Time Discounting Functions as Measures of Radical Uncertainty about the Future?

Conrad Heilmann

Erasmus Institute for Philosophy and Economics (EIPE)
Faculty of Philosophy
Erasmus University Rotterdam
P.O. Box 1738
3000 DR Rotterdam
Email: heilmann@fwb.eur.nl

Abstract

This paper investigates whether time discounting functions can be measures of radical uncertainty about the future. Time discounting functions assign declining values to future time points. This property seems to closely correspond to our intuition that the far future is more uncertain than both the near future and the present. But can the framework of time discounting be used to provide foundations for a measure of radical uncertainty about the future? This paper argues that the standard discounted utility and social discount rate models of time discounting do not lend themselves to do so. While the contribution of this paper is largely negative, it also offers sketches of an alternative way to provide a time discounting characterisation of radical uncertainty about the future via representation frameworks that can be interpreted as providing exclusively epistemic foundations of time discounting.

1 Introduction

The future is uncertain to us in the most fundamental sense: time breeds uncertainty. This proposition captures a deep, twofold intuition concerning uncertainty about the future. One aspect of this intuition is that the future just *is* radically uncertain in an epistemic sense. The other aspect is that our uncertainty about the future is increasing with longer horizons: the further away, the more uncertain the future is to us. In this paper, I take this twofold intuition as a starting point and ask the following question. How can we move beyond such assertions of intuition and give an account of radical uncertainty about the future?

In order to examine one possible answer to this question, I take a closer look at the literature on time discounting (reviewed in Frederick *et al.* (2002)) and ask whether it provides a framework for capturing radical uncertainty about the future. Time discounting functions assign values to time points, with the present assigned the unit weight. Future time points are typically assigned values between zero and the unit weight, and decreasing with time. On the face of it, these properties seem to be well suited to provide a measure of radical uncertainty about the future. Yet, in order for time functions of the kind just described to provide such a measure, it would have to be possible to derive it in the representation framework of time discounting. Thus, I ask in this paper whether the representation frameworks that yield such time discounting functions stand a chance to provide sound foundations of the idea of radical uncertainty about the future. If the latter would be possible, it should not be taken as an indication that we have developed a theory of 'discounting the future for uncertainty' or a justification of reducing the value of future events because of uncertainty. Rather, deriving a discounting function as a measure or representation of radical uncertainty would just mean that it captures the phenomenon in the right way – how such a uncertainty discounting function is to be used in decision-making is a different question altogether.

It turns out that even deriving a discounting function as a measure of radical uncertainty about the future is by no means straightforward. The main contribution of the paper is thus to show that standard representation frameworks of time discounting cannot be interpreted in this way. Standard frameworks in this literature build on the so-called discounted utility (DU-)model of time discounting (e.g. Samuelson (1937), Koopmans (1960), Lancaster (1963), and Fishburn and Rubinstein (1982)). On this account, a discounted utility function represents the impatient time preferences of agents. The discounted utility function, in turn, implies a time discounting function of the kind mentioned earlier. Now, while the latter could be interpreted with a notion of radical uncertainty about the future, I will demonstrate that it is impossible to interpret the axioms of time preference in the DU-model with any such notion. That is, while we can adopt the concept of radical uncertainty about the future as an ad hoc interpretation of the implied time discounting function generated by the DU-model, we cannot interpret its foundations – namely time preference – in this way. Likewise, the framework of the so-called social discount rate, initially proposed by Ramsey (1928), cannot be linked to an interpretation of radical uncertainty about the future.

The contribution of the paper is mainly negative. The two arguments suggest that the conceptual link between uncertainty and time discounting is weaker than is commonly assumed. That is to say, the idea of 'discounting for uncertainty' lacks foundations in standard accounts of time discounting. We should thus be much more careful in suggesting that discounting offers a framework that captures uncertainty, or that we can even discount future outcomes due to uncertainty. On the constructive side, I offer some reasons as to why standard frameworks of time discounting fail in the way described, as well as some suggestions about how non-standard and generalised frameworks could offer the right kind of foundations. These accounts operate (in part) outside the time preference framework, such as Ok and Masatlioglu (2007), and Heilmann (2008). Therefore, they are not linked to value assessments in the way standard accounts are, which makes an exclusively epistemic interpretation of the derived discounting measures much more plausible. I will briefly introduce the main elements of these frameworks and sketch how it could be possible to interpret them as capturing radical uncertainty about the future.

