The University of Manchester

Economics

Discussion Paper

Series

EDP-2303

Dynamic Preference Foundations of Expected Exponentially-Discounted Utility

Craig S. Webb

May 2023

Economics

School of Social Sciences
The University of Manchester
Manchester M13 9PL

Dynamic Preference Foundations of Expected Exponentially-Discounted Utility.

Craig S. Webb [*]

May 24, 2023

Abstract

Expected exponentially-discounted utility (EEDU) is the standard model of choice over risk and time in economics. This paper considers the dynamic preference foundations of EEDU in the timed risks framework. We first provide dynamic preference foundations for a time-invariant expected utility representation. The new axioms for this are called foregone-risk independence and strong time invariance. This class of dynamic preferences includes EEDU as a special case. If foregone-risk independence is strengthened to a new condition called conditional consistency, then an EEDU representation results. Alternative approaches for extending exponential discounting axioms to risk are considered, resulting in five new preference foundations of EEDU.

[*]Economics, University of Manchester, Oxford Road, Manchester M13 9PL, United Kingdom; E-mail: craig.webb@manchester.ac.uk. I am grateful to Hao Bai, Horst Zank, and two anonymous reviewers for helpful comments. The usual disclaimer applies.

1 Introduction.

Risk and time are often treated separately in economics, yet many decisions involve an element of both. The standard model that combines risk and time in economics is *expected exponentially-discounted utility* (EEDU). It combines the most widely used model for choice under risk, expected utility, and for riskless choice over time, exponential discounting. This paper considers the preference foundations of EEDU.

A natural starting point for a foundation of EEDU is to use existing axioms for expected utility (Fishburn, 1970) at each point in time and to use the existing axioms for exponential discounting (Fishburn and Rubinstein, 1982; Attema, Bleichrodt, Rohde and Wakker, 2010) for riskless objects. It is known (Abdellaoui, Diecidue and Onculer, 2011) that this approach does not deliver EEDU, in particular because the utility derived over time and the utility derived for risk need not be cardinally equivalent. The EEDU model, as such, requires its own treatment. A consideration of how risks are judged though time, as opposed to a separate treatment of risk and time.

The foundations of EEDU have been considered in various frameworks, such as lotteries over riskless consumption streams (Epstein, 1983; Hayashi, 2003), streams of independent lotteries (Anchugina, 2017), and uncertain / ambiguous streams of outcomes for EEDU with (possibly sets of) subjective probabilities (Kochov, 2015; Bastianello and Faro, 2022). We consider a framework of *timed risks*. The decision maker receives one outcome, at one point in time, but this timed outcome is risky. The timed risks framework has been used before (Nachman, 1975; Prakash, 1977; Fishburn and Rubinstein, 1982; Dejarnette et al, 2020; Ebert, 2020) and its simplified structure makes it a suitable testing ground for developing preference conditions that clarify the normative and empirical content of EEDU. The timed risk framework lends itself to various economic applications. For example, models of alternating offers bargaining typically use timed outcomes (Rubinstein, 1982), or incorporate a risk of breakdown (Binmore, Rubinstein and Wolinsky, 1986), or do both (see Muthoo, 1999, p.85). The timed risk framework captures each of these. There is also a large literature on optimal stopping problems, applied to the evaluation of various options. The decision time, as well as the random amount and random time of the eventual payoff, are central in the analysis. Such problems fit well with the timed risk framework.

We adopt a dynamic approach, modelling a decision maker as a set of decision-time-indexed preference relations over timed risks. Key to the approach here is that updating occurs at each decision time. The decision maker only considers timed risks that are known not to have paid out before the current decision time. Otherwise, the timed risk is of no interest. Consider the following timed risk, that offers equal chances of £10 at time 1, £20 at time 2, or £30 at time 3:

$$p = \left[\frac{1}{3} \text{ chance, £10 at time 1; } \frac{1}{3} \text{ chance, £20 at time 2; } \frac{1}{3} \text{ chance of £30 at time 3}\right]$$

If the decision maker is considering p at time zero or at time one, then all of the possible outcomes are yet to pass. At time 2, however, the decision maker considering p knows that "£10 at time 1" has passed and did not happen. Updating the probabilities of the remaining outcomes proportionally, as in Bayes rule, gives:

$$p|2 = \left[\frac{1}{2} \text{ chance, £20 at time 2; } \frac{1}{2} \text{ chance, £30 at time 3}\right]$$

Similarly, at decision time 3, the timed risk can be updated again giving:

$$p|3 = \Big[\text{ £30 at time 3}\Big]$$

For all later times, the updated timed risk is not defined. The preference relations at each decision time are defined over the set of those timed risks that remain well-defined after updating. The details are given in Section 2. The implications of this behaviour are captured in a simple and testable preference axiom called *foregone-risk independence*. Consider two timed risks, p and q. Foregone-risk independence requires that, at decision time a, p is preferred to q if and only if $p|a$ is preferred to $q|a$. Section 3 provides the formal details.

