliefs something. If someone believes that something will be the case for sure, then the probability of this event is 1. If someone beliefs that something will not be the case for sure then the probability is 0. In both these cases one is sure of something. If someone is in complete doubt about whether something will be the case then the probability of this event is 0,5. This event, in which someone is completely unsure, is defined by Ramsey as the ethically neutral event. The event is ethically neutral when "the subject is said to have belief of degree ½ in such a proposition[8] p if he has no preference between the options (1) α if p is true, β if p is false, and (2) α if p is false, β if p is true, but has a preference between α and β simply" (Ramsey 73). Suppose that Frank finds that he is indifferent between the following gambles [B if Ev_k; G if not Ev_k] [9] and [G if Ev_k; B if not Ev_k]. Since Ev_k is an ethically neutral event Pr (Ev_k) =0,5.[10]

Frank Ramsey uses the ethically neutral position in order to assign utilities. He first finds the ethically neutral position, then he finds an outcome so that one is indifferent between that outcome for sure and the gamble between the best and the worst option. This outcome is assigned a utility U=0,5. Then he finds the outcome so that one is indifferent between this outcome and the best outcome. This outcome is assigned 0,75 utility. Again a bet is proposed between the outcome with 0,75 utility and the best outcome. This outcome is assigned 0,875. In this manner he continues to assigns utilities to all the outcomes. For this method to work one needs many outcomes available to make sure that there always is an outcome that one is indifferent to when a bet is proposed. Ramsey uses a similar method to calculate the subjective

---

[8] In my text referred to as event
[9] Ev_k is a possible event.
[10] For a more detailed explanation about Ramsey's ethically neutral position see "Truth and probability" by Frank Ramsey.

probability, this is, the degree of certainty, of specific events, using the utilities found using the procedure above.

Assume again: B>A>P>G and U (B) =1 U (A) = 0,7 U (P)= 0,6 U (G)= 0. There are the following events: $Ev_{1,}$ $Ev_{2,}$ Ev_3, $Ev_{4,}$….., $Ev_{n.}$ We want to find the subjective probabilities of Ev_1 to $Ev_{n.}$

Suppose that it is the case that:

A ~ [B if Ev_3 ; G if not – Ev_3]

Then since the expected utility of [B if Ev_3, G if not – Ev_3] = U (B) Pr (Ev_3) + U (G) (1- Pr (Ev_3))

We can infer that:

U(A) = U (B) Pr (Ev_3) + U (G) (1- Pr (Ev_3))

Substituting the values for U (A), U (B), and U (G), we get:

0,7 = 1 Pr (Ev_3) + 0 (1 – Pr (Ev_3))

Therefore Pr (Ev_3) = 0,7. Provided there are enough outcomes and events, we can always find an outcome such that we can infer the probability of an event E_i using this method. If there are more outcomes and more events then the utilities and probabilities can be more precisely assigned. We need many outcomes in order for us to always be able to find the outcome to which one is indifferent, when we propose a bet.

c. Utilities, probabilities and expected utilities

By this way of inferring utilities and probabilities one is able to combine utilities as well as probabilities in one's calculations in the way that standard probability calculus[11] allows. This is the case since utilities are inferred from probabilities. For example since U (A) = [A (p); Z (1-p)] if the p=1 then the utility assigned to A is 1. U (A) in other words, can be written as a bet.

[11] The standard probability calculus is the approach by Andrey Kolmogorov, a Russian mathematician. Kolmogorov, A. N. *Foundations of the Theory of Probability*. New York: Chelsea Pub., 1956. Print.
For an interpretation with explanation read chapter 7 “Kinds of Probability” in Papineau, David. *Philosophical Devices: Proofs, Probabilities, Possibilities, and Sets*. Oxford, England: Oxford UP, 2012.