The paper proceeds as follows. Section 2 discusses the problem of radical uncertainty about the future in greater detail. Section 3 introduces the DU-model of time discounting, and the approach of the social discount rate and shows that these approaches cannot provide foundations for an account of radical uncertainty about the future. Section 4 sketches ways in which alternative representations of time discounting that operate outside the time preference and DU-model frameworks can provide the foundations of a time discounting account of radical uncertainty about the future. Section 5 briefly concludes.

2 Radical Uncertainty about the Future

Prospects that extend far into the future are, by and large, more uncertain than those that play out over a short horizon. For example, I am less certain about what I will be doing in twenty years time than about what I will be doing tomorrow. Why is this the case? Let me offer some initial speculations. It seems to be the case that there is a greater continuity of 'how things are' in the near future when compared to the far future. The status quo provides us with more indications about how things will be tomorrow than how they will be next year. Put more abstractly, the far future is associated with more complex causal branching, which makes us more uncertain about it. Moreover, long horizons might also lead us to factor in highly unlikely events much more than in a short horizon: if I only plan about tomorrow, I might not entertain the possibility of an earthquake in a given region, but If I plan for several decades, I might be more inclined to do so. In such an intuitive and general sense, uncertainty seems to increase with longer horizons, and might well be radical for very long ones.

All this seems plausible enough. Yet, we can easily imagine scenarios in which these intuitions do not seem to hold. For instance, think about a person who is made redundant two years before their retirement age. Surely, there is greater continuity of monthly income when retirement benefits kick in than before that happens. I might also be able to say more clearly what I will be doing next Sunday from the status quo of my Sunday routine that I went through today, with more uncertainty about what happens during the week leading up to that Sunday. The future could also become less complex in a project in which the number of partners involved diminishes over its duration. Finally, in the wake of a catastrophe, such as a Tsunami, highly unlikely events such as a nuclear disaster will be more likely than over a longer horizon. These cases seem to suggest that the far future is not always more uncertain than the near future.

Now, it is easy to see that in the above cases, the lower level of uncertainty about the future just concerned some specific causal and spatiotemporal path about which some information was available, which resulted in less uncertainty about a specific type of event in the future. This kind of uncertainty may well be assessed within standard decision-theoretic frameworks (such as Savage (1972), Jeffrey (1983), Joyce (1999), or Bradley (2007)).

However, the general point concerning radical uncertainty about the future still stands. Reduced uncertainty about some future events – and less uncertainty about some aspect of the far future when compared to the near future – is compatible with there also being increasing radical uncertainty about the future in general. In all of the above scenarios, there are many uncertain aspects about the future, with most of them being more uncertain for the far future. For instance, the redundant person will typically be more uncertain about a multitude of other aspects in the far future than in the near future, such as their personal health, the situation of their family, the stability of the political system, and so on.

We are thus dealing with a much more general problem. We are not concerned with the kind of 'local' uncertainty inherent in a defined causal and spatiotemporal path or specific event, which we may be able to capture by using the standard tool of probabilistic assessment in decision theory. Rather, radical uncertainty about the future concerns uncertainty about all sorts of background factors that are not explicitly modelled in such scenarios. The point is this. By and large, we think that the future is uncertain, and that the further away events and processes, the more uncertain they are. Beyond such a general statement, it is very hard to say how and why exactly this kind of uncertainty about the future arises, what significance it has, and how we can capture it and factor it into our models of decision-making. That is to say, I take it that there is an intuition that beyond any specific aspects we have privileged information about, the future is radically uncertain, and the further away, the more uncertain it is.

The question, then, is the following: can we move beyond stating this kind of general intuition about radical uncertainty about the future? If not, we would be stuck with the following answer: whenever we assess future prospects, we should simply do as best as we can to apply our standard probabilistic methods, so as to quantify uncertainty about the future as far as possible. Indeed, this is what seems to be prescribed by standard decision-theoretic frameworks. However, the longer the horizons, the more untenable it is to capture an agent's uncertainty about the future in this way.