Time invariance, the perception of times as delays relative to decision time, is a well-known property of riskless choice over time (Halevy, 2015). It is assumed by most models of discounting. In riskless choice over time this requires that preferences are not reversed if the decision time and the timed outcomes are subject to common delay. Also in Section 3, we extend this condition to risk. We formulate a *strong time invariance* axiom that requires that preferences are not reversed if the decision time and all except a set of possible timed outcomes, common to both timed risks, are subject to common delay. For example, consider the following timed risks:

$$p = \left[\frac{1}{3}\text{ chance, £10 at time 1; }\frac{1}{3}\text{ chance, £20 at time 2; }\frac{1}{3}\text{ chance of £30 at time 3}\right]$$
$$q = \left[\frac{1}{3}\text{ chance, £10 at time 1; }\frac{1}{3}\text{ chance, £20 at time 2; }\frac{1}{3}\text{ chance of £40 at time 4}\right]$$

The "$\frac{1}{3}$ chance of £10 at time 1" and the "$\frac{1}{3}$ chance of £20 at time 2" are common to both p and q. Let us fix those and apply a one unit of time delay to the remaining outcomes:

$$p' = \left[\frac{1}{3} \text{ chance, £10 at time 1; } \frac{1}{3} \text{ chance, £20 at time 2; } \frac{1}{3} \text{ chance of £30 at time 4}\right]$$
$$q' = \left[\frac{1}{3} \text{ chance, £10 at time 1; } \frac{1}{3} \text{ chance, £20 at time 2; } \frac{1}{3} \text{ chance of £40 at time 5}\right]$$

Strong time invariance requires that p is preferred to q at decision time zero if and only if p' is preferred to q' at decision time one. When combined with foregone-risk independence, and basic assumptions concerning preferences, this condition is sufficient to establish an expected utility representation of dynamic preferences. This general expected utility representation includes EEDU as a special case.

In Section 4 it is shown that strengthening foregone-risk independence to an axiom called *conditional consistency* characterises EEDU. The conditional consistency axiom requires, roughly, that actual behaviour is consistent with planned behaviour. The condition is one way to extend time consistency to risk. Alternative ways of extending exponential discounting axioms to risk are considered. Axioms of strong stationarity and of strong time consistency are formulated and the logical relationships between these axioms are explained. Section 5 summarises and states the main theorem paper. The main theorem provides five equivalent axiom sets, each of which characterise those dynamic preferences that admit EEDU representations. All proofs are contained in the Appendix.

2 Preliminaries.

2.1 Notation for Timed Risks.

Let $\mathcal{X}$, the set of outcomes, be a separable metric space. Time is $\mathcal{T} = [0, \infty)$ and has the usual metric. Timed outcomes, such as (x, t), are elements of $\mathcal{X} \times \mathcal{T}$. Throughout the paper, elements of $\mathcal{T}$ are taken to be calendar times. Timed risks, denoted p, q, $r \ldots$ are elements of $\mathcal{L}$, which is the set of simple probability measures over $\mathcal{X} \times \mathcal{T}$ endowed with the topology of weak convergence. A timed risk provides the decision maker with one outcome at one point in time, but both the outcome and its timing are random. The degenerate timed risk, that assigns probability one to a timed outcome (x, t), is written $\delta_{(x,t)}$. We can write a timed risk as

$$p = \sum_{(x,t) \in \mathcal{X} \times \mathcal{T}} p(x, t) \delta_{(x,t)}$$

where $p(x, t) \geqslant 0$ for all $(x, t) \in \mathcal{X} \times \mathcal{T}$, $\sum_{(x,t) \in \mathcal{X} \times \mathcal{T}} p(x, t) = 1$, and $p(x, t) > 0$ for finitely many $(x, t) \in \mathcal{X} \times \mathcal{T}$.