Another example: U (M) = [Ap; Z1-p] if U (A) = 1 and U (Z) = 0 and p is 0,5 then the utility of M is 0,5 and the utility of M equals the bet between A and Z with a probability of 0,5. Since utilities are defined in terms of probabilities, one is able to use the standard probability calculations and one is able to multiply probabilities with utilities.

In order to calculate expected utility one needs a probability and a utility and one needs to multiply those. Suppose that there is a Pr (P)= 0,7, this is in other words: the probability that you will get P, and a U (P) 0,7 then one is able to calculate the expected utility by multiplying 0,7 with 0,7 which results in an expected utility of 0,49. One is also able to combine two expected utilities, for example EU(A) of 0,2 and EU (B) of 0,1 have a combined EU (AB) of 0,3[12].

4. Climate Change Example

I will give a very simplified example about climate change in order to show that what applies to decision making about fruit baskets also applies to decision making about climate change. The information I use is partly fictional, however it is based on the information from the Stern Review (Stern 67). The Stern Review is a comprehensive review about the economic consequences of climate change. There are analyses of information of many different sources, predictions are discussed and possible actions are examined. In this simplified example, a decision under uncertainty must be made between two possible options. In this example there are two options, two events and four outcomes. The first option is to do nothing and the second is to introduce carbon pricing[13]. In the first event, namely $Ev_{1,}$ the earth will warm up by 3 degrees Celsius and in second event, $Ev_{2,}$ the earth will warm up by 5 degrees Celsius. There are outcomes A, B, C and D. For example, outcome A, as described by Stern consists among other sub-outcomes of the following sub-outcomes: "15-40% loss of species" and "the potential for Greenland ice sheet to begin melting irreversibly, accelerating sea level rise and committing world to an eventual 7m sea level rise" (Stern 67). The table looks as follows:

---

[12] For more information about calculation of utilities and probabilities see: Papineau, David. *Philosophical Devices: Proofs, Probabilities, Possibilities, and Sets*. Oxford, England: Oxford UP, 2012.
[13] This means that for example the use of fossil fuel is taxed to make the user pay for climate change, there are many ideas about how to implement this and whether everyone must be taxed with the same rate or not.

	Ev_1	Ev_2
Nothing	A	B
Carbon pricing	C	D

Table 3

In this example I just want to show that there is no difference in applying the method to fruit baskets and climate change. I will go into more detail about the application of the method on climate change later in this paper. For now I will make many assumptions, for example that I know how to value different climate change outcomes and sub-outcomes. Suppose that I am able to come up with the following scale C > D > A > B. Then I assign 1 to C and 0 to B and I propose bets in order to find the utilities of D and A. Once I have found those, the probabilities of events E_1 and E_2 should be inferred, by again proposing bets. Then one is able to calculate the expected utility for the two options available, namely doing nothing and introducing carbon pricing. Last, one can make the best choice by opting for the option with the highest expected utility.

5. Axioms

There are a few axioms to which one should adhere in order for the above method to work. I will discuss the three most important axioms. I will use the climate change example from above to illustrate what axioms are needed for. First of all what one necessarily should be able to do is to order all the outcomes from best to worst. *Ordering* consists, among other things, in *completeness* and *transitivity*. *Completeness* means that one should be able to compare any two outcomes, so that one outcome A should be better than, worse than or equal to outcome B. *Transitivity* means that if C is better than D, and D is better than A, then C should also be better than A. The axiom of *ordering* ensures that one is able to rank all the outcomes on an ordinal scale, as in the example above: C > D > A > B.

To be able to assign utilities to outcomes one should satisfy *continuity*. *Continuity* means that for any outcome there should be a gamble between a more and a less preferred outcome to which the outcome is indifferent. If we take the ordering as it was above, then for every outcome A, B, C and D there is a gamble between two other outcomes to which this outcome is indifferent.

It might be hard to find the outcome to which one is indifferent. It may not be possible when there are just four outcomes available, one needs more outcomes in order to satisfy this axiom.