To see why, imagine we want to specify possible paths that result from taking alternative actions today. Surely, our ability to specify 'full world histories' and assess them probabilistically, as Savage (1972) has it, will break down at some point. For instance, individuals will make decisions about longterm investments or savings that will have consequences that are spread out over several decades. Collectives will have to make decisions with consequences that span several hundred years, such as when deciding how to deal with the problem of climate change. In such contexts, uncertainty about the future may well be radical: long horizons can give rise to considerations that lie outside what we can reasonably include in any 'objective' probability framework or subjective probability assessments. We will have to make choices what kinds of possible paths to model, and what kinds of factors to include in their specification that will leave out some important considerations. Can an individual agent reasonably be expected to give a subjective probability of whether she will still live in a liberal democracy in fifty years? Can we expect experts to assess the interplay of large-scale parameters in climate models over hundreds of years?

Now, some methods have been developed that deal with such problems, such as non-standard probability frameworks, including imprecise probabilities (Halpern (2001) provides an overview). The question I am pursuing in this paper is the following. Is there a way of capturing radical uncertainty about the future over and above of what we may be able to quantify in such standard and non-standard probabilistic frameworks? In other words, is there a chance that we can give foundations to the idea that there is general radical uncertainty about the future? The next section reviews two families of existing time discounting approaches that could be candidates for capturing general radical uncertainty about the future and shows that they cannot provide foundations for general radical uncertainty about the future.

3 Time Discounting

3.1 The Intuition

In decision theory and microeconomic theory, time discounting offers a framework in which intertemporal aspects of prospects can be evaluated (Frederick *et al.*, 2002). The question of this paper is whether such time discounting approaches can capture radical uncertainty, that is: can we use time discounting methods in contexts where we want to take radical uncertainty about the future into account?

On the face of it, time discounting functions would be well suited to do so. Typically, a time discounting function is an assignment of weights to time points, such that the present is given full weight and thus not discounted at all, with future time points assigned declining values between zero and 1. If we have the intuition that uncertainty increases with time, then this type of function seems to be of the right kind to accommodate it. More formally, discounting functions $D : T \longrightarrow (0, 1]$ are decreasing mappings from a set of time points T to the real interval $(0, 1]$, such that $D(0) = 1$. Often, this is achieved by an assignment $D(t) = \delta^t$, where a discounting factor δ is used to generate a discounting factor for every time point.[1] Assessments of this kind are frequently integrated into decision theory, most commonly by directly weighting utilities with it. Of course, the exact motivation for the use of discounting functions will depend on how it was derived. For the task at hand, we need to investigate how 'uncertainty time discounting functions' can be derived in order to decide whether they are well suited to capture the intuition that time breeds uncertainty.

At first glance, the idea of using this framework to capture radical uncertainty about the future seems to be simple enough. The discounting function and the discounting factor would, in short, represent the intuition that there is more uncertainty about the far future than the near future, and none about the present. That is, $D(0) = 1$ would capture the intuition that there is no uncertainty about the future with regards to the present. The decreasing values for future time points, such as, for instance, $D(1) = .99$, $D(2) = .98$, $D(3) = .97$, and so on, express increasing uncertainty about the future. In the following sections, I will analyse whether this stance can be maintained and substantiated when looking at how time discounting functions are commonly constructed.

I will ask the following questions. Can the discounting functions really be taken as an expression of radical uncertainty about the future? Are there representation frameworks that can be interpreted such that they provide a proper foundations for such a view? Will a representation of radical uncertainty thus understood be sound, and can it be applied in conjunction with standard probability assessments that are part of decision-theoretic approaches? Answering these questions is greatly complicated by the fact that time discounting and its foundations are notoriously ambiguous areas of research in themselves. Therefore, we will distinguish between two different approaches on time discounting, and will ask whether they can provide detailed and comprehensive foundations for the intuition that radical uncertainty about the future can be captured by timed discounting functions.

In the following, I will first ask whether the time preference and discounted utility framework (or DU-model) can be interpreted as capturing radical uncertainty. Secondly, I will review the proposals around the problem of risk and uncertainty that have been made in the behavioural economics literature and the literature on the social discount rate.

3.2 Time Preferences and Discounted Utility (DU-model)

In the DU-model, the discounting factor δ that is used to generate discounting factors for each time point via the discounting function $D(t) = \delta^t$ is a representation of time preferences.

In this context, it has been proposed to interpret the discounting factor as the probability of survival. More precisely, $\delta^t = p(\text{survival in period } t + 1)$. This, is for example explicitly mentioned as a possible conceptual motivation for time discounting of utility in the influential graduate microeconomic theory textbook by Mas-Colell *et al.* (1995, 734). Before returning to the question whether this interpretation can indeed be linked to the time preference and discounted utility framework, we will introduce the latter formally.