We will be interested in cases where the timed outcomes of a given timed risk are subject to a common delay. Given a timed risk $p \in \mathcal{L}$ and $\Delta \geqslant 0$, we write

$$p_{\Delta} = \sum_{(x,t) \in \mathcal{X} \times \mathcal{T}} p(x, t) \delta_{(x,t+\Delta)},$$

which is the timed risk p with all possible timed outcomes delayed by Δ. We will also be interested in cases where some, but perhaps not all, timed outcomes are subject to a common delay. Given $\mathcal{S} \subseteq \mathcal{T}$, we write $[p, \mathcal{S}, r]$ to denote the timed risk such that

$$[p, \mathcal{S}, r]\,(x,t) = \begin{cases} p(x,t) & \text{for all } (x,t) \in \mathcal{X} \times \mathcal{S}, \\ r(x,t) & \text{for all } (x,t) \in \mathcal{X} \times \mathcal{T} \setminus \mathcal{S}. \end{cases}$$

That is, the timed risk that coincides with p for all times in $\mathcal{S}$ and coincides with r elsewhere. Note that $[p, \mathcal{S}, r] \in \mathcal{L}$ if and only if $p(\mathcal{X} \times \mathcal{S}) = 1 - r(\mathcal{X} \times \mathcal{T} \setminus S)$.

Given $p \in \mathcal{L}$, we denote the marginals on $\mathcal{X}$ and $\mathcal{T}$ as $p_{\mathcal{X}}$ and $p_{\mathcal{T}}$, respectively. Let $\mathcal{T}(p) = \{t : p_{\mathcal{T}}(t) > 0\}$. Given a subset $\mathcal{S} \subseteq \mathcal{T}$, we let $\mathcal{S}(p) = \mathcal{T}(p) \cap \mathcal{S}$. That is, $\mathcal{S}(p)$ is the times in $\mathcal{S}$ of those timed outcomes to which p assigns strictly positive probability. Combining this with the above notations, if $[p, \mathcal{S}, r] \in \mathcal{L}$ then $[p_\Delta, \mathcal{S}, r]$ is well-defined for all Δ such that $t \in S(p)$ implies $t + \Delta \in \mathcal{S}$.

2.2 Dynamic Preferences and Foregone-Risk Independence.

At time zero, we will define initial preferences $\succcurlyeq_0$ over the set of timed risk $\mathcal{L}$. To specify a dynamic preference, we want to consider a preference relation $\succcurlyeq_a$ at each decision time $a \in \mathcal{T}$. But, for decision times $a > 0$, consider the following question: for what set of timed risks should preferences $\succcurlyeq_a$ be defined? Two extreme cases are evident. First, timed risks with possible outcomes only at time a or later. It seems a basic requirement that the decision maker at time a can rank such objects. Second, consider timed risks with possible outcomes strictly before time a. At time a the outcomes of such timed risks have passed, and so we will not require that the decision maker can rank such objects. We could assume that the decision maker each time is indifferent between all passed timed outcomes. Instead, we will say that $\succcurlyeq_a$ is simply not defined for such objects. We will define each $\succcurlyeq_a$ only on a relevant subset of $\mathcal{L}$.

Given decision time $a > 0$, consider timed risks with possible outcomes both before and after time a. We will require that the decision maker can rank such objects at time a. These rankings, however, need not agree with $\succcurlyeq_0$ because outcomes that occur before a are, in this model, irrelevant at decision time a. To this end, we will consider how timed risks are updated as time passes. The decision maker only considers timed risks that are known not to have paid out before the current decision time. Hence, as time passes, the decision maker can update the probabilities attached to the remaining outcomes. Let $p|a$ denote the timed risk p conditional on decision time a. Let $t_p = \max\{t : p(x,t) > 0\}$ denote the latest time at which p can possibly pay out. For $0 \leqslant a \leqslant t_p$, we define $p|a$ so that:

1. $p|a\,(x,t) = 0$ for all (x,t) with $t < a$.

2. $p|a\,(x,t) = \frac{p(x,t)}{p(\mathcal{X}\times[a,\infty))}$ for all (x,t) with $t \geqslant a$.