The last axiom one should adhere to is *independence*. This means that if C > D, then [C (p) +A (1-*p*)] > [D (*p*) + A (1-*p*)]. This means that if C is better than D there cannot be a third outcome A that can change this. It is not possible that because of the introduction of an irrelevant alternative A, [D or A] will be better than [C or A] with identical probabilities. An irrelevant alternative is an outcome that cannot change the utilities of other outcomes and also does not change the interval between two other outcomes. All outcomes should be assigned independently. Furthermore, *independence* also means that the gamble [C 1-*p*; B *p*] should only be preferred to the gamble [C 1-*q*; B *q*] if *p* > *q*. Moreover *independence* implies that it should not matter how probable an outcome is, for assigning the utility to an outcome.

Ordering and *continuity* can both possibly be satisfied. What could be problematic about *ordering* and *continuity* is that outcomes are sometimes hard to value due to their complexity or due to the lack of information. Especially in regard to climate change it will be very hard to value outcomes accurately and precisely. For example it is hard to compare the value of protecting nature with the value of human well-being. Expected utility theory does not prescribe how to value. One might need a theory to value outcomes. Broome writes that in the case of climate change there are different kinds of value at stake and that they might be incommensurable ("Valuing Policies" 11). Again, I will not go into detail about how to value outcomes.

If these two axioms are satisfied, which might not be possible, but for now we assume that we can, then still one other axiom should be satisfied, namely the axiom *independence*. In the next section I will discuss this axiom in more detail and I will argue that satisfying this axiom might problematic in the context of climate change policy.

6. Independence and the Allais Paradox

I will explain why the *independence* axiom is problematic in the case of expected utility theory. In order to explain this I will first discuss the paradox described by Maurice Allais. Allais criticized the plausibility of the *independence* axiom in 1953 in “Le Comportement de l’Homme

Rationnel devant le Risque: Critique des Postulats et Axiomes de l'Ecole Américaine". I will explain the paradox by the use of an example about betting with money as the price.

Suppose that we ask someone to choose between the two options available for our subject, namely option A and option B. Furthermore suppose that there are three possible events with a possibility of occurring, namely event p, event q and event r. There are six possible outcomes. In Table 4, the information is shown and the dollars are in millions.

	Event p Pr (0,01)	Event q Pr (0,1)	Event r Pr (0,89)
Option A	$ 0	$5	$1
Option B	$ 1	$ 1	$ 1

Table 4

Option A is a lottery [$5 million Pr (0,1); $1 million Pr (0,89); $0 million Pr(0,01)] and option B is $1 million for sure. The subject thinks about these options for a while and comes to the conclusion that $1 million for sure is the better option[14]. The outcome of $5 million would of course be the best outcome, however the chance of getting this outcome in option A is not very high. Furthermore, if the subject chooses to opt for option A, there is chance that he will not get anything. The chance of getting nothing is very small, namely 0,01, however it would be very disappointing if this would be the result, since he could have had $1 million for sure. So, the subject concludes that option B is the better option for him.

Suppose another choice is offered to the subject. The choice is between option C and option D. The same probabilities hold for the events p, q and r. Take a look at Table 5, the dollars are in millions.

	Event p Pr (0,01)	Event q Pr (0,1)	Event r Pr (0,89)
Option C	$0	$5	$0
Option D	$ 1	$ 1	$ 0

Table 5

[14] See Kahneman, *Thinking Fast and Slow,* p. 281 for more elaboration on why most people rank options in this way.

Option C is a lottery [$5million Pr (0,1); $0 Pr (0,9)] and option D is a lottery [$1 million Pr (0,11); $0 million Pr (0,89)]. Again the subject takes some time to think about the options. He thinks the following: the chance of getting $5 million in option C is almost the same as getting $1 million in option D'. The subject prefers $5 million over $1 million. The outcome of $1 million is just 1 percent less likely to happen than $5 million. The subject chooses to opt for option C.