The canonical interpretation of time discounting in economics depicts the discounting factor δ in a discounting function as a representation of pure and positive time preference. The concept of *pure and positive time preference* implies that an agent has a preference for utility at earlier points in time over later points in time, in addition to her preferences over prospects. More specifically, *pure* time preference refers to the fact that the betterness which is expressed by time preference is associated with distance in time. *Positive* time preference refers to the fact that time preferences capture the idea that 'earlier' is better than 'later', with the present being the earliest point in time considered such that positive outcomes at it are preferred to all later points in time.

More formally, in the DU-model and its variants, the new concept of time preference is introduced to deal with the subjective evaluation of intertemporal prospects. Instead of introducing a valuation on some set of outcomes X which gives a utility function, time preferences compare prospects that are combinations of outcomes and times, i.e. the domain of preference becomes $X \times T$, and time preferences $\succsim_{TP}$ over this domain can be numerically represented as *discounted utility.* Fishburn and Rubinstein (1982) provide an axiomatisation of pure and positive time preference in difference structures that gives exponential discounting in similar spirit to standard representations by Samuelson (1937), Koopmans (1960) and Lancaster (1963). The formal framework in Fishburn and Rubinstein (1982) is based on an outcome-time structure $\langle X \times T, \succsim_{TP} \rangle$, where $\succsim_{TP}$ is a preference relation on pairs of time points and outcomes (note that the outcomes are already valued, i.e. the set of outcomes is a set of utilities). Fishburn and Rubinstein (1982, 680) use the standard conditions on rational preferences as a starting point and add further conditions to prove that the preference relation $\succsim_{TP}$ can be represented by a discounted utility function.

Outcome-time structure (Fishburn and Rubinstein, 1982, 680). X is a non-degenerate real interval, T is either a set of successive non-negative integers or an interval of non-negative numbers, and $0, 1 \in T$. For $\succsim_{TP}$ on $X \times T$, for all $x, y \in X$ and all $s, t \in T$, the following conditions hold:

Weak order. $\succsim_{TP}$ is a weak order on $X \times T$;

Monotonicity. If $x > y$ then $(x, t) > (y, t)$;

Time impatience. If $s < t$ then

(i) if $x > 0$ then $(x, s) \succ_{TP} (x, t)$,
(ii) if $x = 0$ then $(x, s) \sim_{TP} (x, t)$,
(iii) if $x < 0$ then $(x, t) \succ_{TP} (x, s)$;

Continuity. $\{(x, t) : (x, t) \succsim_{TP} (y, s)\}$ and $\{(x, t) : (y, s) \succsim_{TP} (x, t)\}$ are closed in the product topology on $X \times T$;

Stationarity. If $(x, t) \sim_{TP} (y, t + \mu)$ then $(x, s) \sim_{TP} (y, s + \mu)$.

The key conceptual assumptions in this framework are the conditions of time impatience and stationarity. The time impatience condition states that agents prefer to receive positive utility earlier and negative utility later (Fishburn and Rubinstein, 1982, 680). The stationarity condition asserts that indifference between two time-dependent outcomes depends only on the difference (μ) between the times and not on the actual time points $s, t \in T$ (Fishburn and Rubinstein, 1982, 681). Together, these constraints on the outcome-time structure $X \times T$ assure that time preferences can be represented by a discounted utility function.

Theorem (Fishburn and Rubinstein, 1982, 682). Suppose time preferences $\succsim_{TP}$ on an outcome-time structure $X \times T$ that satisfy the conditions of weak order, monotonicity, continuity, time impatience, and stationarity. Then, given any $0 < \delta < 1$, there is a continuous, increasing real-valued function u on X such that:

(i) for all (x, t), $(y, s) \in X \times T$, $(x, t) \succsim_{TP} (y, s)$ iff $\delta^t u(x) \geq \delta^s u(y)$,

(ii) for all $x \in X$, $u(x)$ is zero if $x = 0$, positive if $x > 0$, and negative if $x < 0$,

(iii) if T is an interval then u is unique (given δ) up to multiplication by positive constants on $\{x \in X : x > 0\}$ and on $\{x \in X : x < 0\}$.

Accordingly, upon assuming an outcome-time structure, time preferences are representable by a discounted utility function. The main advantage of this method is that one can motivate a discounting function in a representation framework, which yields proper foundations. Yet, for the task at hand we need to ask whether these foundations are compatible with an interpretation of fundamental or radical uncertainty about time.