The notation $p|a$ is not defined if $a > t_p$ and we write $p|a \in \mathcal{L}$ if and only if $p|a$ is well defined. We can now define dynamic preferences. For all decision times $a \in \mathcal{T}$ a preference relation $\succcurlyeq_a$ is defined over the set $\mathcal{L}_a = \{p : p|a \in \mathcal{L}\}$. That is, $\succcurlyeq_a$ is defined over all timed risks p such that $p|a$ is well defined. This captures the idea of defining each $\succcurlyeq_a$ only on a relevant subset of $\mathcal{L}$. If one updates at time a and a timed risk is still well defined, then it still has something to offer and so is relevant. Note that $\mathcal{L}_0 = \mathcal{L}$ and that $a \leqslant b$ implies $\mathcal{L}_a \supseteq \mathcal{L}_b$. Furthermore, for all $a \in \mathcal{T}$, $\mathcal{L}_a$ is a nonempty convex subset of $\mathcal{L}$. A dynamic preference is a set of such preference relations $\{\succcurlyeq_a\}_{a\in\mathcal{T}}$.

Our first axiom, foregone-risk independence, captures the preference implications of a decision maker who updates in the way described above:

Axiom 1 (Foregone-Risk Independence). *For all* $p, q \in \mathcal{L}_a$ *we have* $p|a \succcurlyeq_a q|a$ *if and only if* $p \succcurlyeq_a q$.

If the decision maker knows the decision time a and knows that the timed risks under consideration have not yet paid out, then choosing between $p|a$ and $q|a$ really is the same as choosing between p and q. This simply suggests that the decision maker makes use of available information and that they understand probability calculus. A violation of foregone-risk independence would mean not using available information in an appropriate manner and the condition warrants normative status. A caveat is that the assumptions of the framework must be appropriate for the problem at hand. In a framework of sequences, where the decision maker receives an outcome at each point in time, it is not so clear that independence of previous outcomes is appropriate. In the sequences framework, where previous outcomes did actually happen, the memory of such consumption may well be relevant (Gilboa, Postlewaite and Samuelson, 2016). In the timed risks framework, however, the decision maker will receive only one outcome at one point in time. Foregone-risk independence here means that the decision maker is not affected by timed outcomes that may have previously been possible but never actually happened.

To give a simple example, consider a business owner who is ordering some inventory that can be immediately sold. There are three suppliers, A, B, and C. Orders can be placed up to the day before delivery, and each supplier charges a different price. Supplier A delivers Tuesday or Friday, with equal probabilities, and is medium priced. Supplier B delivers Tuesday, but is expensive. Supplier C delivers Thursday, and is cheap. Suppose that the business owner ordered from supplier A on Monday, but finds after Tuesday that the inventory has not arrived. The business owner updates and re-evaluates their choice. Supplier A is now certain to deliver Friday. The previous possibility that supplier A might deliver Tuesday is not relevant. The owner knows this did not happen. Supplier B is also now irrelevant, as Tuesday has passed. The business owner could, and likely should, switch their order to supplier C.

2.3 Dynamic Preference Conditions.

A dynamic preference is a weak order if, for all $a \in \mathcal{T}$, $\succcurlyeq_a$ is complete and transitive on $\mathcal{L}_a$. It is continuous if, for all $a \in \mathcal{T}$ and all $p \in \mathcal{L}_a$, the sets $\{q \in \mathcal{L}_a : q \succcurlyeq_a p\}$ and $\{q \in \mathcal{L}_a : p \succcurlyeq_a q\}$ are closed. An outcome $x_0 \in \mathcal{X}$ is null if, for all $a \leqslant s, t$, we have $\delta_{(x_0,s)} \sim_a \delta_{(x_0,t)}$. A set of outcomes $\mathcal{X}$ with a non-empty subset of null outcomes $\mathcal{X}_0 \subset \mathcal{X}$ is non-negative if for all $x \in \mathcal{X}$, $x_0 \in \mathcal{X}_0$ and $a \leqslant t$ we have $\delta_{(x,t)} \succcurlyeq_a \delta_{(x_0,t)}$, and is non-trivial if $\mathcal{X} \setminus \mathcal{X}_0$ is non-empty. If $\mathcal{X}$ is non-negative and non-trivial, a dynamic preference satisfies impatience if for all $x \in \mathcal{X} \setminus \mathcal{X}_0$ and all $a, t, s \in T$ we have $a \leqslant t < s$ if and only if $\delta_{(x,t)} \succ_a \delta_{(x,s)}$. The following assumptions will be used throughout the paper:

Definition 1 (Basic Assumptions)**.** *The dynamic preference $\{\succcurlyeq_a\}_{a \in \mathcal{T}}$ is a continuous and impatient weak order. The set of outcomes X has a non-empty subset of null outcomes, is non-negative, and is non-trivial.*

The non-negativity assumption is purely for convenience and everything here can be extended to include negative outcomes. Non-triviality is necessary for non-constant utility representations. That X includes at least one null outcome is, however, an assumption we cannot dispense with unless the representations obtained are altered. See, for example, Fishburn and Rubinstein (1982, 688-690).