This is how many people will choose between the above options. The problem is that all these people violate the *independence* axiom. This is the case since the difference in valuing compound lotteries can only be in the different elements of the lotteries. Therefore, between option A and B one might only look at the outcomes $0 million Pr (0,01) and $5 million Pr (0,1) for option A and $1 million Pr (0,01) and $1 million Pr (0,1) for option B. To make this point more clear, take a look at Table 6.

	Event p Pr (0,01)	Event q Pr (0,1)	Event r Pr (0,89)
Option A	$ 0	$ 5	$ 1
Option B	$ 1	$ 1	$ 1
Option C	$ 0	$ 5	$ 0
Option D	$ 1	$ 1	$ 0

Table 6

The difference between option A and B are in the events of p and q. In the event of r, option A and B are the same, namely $1 million. The same holds between the options C and D, in the event of r the outcomes are the same in options C and D. The difference between option C and D should be under events p and q.

If it is the case that option B > option A, meaning:

[$1 million Pr (0,11)] > [$0 Pr (0,01); $5 million Pr (0,1)]

Then it cannot be the case that option C is better than option D, since that would mean that:

[$0 Pr (0,01); $5 million Pr (0,1)] > [$1 million Pr (0,11)]

It cannot be the case that it is true that the one outcome is better than the other outcome and that it is also the case that the opposite of this is true. If Q > Z it cannot also be the case that Z > Q. The same options under event p and q, with the same probabilities in both cases, cannot generate this difference in choice. Therefore, it must be the case that by taking into account the outcomes under event r in option A and B and also in option C and D, the value of the outcomes changes in the events of p and q. The value of the outcomes in event r, although being the same in both options does change the option chosen. The outcomes in event r should be irrelevant alternatives, according to the *independence* axiom, as was the case above with outcome A as explained in the section about axioms. The value of the outcomes under event r should be independent from the value of the outcomes under the events p and q. The subject who prefers B over A and C over D violates the *independence* axiom.

In the above case the probabilities were known. The violation of the *independence* axiom is not only the case when the probabilities are known, as in the example above. The violation of *independence* can also be the case when the probabilities are not known, as in case of climate change. In the example above the outcomes of one event were not independent of the outcomes of another event. If we are unsure about the probabilities of events then still *independence* can be violated. With 'being unsure' about probabilities I do not mean that someone does know nothing about the probability of an event. In most cases one does have an idea about whether something is probable, improbable or highly probable. Take a look at Table 7, in which the probabilities of the events are not known exactly.

	Event p	Event q	Event r
Option A	$ 0	$5	$1
Option B	$ 1	$ 1	$ 1
Option C	$0	$5	$0
Option D	$ 1	$ 1	$ 0

Table 7

Again, if the subject prefers option B over option A and option C over option D, then this subject probably has a different judgement about the probability of event p in the two different choices available. In the first case, the choice between A and B, the subject judges the event p probable to the extent that he opts for option B instead of option A. If the subject would have thought that p is very unlikely then he might have taken the gamble and would have chosen for option A, since in options A there is the probability of getting $5 million. However, since in option A there is a probability of not getting anything, while in option B $1 million will be won for sure, one opts for option B.

Now the same should be the case for the choice between option C and D. Since there is a chance of not winning anything in option C in two events (namely p and r), and there is just one event in which one does not win anything in option D (namely event r), one could opt for option D, if one thinks that this is a safe option. However this is not what most people choose. Now, since event r is expected to be most likely, and therefore the chance of not getting anything are already high in both cases, one chooses to take the risk of not getting anything and opts for option C. In this case the probability of event p is not judged to be very probable, because if it was judged to be probable, then one would have opted for option D instead.

So, since there is this third event r, with different outcomes between option A and B, and options C and D, there is a difference not only in assessing the utilities of outcomes, but also in assessing the probabilities of the events. For the first choice, between option A and B, the event p is judged to be quite probable. For the second choice, between option C and D, event p is judged to be not quite probable. So, if this is the case, one should see that the utilities of the different outcomes are not independent of the assessment of the probabilities, in the sense that one and the same event is judged to have different probabilities. So, our judgment about the probability of an event depends on what the alternative outcomes are.