It is hard to see how this would be possible. While the framework yields a discounting function that has the desired properties – which, in turn, correspond well enough to our intuitions concerning radical uncertainty about the future – it is implausible to interpret the concept of time preference with it. To see why, consider the so-called axiom of time impatience in the above framework. It consists of three parts, detailing how time impatience behaves with regards to outcomes that are evaluated (i) positively, (ii) neutrally, and (iii) negatively, assuming that time point s is earlier than time point t:

(i) outcomes that are evaluated positively (i.e. $x > 0$) are strictly preferred at an earlier time, i.e. $(x, s) \succ_{TP} (x, t)$.

(ii) outcomes that are evaluated neutrally (i.e. $x = 0$) induce indifference with regards to the time they occur, i.e. $(x, s) \sim_{TP} (x, t)$,

(iii) outcomes that are evaluated negatively (i.e. $x < 0$) are strictly preferred at a later time, i.e. $(x, t) \succ_{TP} (x, s)$.

Now, in order to interpret these requirements with radical uncertainty, the latter would need to coincide in some fashion with the time preferences. Yet, we said earlier that radical uncertainty with regards to the future is increasing, which would mean that it behaves regularly with regards to the time, i.e. for any time point s that is earlier than any time point t, we would assume that there is less radical uncertainty associated with s than with t. However, it is implausible that radical uncertainty would vary with value in the same way as was assumed about time impatience.

The main problem with motivating the axioms is that our intuitions about radical uncertainty are independent of the valuation. We have just seen that this does not go together with the axiom of time impatience, which distinguished between positive, neutral, and negative outcomes. In contrast, radical uncertainty about the future of the kind we are interested in here has nothing to do with the fact what kinds of valuations are in the picture: it is simply not the case that the future is less uncertain for outcomes with negative evaluation, or that there is more uncertainty when there are positive valuations involved. (For an account that seems to capture scenarios in which this is the case, see the next section.)

Furthermore, the idea of general radical uncertainty about the future – i.e. that independent of a given specified path there is just more uncertainty associated with the far future than the present or the near future – does not go well with singling out evaluations of single outcomes. As discussed in section 2, uncertainty about those kinds of specific scenarios might well be representable by standard decision-theoretic tools or non-standard probabilistic frameworks.

Therefore, the DU-model cannot be interpreted as providing a foundation for time discounting that is linked to radical uncertainty about the future. While the discounting function that is implied in the above framework can be linked to such an interpretation, the axioms in the representation theorem itself cannot. We need to look elsewhere.

3.3 Risk Aversion and the Social Discount Rate

The concept of the social discount rate is a different framework which is sometimes used to provide foundations for time discounting. Even though those foundations are not thoroughly axiomatic, they nevertheless provide – together with the DU-model – the canonical approach to time discounting.

The social discount rate goes back to the idea that a discounting factor δ can also be generated by a discount rate, such that the discounting factors for each time point are generated as follows: $D(t) = \left(\frac{1}{1+r}\right)^t$. Now, the social discount rate approach postulates that r is a composite concept that captures a number of different motivations for time discounting. Standardly, the social discount rate is given by the so-called 'Ramsey equation' as $r = \rho + \eta g$ where ρ is the (social) rate of time preference, η the elasticity of marginal utility of consumption or the coefficient of relative risk aversion, and g is the growth rate of consumption per capita. Ignoring rho and g, we can ask whether η can be interpreted as reflecting radical uncertainty about the future.

Here, Medvecky (2012) shows that it cannot – and for similar reasons that where given in the case of the axioms of the DU-model. Firstly, it is important to recognize that the standard interpretation of η in the social discount rate approach is that it is a coefficient of relative risk aversion. What this means is in fact that again, uncertainty is tied together with valuational concerns: 'According to proponents of discounting for uncertainty, uncertainty leads us to discount because we prefer more certain returns to less certain returns, and the present is more certain than the future.' (Medvecky, 2012, 4). This shows that the latter concern, i.e. the idea that the present is more certain than the future, is present in the risk aversion coefficient, but tied in with the consequences it is alleged to have for the valuation. Medvecky (2012) goes on to argue that in the case of uncertainty, we should discount positive outcomes (benefits) and negative outcomes (costs) separately, but that argument is irrelevant for the purposes of our discussion here.