2.4 Dynamic Models.

A dynamic model $\{U_a\}_{a\in\mathcal{T}}$ is a set of real-valued utility functions, $U_a : \mathcal{L}_a \to \mathbb{R}$. A dynamic preference $\{\succcurlyeq_a\}_{a\in\mathcal{T}}$ is represented by the dynamic model $\{U_a\}_{a\in\mathcal{T}}$ if, for all $a \in \mathcal{T}$ and all $p, q \in \mathcal{L}_a$, we have $p \succcurlyeq_a q$ if and only if $U_a(p) \geqslant U_a(q)$. That is, a utility representation of each of the preference relations in the dynamic preference. It has been known since Debreu (1954, 1964) that the basic assumptions are necessary and sufficient for $\{\succcurlyeq_a\}_{a\in\mathcal{T}}$ to be represented by a dynamic model $\{U_a\}_{a\in\mathcal{T}}$ with each U_a being continuous and strictly decreasing with respect to time. We will consider two special cases below.

3 Expected Time Invariant Utility.

A timed risk p can be identified with a random timed outcome (X, T) that is distributed such that $Pr\left((X, T) = (x, t)\right) = p(x, t)$ for all $(x, t) \in \mathcal{X} \times \mathcal{T}$. Given a utility function for timed-outcomes, $U : \mathcal{X} \times \mathcal{T} \to \mathbb{R}$, we will use the notation

$$\mathbb{E}_p[U(X, T)] = \sum_{\{(x,t):p(x,t)>0\}} p\,(x, t)\, U(x, t)$$

to denote the expected utility of a timed risk p. A dynamic model conforms to expected time invariant utility (ETIU) if there exists a utility function for timed-outcomes $U : \mathcal{X} \times \mathcal{T} \to \mathbb{R}$ such that, for all $a \in \mathcal{T}$ and all $p, q \in \mathcal{L}_a$, we have:

$$p \succcurlyeq_a q \;\Leftrightarrow\; \mathbb{E}_{p|a}[U(X, T - a)] \geqslant \mathbb{E}_{q|a}[U(X, T - a)].$$

Three things are apparent in this representation: the expected utility form, the updating of timed-risks at each decision time, and the treatment of time relative to decision time. That is, the decision maker is concerned with outcomes and their respective delays. As expectation is defined above, a timed outcome (x, t) will be included in the calculation of $\mathbb{E}_{p|a}[U(X, T-a)]$ only if $p|a(x,t) > 0$, which is possible only if $t \geqslant a$. The updating at each decision time implies that ETIU preferences necessarily satisfy foregone-risk independence.

A dynamic preference $\{\succcurlyeq_a\}_{a \in \mathcal{T}}$ satisfies the independence axiom if, for all $a \in \mathcal{T}$, $p, q, r \in \mathcal{L}_a$ and $0 \leqslant \alpha \leqslant 1$ we have:

$$p \succcurlyeq_a q \quad \text{if and only if} \quad \alpha p + (1-\alpha)\, r \succcurlyeq_a \alpha q + (1-\alpha)\, r.$$

If preferences can be represented by ETIU, then they must satisfy the independence axiom. Independence is a static axiom, but most convincing normative defences invoke dynamic arguments in atemporal frameworks using compound lotteries (Hammond 1988, Karni and Schmeidler, 1991). We replace the independence axiom with dynamic conditions appropriate for the timed-risk framework.

3.1 Strong Time Invariance.

Delays are differences between two calendar times. Calendar times are, therefore, delays from a specified time zero. It may, however, be that the decision maker finds the delay from decision time to be more relevant than the delay from time zero. If, for example, (x, t) refers to receiving an outcome x on 1st January 2040, and the current decision time is 1st of January 2023, then (x, t) refers to receiving an outcome x with a delay of 17 years. As the decision time changes, the delay associated with

(x, t) changes, but its calendar time does not. A decision maker has time invariant preferences if, essentially, they focus on delays when assessing timed risks.