What also could be taken from this experiment, emphasized by Kahneman and Tversky, is that if it is highly probable that one does not win anything one is more risk seeking and if it is highly probable that one will win or when there are many ‘good’ options available, then one is risk aversive[15]. Furthermore, Kahneman and Tversky write that the value assessment of an outcome (of, for example $ 0) can include the disvalue or value of what could have been the case.

---

[15] Read Kahneman and Tversky “Prospect Theroy: An Analysis of Decision under Risk” and Kahneman, Daniel. *Thinking, Fast and Slow*.

If in the case above one gets the $ 0, this is not to be seen as a value neutral outcome. The value of $ 0 includes in a way the disappointment of not getting the $ 1 million for sure. The disappointment of getting $ 0 is the result of this outcome being part of option A which includes other better outcomes; $ 1 million and $ 5 million, which the subject could have attained. This means that the value of the outcome $ 0 is not determined by the value of the outcome alone. It means that there is some 'interaction' between the outcomes. If the outcome should be *independent*, then we should ignore the other outcomes completely when valuing a particular outcome. However, this is not how we determine the value of an outcome. I will now turn to an example about climate change policy making in order to show that the *independence* axiom might not be satisfied in many climate change cases.

7. Climate Change Example

Let me again give an example about climate change, in which I will show that the *independence* axiom might not be satisfied in cases of climate change policy. The events are uncertain and outcomes are hard to value, when making a decision about what to do against climate change. I will extend the climate change example I used above. Suppose that there are three possible actions, namely doing nothing, replacing all cars powered by fossil fuels with electric cars or introducing carbon pricing. Suppose that there are four possible events, namely it will warm 2, 3, 4, or 5 degrees Celsius. All are expressed for a band of 1 degree Celsius, so 2 degree Celsius is the range between 1,5 and 2,5 degree Celsius. Suppose there are possible outcomes A, B, C, D, E, F, G, H and I. Take a look at Table 7:

	Ev +2 C	Ev +3 C	Ev +4 C	Ev +5 C
Nothing	A	D	G	J
Electric cars	B	E	H	K
Carbon pricing	C	F	I	L

Table 7

To remind ourselves, what should be done in order to make a choice according to expected utility theory is the following: first, all the outcomes should be ordered from best to worst. Secondly for every outcome there should be a gamble between two other outcomes to which one is indifferent. By proposing bets one assigns utilities to the outcomes. Thirdly the probabilities are inferred by again proposing bets. Once all the utilities and probabilities are assigned, one is able to calculate the total expected utility for the possible options available. One should choose the option with the highest expected utility, since is considered the best option.

The outcomes in the table are compound outcomes, this means that for example outcome A consists of different sub-outcomes. The Stern Review lists the sub-outcomes of outcome A, two of these sub-outcomes are:

(1) “15-40% loss of species” (Stern 67).
(2) “the potential for Greenland ice sheet to begin melting irreversibly, accelerating sea level rise and committing world to an eventual 7m sea level rise”(Stern 67).

These sub-outcomes should be taken together with the other sub-outcomes into the total outcome A and then outcome A needs to be assigned a utility. To only value the loss of 15 – 40% of species, might be a highly difficult task. Although this might be problematic, suppose for now that we can, suppose that we have enough information and a good theory for valuing very different sub-outcomes. Suppose that we can value the compound outcomes in an accurate and precise way. Furthermore, suppose that we can compare every compound outcome with every other compound outcome and that we are able to come up with an ordinal utility scale looking as follows:

$$A > B > C > E > F > D > I > H > L > G > K > J$$

Once this ordinal utility scale is created the axiom of *ordering* and with that of *completeness* and *transitivity* is satisfied.