As in time preference, since the different elements of the social discount rate are also linked to considerations of value, we cannot get around the fact that uncertainty works in 'both directions' – i.e. negative and positive value, but that the interpretations of the above frameworks rests on other concepts that do not. Standard frameworks of discounting can hence not serve as a proper foundation of the idea that discounting functions could reflect radical uncertainty about the future.

4 Towards Epistemic Foundations of Time Discounting

4.1 Value-Time Entanglements

We have seen that standard frameworks do not lend themselves to provide the desired foundations for a measure of radical uncertainty about the future. While showing that is the main message of this paper, I now offer a discussion of why the two approaches fail to do so. I will suggest that their main problem is the entanglement of time and value in the derivation of the time discounting function.

Consider again the target of any representation of time discounting, namely a decreasing function $D : T \longrightarrow (0, 1]$, such that $D(0) = 1$. What we aim at is an axiomatic framework in which we can derive such a function, and interpret the function as a numerical representation of radical uncertainty about the future. The two standard approaches do not directly derive such a function. Rather, the discounted utility approach starts from a standard preference account of value and introduce additional structure, such that a time discounting function is derived as an addition to the utility function. In similar vein, the social discount rate approach introduces different concepts in which time and value considerations have a joint influence on its derivation.

In the previous section, I have shown that the value dimension makes it impossible to interpret the standard frameworks with radical uncertainty about the future, as the latter is independent of the kinds of value considerations that are inherent in them. As per our initial discussion, we simply desire a measure of radical uncertainty about the future that captures the twofold intuition that the future is more uncertain than the present, and that the further away, the more radical uncertainty there is.

At this point, it is useful to remind ourselves of the important distinction between 'discounting for radical uncertainty about the future' on the one hand and the idea of 'providing a measure of radical uncertainty about the future'. It seems that the former goal is much more ambitious, as it aims at taking an evaluation and then weighting it with a factor that is related to radical uncertainty about the future. We can suspect that the standard approaches, with their entanglements of time and value, are aimed at this much more ambitious goal. Yet, as we have seen, the derivation of the discounting function does not lend itself to such an interpretation in either framework. The question is thus whether a more modest goal, namely viewing discounting functions as measures of radical uncertainty about the future, can be achieved.

The disadvantages of the two aforementioned approaches thus leads us to the question whether there are other ways to derive a time discounting function that avoid their problems. That is, we are facing the challenge of deriving a measure of radical uncertainty about the future that has the form of a time discounting function, but is not derived on or in relation to the value domain, and therefore motivated exclusively epistemically. Indeed, the best prospect of achieving this goal is by introducing a radical uncertainty ordering of any future event, and impose the condition that the ordering has to correspond to the time index, such that the values of the time discounting functions are decreasing of that event. How can we do so? I will now turn to evaluating two kinds of proposals that can be interpreted in this way, and contend that they at least possibly capture the idea of radical uncertainty about the future.

4.2 Delay Representations

In a first step, we ask whether we can interpret the representation theorem by Ok and Masatlioglu (2007) can be interpreted with radical uncertainty about the future. In this alternative axiomatisation of time discounting, the representation was (partly) achieved by a 'delay ordering' on a time domain. This representation is compatible with the kinds of approaches to delay or interval discounting given by, amongst others, Ainslie (1992, 2001) and Scholten and Read (2006).

The approach assumes a set of time points T which are ordered according to the perception of the delay by an agent. That is, pairs of $T \times T$ are ordered according to how relevant the agent perceives the delay. In Ok and Masatlioglu (2007), it is assumed that the ordering of the time points reflects the perception of delay by an agent. That is, in an elicitation exercise, we can ask agents how much delay they associate with each externally given time point that is in T. What this can yield are patterns of horizon perception, for instance, a delay of one or two days may make a difference in the short horizon, but not so much in the long run.

Ok and Masatlioglu (2007) use such a delay ordering in conjunction with other motivations that may play a role in time discounting to provide foundations for time discounting functions, but it would also be possible to isolate this part of the foundations. That is, we could use an ordering of time points according to a specific conception to provide foundations for a time discounting function. In the context of radical uncertainty about the future, we simply could change the interpretation of the delay ordering to one of a radical uncertainty ordering. Thus, we would be asking how much radical uncertainty about the future is associated with each time point.