Time invariance is a well-known condition in intertemporal decision making. The property was often implicitly assumed in the literature on time discounting, until Halevy (2015) provided a preference definition. For riskless timed outcomes, this states that preferences are not reversed if all timed outcomes and the decision time are delayed by a common amount. For timed risks, one could ask that preferences are not reversed if all *possible* timed outcomes and the decision time are delayed by a common amount. If a possible timed outcome is common to both timed risks, then we ask that preferences are unaffected if everything except these are delayed by a common amount. We call this strong time invariance:

Axiom 2 (Strong Time Invariance)**.** *For all* $a \in \mathcal{T}$*,* $[p, \mathcal{S}, r], [q, \mathcal{S}, r] \in \mathcal{L}_a$*, and* $\Delta \geqslant 0$ *such that* $t \in \mathcal{S}$ *implies* $t \geqslant a$ *and* $t \in \mathcal{S}(p) \cup \mathcal{S}(q)$ *implies* $t + \Delta \in \mathcal{S}$*, we have* $[p, \mathcal{S}, r] \succcurlyeq_a [q, \mathcal{S}, r]$ *if and only if* $[p_\Delta, \mathcal{S}, r] \succcurlyeq_{a+\Delta} [q_\Delta, \mathcal{S}, r]$*.*

Suppose that $[p, \mathcal{S}, r], [q, \mathcal{S}, r] \in \mathcal{L}_a$. The timed outcomes that are not common to both of these timed risks occur only at times in $\mathcal{S}$. Specifically, they occur at times in $\mathcal{S}(p) \cup \mathcal{S}(q)$. If we wish to delay those timed outcomes by $\Delta \geqslant 0$, whilst preserving the common timed outcomes outside of $\mathcal{S}$, then we can ensure this by choosing Δ such that the delayed times in $\mathcal{S}(p) \cup \mathcal{S}(q)$ remain in the subset $\mathcal{S}$. That is, choose $\Delta \geqslant 0$ such that $t \in \mathcal{S}(p) \cup \mathcal{S}(q)$ implies $t + \Delta \in \mathcal{S}$. If the subset $\mathcal{S}$ occurs after decision time a, so that $t \in \mathcal{S}$ implies $t \geqslant a$, and Δ is chosen appropriately, then $[p_\Delta, \mathcal{S}, r]$ and $[q_\Delta, \mathcal{S}, r]$ are well defined and both belong to $\mathcal{L}_{a+\Delta}$. The axiom then requires that preferences are not reversed if the decision time and those timed outcomes in $\mathcal{S}$ are all delayed by a common amount.

Restricting strong time invariance to degenerate timed risks gives that condition that, for all $\delta_{(x,t)}, \delta_{(y,s)} \in \mathcal{L}_a$ we have $\delta_{(x,t)} \succsim_a \delta_{(y,s)}$ if and only if $\delta_{(x,t+\Delta)} \succsim_{a+\Delta} \delta_{(y,s+\Delta)}$, which is time invariance. Strong time invariance extends this to risk. We state as a Proposition that strong time invariance is necessary for ETIU maximisation:

Proposition 1. *If a dynamic preference $\{\succsim_a\}_{a\in\mathcal{T}}$ can be represented by expected time invariant utility then it must satisfy strong time invariance.*

3.2 ETIU Representation.

In the above, we have introduced the axioms of foregone-risk independence and of strong time invariance. If dynamic preferences satisfy both of these axioms, then the following result, which is central to our foundation of ETIU, can be obtained :

Proposition 2. *Consider a dynamic preference $\{\succsim_a\}_{a\in\mathcal{T}}$ that satisfies the basic assumptions. If $\{\succsim_a\}_{a\in\mathcal{T}}$ satisfies foregone-risk independence and strong time invariance, then it satisfies the independence axiom.*

The main idea of the proof of Proposition 2 can be explained as follows. Consider two timed risks $p = p_1\delta_{(x,0)} + p_2\delta_{(y,2)}$ and $q = q_1\delta_{(x',0)} + q_2\delta_{(y',2)}$ that have possible timed outcomes that occur at times 0 and 2, and let $r = \delta_{(z,1)}$. Taking $\mathcal{S} = [0,1) \cup (1,\infty)$ and $\Delta = 2$ we can apply strong time invariance and $\alpha p + (1-\alpha)r \succsim_0 \alpha q + (1-\alpha)r$ is equivalent to:

$$\alpha\left(p_1\delta_{(x,2)} + p_2\delta_{(y,4)}\right) + (1-\alpha)r \succsim_2 \alpha\left(q_1\delta_{(x',2)} + q_2\delta_{(y',4)}\right) + (1-\alpha)r.$$