One must assign utilities to all the possible outcomes. It does not matter what utility one assigns to U (A) or U (J) as long as U (A) > U (J). All the outcomes of climate change are negatively phrased. Outcomes are about loss, deaths, non-voluntary migration etc. This means that although outcome A is our best option, still it is not a positive one. Best in the case of climate change is least bad. Option A is the best option available given the other options we have available. This also means that the 0 utility that will be assigned to our worst option J, is not a neutral 0. Wrongly, the utility 0 can be seen as neutral or as an outcome in which nothing is won or lost, remember the case with 0 dollar above. We have to keep in mind that 0 dollar or 0 utility is not neutral. The 0 utility that will be assigned to J is highly negative. One could for example choose to assign -1 to U (A) and -100 to U (J). However, I will just follow the procedure of assigning utilities as explained above.

Once we have created the ordinal utility scale, we are able to do the ‘trick’ of assigning utilities. We assign utility 1 to our best option A and 0 to our worst option J. We will continue by proposing bets to find out the utilities of the other outcomes.

First, reconsider the Allais paradox. This paradox is an example of a broader phenomenon in which what is considered good about an option cannot be reduced to the goodness or badness of individual outcomes of this option. In order for me to be able to argue that expected utility theory is an improper way of making a choice under uncertainty about climate change, I need to show or at least make plausible that there is some kind of ‘interaction’ between the possible individual outcomes that partly determines the badness rate of the option. In other words, it needs to be plausible that the option is sometimes more than its components combined. I will show that this will be the case in some climate change cases.[16]

8. Independence and Climate Change Policy Making

Suppose, to make the example more clear, that we disregard the first option of doing nothing. We have two options left: the electric car option and the carbon pricing option. I will show, by the use of an idea by Peter Diamond, that sometimes it will be impossible to reduce the value of an option to the sum of the value of the individual outcomes of this option (Diamond 766). This means that the badness rate of a policy is sometimes not just the badness rates of the different

[16] For the idea of the example below I heavily depend on Peter Diamond’s idea that the total value of an option is sometimes not the sum of the value of the individual outcomes, which he describes in "Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons: Comment." *Journal of Political Economy* Oct 75.5 (1967): 765-66.

possible outcomes combined, but that there is some kind of ‘interaction’ between the different possible outcomes that partly determines the badness rate of the option.

Still, I will keep the example simple, but to make this example more realistically regarding climate change, let me again introduce two outcomes from the Stern Review. Suppose that in the case of the electric car policy (A) “15-40 % of species facing extinction.”(Stern 67). Suppose that in the case of carbon pricing (B) “1-3 million more people die from malnutrition” (Stern 67).

Option A and option B are both bad. Suppose that one cannot choose between these options and therefore one decides that option A and option B are equally bad. This means A ~ B. A third option C is offered. Option C is the possibility to make a choice between A and B by tossing a fair coin. In this case both option A and B have an equal chance of occurring. Option C seems to be the better option, since it is a fairer option. Therefore: option A is equally good as option B and option C is better than option A as well as better than option B. So, A ~ B < C.

In this example there are two events, namely head or tails. There are three options, namely A, B and C. There are two outcomes per option.

	Head	Tails
A: electric cars	15-40 % of species facing extinction	15-40 % of species facing extinction
B: carbon pricing	1-3 million more people die from malnutrition	1-3 million more people die from malnutrition
C: toss a fair coin	15-40 % of species facing extinction	1-3 million more people die from malnutrition

Table 8

Option A is equally bad as option B. However, option C > option A, which means;

> [15-40 % of species facing extinction if ‘head’ ; 1-3 million more people die from malnutrition is ‘tails’] > [15-40 % of species facing extinction if ‘head’ ; 15-40 % of species facing extinction if ‘tails’]

There is no difference between option C and option A when ‘head’ occurs, since in both cases ‘15-40 % of species face extinction’. It follows that if option C is better than option A, then the difference between the both is because of the different outcomes under ‘tails’. When ‘tails’ occurs in the case of C ‘1-3 million more people die from malnutrition’ and in the case of A ‘15-40 % of species face extinction’. In short, it must be the case that ‘1-3 million more people die from malnutrition’ is better than ‘15-40 % of species facing extinction’. If C >A, the forgoing would follow from *independence*. However, problematically this would also mean that B > A. This is problematic since we just decided that those are equally bad.