While this seems to be a possible way of proceeding, we do have to note that the foundations of time discounting of Ok and Masatlioglu (2007) still operate in a a standard time preference framework: the representation theorem specifies a discounted utility function, based on both time preferences on an outcome time domain and the delay ordering on the time domain. One consequence of this entanglement with time preferences is that we are not given completely separate axioms for the delay ordering; rather, the conditions on the delay ordering are implied by the specific way in which the time preference framework is set up. This means that the framework does only characterise the delay ordering indirectly, by specifying a degree of freedom in which the time preferences can vary under the various delays. This feature of the framework, while attractive for other reasons, means that it is rather difficult to use for characterising radical uncertainty about the future. What is important to note, though, is that the delay feature of the framework makes it in principle compatible with such an interpretation.

4.3 Generalised Representations

We now turn to a more general representation framework for time discounting, given in Heilmann (2008), that operates completely outside the time preference. The framework proceeds in three steps. Firstly, it introduces a domain of externally given time points T. All this means is simply that we can assume that a set of time points has an inherent, given ordering: for instance, we can assume a set of discrete time points such as $T = \{0, 1, \ldots, t\}$. Secondly, the framework introduces a domain of events Q which happen both at a specified time T and can be ordered according to a specific conception. Thirdly, if the ordering according to the conception corresponds to the externally given time points T in the right kind of way, then there is a time discounting function (with similar properties as the one specified earlier in this paper) that represents the ordering of the events Q.

Thus, what drives time discounting in this general framework it is not that the time points that are ordered directly, but rather timed events are ordered according to a specific conception. In order for this framework to conceptually motivate a time discounting function, one needs to fix a specific conception according to which the events are ordered. In order to provide foundations for time discounting, one interpretation of this ordering has to be fixed, and the ordering has to correspond to the given external ordering of the time points. For instance, the events can be compared with regards to how much impatience, or uncertainty, or delay perception they induce. In the context of radical uncertainty about the future, the events could be ordered according to how much radical uncertainty about the future they are associated with.

It is easy to see that this general representation is completely unconnected to the value domain, and thus does not face the kinds of difficulties noted for the standard framework. Moreover, given that the representation framework assumes that events that happen at the same time are equivalence classes, it rules out that different events that can happen at the same time can induce a difference in the discounting function. It seems thus that the general representation could be combined with an interpretation of radical uncertainty about the future, such that the time discounting function is a measure of it.

4.4 Interpretation

The two frameworks just introduced offer some hope that time discounting functions can be a measure of radical uncertainty about the future, as the representation theorems derive time discounting functions as (partly) independent of the value domain. That is, these frameworks may be compatible with an interpretation by a notion of radical uncertainty: delay orderings can be seen as expressing how much uncertainty they are associated with in the Ok and Masatlioglu (2007) framework, and equivalence classes that all happen at the same point can be ordered according to how much uncertainty they are associated with in the Heilmann (2008) framework. There are two problems with these interpretations. (i) Firstly, how exactly can those interpretations capture radical uncertainty about the future? (ii) Secondly, how could discounting functions thus derived be used together with standard decision-theoretic frameworks?

Concerning the (i) interpretation, there are a number of options. Let me elaborate on one. We could simply say that radical uncertainty about the future that corresponds to time and increases with it is a meta-judgement on the confidence in the representation of an agent's epistemic state. That is, we can leave any subjective probability framework in place and simply ask the agent how much her confidence in her judgements varies with time. This interpretation is possible in both the Ok and Masatlioglu (2007) and the more general one of Heilmann (2008), where we assume that many events could happen at each time. We could present the agent with the task of ordering judgements according to how confident she is in them, and then check whether these kinds of orderings can be represented by a time discounting function. Concerning the interpretation, we would then be able to use the discounting measure as giving us an indication of the reliability of probability assessments, telling us to what extent we should depart from precise probabilities, and how big the probability range should be. In this interpretation, we might even design an elicitation exercise that asks agents to order the reliabilities of their probability judgements, which may well decrease with time. This is a way in which discounting functions could be a measure of radical uncertainty about the future.

Concerning the (ii) use of the discounting function together with standard decision-theoretic accounts, this concerns the question of how we are, if at all, to 'discount for radical uncertainty'. The discussion has shown so far, that only purely epistemic frameworks can provide sound foundations for time discounting functions that are to be interpreted as a measure of radical uncertainty about the future. The one route open, then, for using such a discounting function in a standard decision-theoretic framework would be to use it alongside a standard probability function as an additional measure of an agent's degree of belief or an objective probability. How such an additional measure should be employed, though, depends on the exact interpretation of radical uncertainty is used in the derivation and elicitation. If we interpret the measure of radical uncertainty as capturing decreasing reliability of probability assessments of agents, then it might not be plausible to use these measures as weightings (as in standard discounting). They could still be used to, for instance, inform the calibration of probability ranges.