Updating these timed risks at decision time 2 gives

$$\left(\alpha\left(p_1\delta_{(x,2)} + p_2\delta_{(y,4)}\right) + (1-\alpha)r\right)|2 = p_1\delta_{(x,2)} + p_2\delta_{(y,4)}$$

and

$$\left(\alpha\left(q_1\delta_{(x',2)}+q_2\delta_{(y',4)}\right)+(1-\alpha)r\right)|2=q_1\delta_{(x',2)}+q_2\delta_{(y',4)},$$

and so applying foregone-risk independence, the above holds if and only if:

$$p_1\delta_{(x,2)}+p_2\delta_{(y,4)}\succcurlyeq_2 q_1\delta_{(x',2)}+q_2\delta_{(y',4)}.$$

Then, we must have $p\succcurlyeq_0 q$, because if $p\prec_0 q$, then taking $\mathcal{S}=\mathcal{T}$ and $\Delta=2$ and applying strong time invariance would yield a preference contradicting the above. In this way, the independence axiom is established. The proof of Proposition 2 applies this idea more generally. In the proof, continuity of preferences with respect to time (as given in the basic assumptions) is used explicitly. The above argument assumes that the possible outcomes of r that occur at different times to those of p and of q. If there are possible outcomes with the same timing, there may not exist an $\mathcal{S}$ such that the argument applies. This is handled by perturbing the timings of r's possible outcomes and appealing to a continuity argument.

The following Proposition provides a dynamic preference foundation for ETIU:

Proposition 3. *For a dynamic preference $\{\succcurlyeq_a\}_{a\in\mathcal{T}}$ that satisfies the basic assumptions, the following statements are equivalent:*

1. *$\{\succcurlyeq_a\}_{a\in\mathcal{T}}$ satisfies foregone-risk independence and strong time invariance.*

2. *$\{\succcurlyeq_a\}_{a\in\mathcal{T}}$ can be represented by the dynamic model $\{U_a\}_{a\in\mathcal{T}}$ such that, for all $a\in\mathcal{T}$ and all $p\in\mathcal{L}_a$ and , we have:*

$$U_a(p)=\mathbb{E}_{p|a}[U(X,T-a)],$$

where $U : \mathcal{X} \times \mathcal{T} \to \mathbb{R}$ is continuous and is strictly decreasing in $t - a$ for all non-null x.

If $\{\succcurlyeq_a\}_{a\in\mathcal{T}}$ is represented by $\{U_a\}_{a\in\mathcal{T}}$ as in statement 2 then $\{\succcurlyeq_a\}_{a\in\mathcal{T}}$ is also represented by $\{V_a\}_{a\in\mathcal{T}}$ such that $V_a(p) = \mathbb{E}_{p|a}[V(X, T - a)]$ for all $a \in \mathcal{T}$ and all $p \in \mathcal{L}_a$ if and only if there exists $\lambda > 0$ and $\kappa \in \mathbb{R}$ such that $V = \lambda U + \kappa$. That is, utility is a cardinal scale.

4 Expected Exponential Discounting.

A dynamic model conforms to expected exponentially-discounted utility (EEDU) if there exists a utility function for outcomes $u : \mathcal{X} \to \mathbb{R}$ and a discount factor $0 < \beta < 1$ such that, for all $a \in \mathcal{T}$ and all $p, q \in \mathcal{L}_a$, we have:

$$p \succcurlyeq_a q \iff \mathbb{E}_{p|a}[\beta^T u(X)] \geqslant \mathbb{E}_{q|a}[\beta^T u(X)].$$

Hence, EEDU is the special case of ETIU where $U(x, t - a) = \beta^t u(x)$ for all $(x, t) \in \mathcal{X} \times \mathcal{T}$. Neither the discount factor nor the utility function depend here on the decision time a. Although we apply the discount factor β^t to utility at time $t \geqslant a$, we could equivalently apply the discount factor β^{t-a} at each decision time a and represent the very same preferences. This section provides a dynamic preference foundation for EEDU. It has been shown above that, when combined with the basic assumptions, the axioms of foregone-risk independence and strong time invariance characterise ETIU. EEDU is a special case of ETIU, so these axioms remain necessary, but they are not sufficient.