If we want to argue that C is indeed better than A and B, then we cannot adhere to the *independence* axiom. The point of this example is to show that some of the options cannot be reduced to the individual outcomes alone. If we would have only taken the individual outcomes in account, we would have had a different result. Suppose that we had followed the expected utility theory and we had assigned U = 0,5 to ‘15-40 % of species facing extinction’ as well as to ‘1-3 million more people die from malnutrition’, since we were indifferent between those outcomes. If we would have done so, we would not have been able to tell the difference between A, B and C. The table would have been as follows:

	Head Pr = 0,5	Tails Pr = 0,5	Expected Utility
A: electric cars	U = 0,5	U = 0,5	0,5
B: carbon pricing	U= 0,5	U =0,5	0,5
C: a fair toss	U= 0,5	U =0,5	0,5

Table 9

The way in which the utilities are assigned in the table above is not how we tend to assign them, however it is how we would assign values if we intent not to violate the independence axiom. The problem with the expected utility theory is that we assume that the value of the option is the total value of the individual outcomes. In *Weighing goods* Broome writes that Peter Diamonds claim that the total value of an option is sometimes not the sum of the individual outcomes is very well plausible, when determining subjective goodness (111). Broome also writes that this is not the case when determining objective goodness (*Weighing goods* 115). However, as shown above it could be the case that badness of the option, as in option C above, is not only the badness of the different outcomes.

The choice above between the different policies, is a choice about which policy is better according to expected utility theory. However, what makes C the better option is a consideration of fairness. What is shown in the example above is that this consideration about fairness cannot be reduced to the combination of different individual outcomes. In other words, just by examining the value of the outcomes alone one will not be able to show that option C is a better option. It is probably also the case that an option is sometimes more than its components alone, since the value of an option also includes the value of what could have otherwise happened. There is 'an interaction' between the outcomes. Since when making a decision about climate change one will often be confronted with these kind of considerations, as fairness, many of the climate change problems will not satisfy the *independence* axiom.

From the Allais paradox we learned that sometimes an outcome is valued differently because of the other outcomes available. From the example above we learned that it might be the case that the 'sum' of the outcomes is not always the total value of an option. It follows that in practice expected utility theory might be unfavourable, since the way in which expected utility theory requires us to value outcomes and options is not the way in which we tend value outcomes and options. Theoretical economist and philosophers who think about decisions under uncertainty, see the problem discussed above, however it seems to be the case that the practical implications of this problem have not yet reached these 'climate-mathematicians' as Stern and others.

9. Interference

There is one further point that I shortly want to discuss. There might be another problematic point, when applying expected utility theory to climate change. This point is related to the independence critique, however it is slightly more specific for the climate change example I use in which the events are the possible degrees of the earth warming up. Suppose that a policymaker sits down and starts to think about what to do against climate change. This policymaker is told about expected utility theory and he tries to apply expected utility theory to the climate change problem in order to make a decision about what to do. This policymaker does this in order to find out what the best action is. This best action hopefully will be something that reduces the chances of a catastrophe, or at least reduces the chance of very unfavorable situations.

The problem might be that the goal or the end of the whole project of applying expected utility theory to the climate change problem is to minimize the chance of the world warming up.

When battling climate change the purpose is to try to reduce the chances of the earth warming up 'too much'. It is the case that one's goal is to try to change the probabilities of the events. It might be the case that by choosing for one option the probability of an event will be different. The probabilities of events are interfered by the available outcomes.