There are thus some prospects for discounting measures in connection with radical uncertainty about the future. If the latter is genuinely general, that is, not connected to specific (causal) paths that can be evaluated by standard probability measures, then a function that assigns decreasing values to future time points might be capturing a general and radical uncertainty about the future. These speculations should, however, not distract too much from the general message of caution regarding discounting function capturing uncertainty about the future. The standard frameworks cannot serve this role; some general frameworks might be compatible with it.

5 Conclusions

We started with the twofold intuition that the future is radically uncertain and that it increases: the further away we look into the future, the more uncertain we are about it. We then asked the question whether there is any hope to make this intuition more precise, and looked at the literature on time discounting in economics to see whether their frameworks can be used in such a way. The question was thus whether time discounting functions can be a measure of radical uncertainty about the future.

Discussing standard frameworks in time discounting, we have seen that they cannot provide foundations for a measure of radical uncertainty about the future. The way the representation frameworks are set up does not permit us to do so, chiefly because they are defined closely related to value. We have seen, however, that those frameworks that derive the discounting function separately, offer some prospects for providing a measure of radical uncertainty about the future, provided they solve two problems: firstly, how exactly the notion of radical uncertainty about the future is cached out in their frameworks, and secondly, how the discounting measure is to be applied in conjunction with standard decision-theoretic frameworks.

The main message of this paper is thus one of caution: ideas such as 'discounting for uncertainty' or the notion of 'discounting functions expressing uncertainty about the future' lack foundations in standard time discounting frameworks. There are only some discounting frameworks are in principle open to interpretations of radical uncertainty about the future, which provides some prospects for deriving a measure of radical uncertainty about the future.

Notes

[1]Often, a discount rate r is assumed which expresses how much weight the period $t+1$ is assigned to in period t which givces a discounting factor $\delta = \left(\frac{1}{1+r}\right)$. Thus, time discounting factors are frequently derived in the following way: $D(t) = \left(\frac{1}{1+r}\right)^t$.

References

Ainslie, G. (1992). *Picoeconomics: the Interaction of Successive Motivational States within the Individual*. Cambridge University Press.

Ainslie, G. (2001). *Breakdown of Will*. Cambridge University Press.

Bradley, R. (2007). A unified Bayesian decision theory. *Theory and Decision*, **63**, 233–263.

Fishburn, P. C. and Rubinstein, A. (1982). Time preference. *International Economic Review*, **23**(3), 677–94.

Frederick, S., Loewenstein, G., and O'Donoghue, T. (2002). Time discounting and time preference: A critical review. *Journal of Economic Literature*, **40**(2), 351–401.

Heilmann, C. (2008). Measurement-theoretic foundations of time discounting in economics. *LSE Choice Group Working Papers*, **4**(6).

Jeffrey, R. C. (1983). *The Logic of Decision*. University of Chicago Press.

Joyce, J. (1999). *The Foundations of Causal Decision Theory*. Cambridge University Press.

Koopmans, T. (1960). Stationary ordinal utility and impatience. *Econometrica*, **28**(2), 287–309.

Lancaster, K. (1963). An axiomatic theory of consumer time preference. *International Economic Review*, **4**(2), 221–231.

Mas-Colell, A., Whinston, M., and Green, J. (1995). *Microeconomic theory*. Oxford University Press New York.

Medvecky, F. (2012). Valuing environmental costs and benefits in an uncertain future: risk aversion and discounting. *Erasmus Journal for Philosophy and Economics*, **5**(1), 1–23.

Ok, E. A. and Masatlioglu, Y. (2007). A theory of (relative) discounting. *Journal of Economic Theory*, **137**, 214–45.

Ramsey, F. P. (1928). A mathematical theory of saving. *Economic Journal*, **38**(152), 543–59.

Samuelson, P. (1937). A note on measurement of utility. *Review of Economic Studies*, **4**, 155–61.

Savage, L. J. (1972). *The Foundations of Statistics*. Dover Publications.

Scholten, M. and Read, D. (2006). Discounting by intervals: A generalized model of intertemporal choice. *Management Science*, **52**, 1424–1436.