For example, if one decides to take action and introduces carbon pricing, it might be the case that the chances of the world warming up with four and five degrees Celsius will be smaller than when we would have decided to do nothing. The question is whether the probabilities of the events are the same in all the options available. The probability of the event of 5 degrees Celsius warming might be smaller in the case in which one does introduce carbon pricing instead of doing nothing. This might be problematic since the options should not interfere with the probabilities of the events, according to expected utility theory.

10. Conclusion

Expected utility theory is an established method for making rational choices under uncertainty. John Broome proposes to use expected utility theory to make a choice on how to act on the climate change issue. He is mainly concerned with the problem of valuing and does not propose a method for assigning probabilities. Therefore I used the ‘probability as partial belief’ theory by Frank Ramsey, to be able to infer subjective probabilities. The betting method, as Ramsey proposes is useful for determining utilities, but since most of us are inconsistent in assigning utilities it is not a favourable method for making climate change policy.

If the method is used for a rough estimation about what action is best to perform in the case of an unimportant decision, for example whether is it better to take an umbrella with you when going outside or not, then there is no problem. Although it seems a good method for ordering one’s thoughts schematically for an unimportant decision, it is not a proper or suitable method for making a decision about climate change. Even if one is able to satisfy the axioms of *ordering* and *continuity,* I showed in this paper that expected utility theory is improper in its practical application.

First, by the use of the Allais paradox, I have made clear that when people apply expected utility theory, most people do not value outcomes consistently and most people do not assign probabilities consistently. The value of an outcome and the probability of an event are both partly determined by the values of the other outcomes available. Secondly, by the use of another example I showed that the Allais paradox is part of a broader phenomenon, in which in some cases the value of an option is not the combined value of the individual outcomes. Furthermore, the badness rate of the different policy options are sometimes not only the sum of the different outcomes, there is some interaction between the outcomes that contribute to the total badness of the policy options. A consideration as fairness cannot be reduced to the sum of the individual outcomes. When applying expected utility theory, however, one assesses every outcome on its own and one assumes that the sum of the individual outcomes is the total value of an option. Furthermore, since the goal of applying expected utility theory is to reduce the chance of the earth warming ‘too much’, it might be the case that the possible actions do interfere with the probabilities of the events.

I conclude, in opposition to John Broome, that the expected utility theory is an improper method for decision-making about climate change policy.

Bibliography

Allais, Maurice, 1953, “Le Comportement de l’Homme Rationnel devant le Risque: Critique des Postulats et Axiomes de l’Ecole Américaine”, *Econometrica*, 21: 503–546.

Broome, John. “The Most Important Thing About Climate Change.” *Public Policy: Why Ethics Matters.* Canberra: ANU E Press, 2011. 101-116.

Broome, J. "Valuing Policies in Response to Climate Change: Some Ethical Issues." *The Economics of Climate Change: The Stern Review*. Cambridge, UK: Cambridge UP, 2007.

Broome, John. *Weighing Goods: Equality, Uncertainty, and Time*. Cambridge, MA: Basil Blackwell, 1991.

Broome, John. *Weighing Lives*. Oxford: Oxford UP, 2004.

Diamond, Peter. "Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons: Comment." Journal of Political Economy Oct 75.5 (1967): 765-66.

Kahneman, Daniel. *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux, 2011.

Kahneman, Daniel, and Amos Tversky. "Prospect Theory: An Analysis of Decision under Risk." *Econom etrica* 47.2 (1979): 263-98.

Kolmogorov, A. N. *Foundations of the Theory of Probability*. New York: Chelsea Pub., 1956.

Papineau, David. *Philosophical Devices: Proofs, Probabilities, Possibilities, and Sets*. Oxford, England: Oxford UP, 2012.

Ramsey, Frank Plumpton, and D. H. Mellor. *Philosophical Papers*. Cambridge: Cambridge UP, 1990.

Sahlin, Nils-Eric. *The Philosophy of F.P. Ramsey*. Cambridge: Cambridge UP, 1990.

Stern, N. H. *The Economics of Climate Change: The Stern Review*. Cambridge: Cambridge UP, 2006.