It is well-known that, in riskless choice over time, preferences represented by dynamic exponential discounting must satisfy time consistency. Translating that condition to degenerate timed risks, time consistency is satisfied if, for all $a \leqslant b$ and all $\delta_{(x,s)}, \delta_{(y,t)} \in \mathcal{L}_b$, we have $\delta_{(x,s)} \succcurlyeq_a \delta_{(y,t)}$ if and only if $\delta_{(x,s)} \succcurlyeq_b \delta_{(y,t)}$. One might conjecture that appending this condition to the axiom set in Proposition 3 is sufficient for an EEDU representation. However, as has been observed several times (Abdellaoui, Diecidue and Onculer, 2011; Dejarnette et al, 2020), this is not the case. One can obtain an ETIU representation with von Neumann-Morgenstern utility U for timed risks, and one can obtain an exponential discounting representation $V(x,t) = \beta^t u(x)$ for degenerate timed risks, but there is nothing that requires $U = V$. Indeed, they need not even be cardinally equivalent. As U and V both represent preferences over degenerate timed risks, the most that can be said is that they must be ordinally equivalent. That is, there exists a strictly increasing function $\phi : \mathbb{R} \to \mathbb{R}$ such that $U = \phi \circ V$. We are concerned with how the time consistency condition can be extended to timed risks to obtain EEDU. We now strengthen the foregone-risk independence axiom in a way that encapsulates time consistency.

4.1 Conditional Consistency and EEDU.

Consider the following axiom called *conditional consistency*:

Axiom 3 (Conditional Consistency)**.** *For all $p, q \in \mathcal{L}_b$ and $a \leqslant b$ we have $p|b \succcurlyeq_a q|b$ if and only if $p \succcurlyeq_b q$.*

The preference $p|b \succcurlyeq_a q|b$ could be interpreted as a plan, made at time a, to choose p over q when time b arrives. The preference $p \succcurlyeq_b q$ means p is actually chosen over q at time b. Hence, conditional consistency captures the idea that actual behaviour is consistent with planned behaviour. Conditional consistency strengthens foregone-risk independence, which can be seen by taking $a = b$. Restricting conditional consistency to degenerate timed risks, noting that $\delta_{(x,s)}|b = \delta_{(x,s)}$ for all $\delta_{(x,s)} \in \mathcal{L}_b$, we have time consistency. Conditional consistency is the precise strengthening of foregone-risk independence that delivers a dynamic preference for EEDU:

Proposition 4. *For a dynamic preference $\{\succsim_a\}_{a\in\mathcal{T}}$ that satisfies the basic assumptions, the following statements are equivalent:*

1. *$\{\succsim_a\}_{a\in\mathcal{T}}$ satisfies conditional consistency and strong time invariance.*

2. *$\{\succsim_a\}_{a\in\mathcal{T}}$ can be represented by the dynamic model $\{U_a\}_{a\in\mathcal{T}}$ such that, for all $p \in \mathcal{L}$ and all $a \in \mathcal{T}$, we have:*

$$U_a(p) = \mathbb{E}_{p|a}[\beta^T u\,(X)],$$

where $u : X \to \mathbb{R}_+$ is continuous, $u(x_0) = 0$ for all null $x_0 \in X$, and $\beta \in (0, 1)$.

Denote the representation in statement 2 as (u, β). Then $\{\succsim_a\}_{a\in\mathcal{T}}$ is also represented by (v, γ) if and only if $\beta = \gamma$ and $v = \lambda u$ for a constant $\lambda > 0$. That is, the discount factor is unique and utility is a ratio scale.

If we restrict the framework here to include only degenerate timed risks (timed outcomes), then a representation (u, β) can be obtained. In the riskless case, however, it is known that u and β are unique only up to joint power. That is, (v, γ) also represents preferences if and only if there are strictly positive a, b such that $v = au^b$ and $\gamma = \beta^b$. An implication of uniqueness up to joint power is that one is free to choose the discount factor arbitrarily, provided that utility is suitably adjusted. So one learns nothing, for example, by comparing two individuals' discount factors. In the timed risks framework, however, the discount factor obtained in the above Proposition is uniquely determined. Whether it is an appropriate measure of impatience is not clear, but interpersonal comparisons are at least meaningful. See Benoit and Ok (2007) for more on this.

4.2 Alternative Axioms.

In the previous section, conditional consistency was introduced as an extension of time consistency to timed risks. The idea of time consistency in riskless choice is that the decision maker does not reverse previously expressed preferences. For timed risks, if the decision maker prefers p to q at decision a, then it is not immediate that p should be preferred to q at a later time b. It is possible that there are timed outcomes under p that are attractive at time a but that have passed by time b. The timed risks $p|b$ and $q|b$ have no possible timed outcomes in the interval between a and b, and so this does not present as a problem for the conditional consistency condition. One might consider cases where there are possible timed outcomes of p and q between decision times a and b, but these are common to both timed risks. Then, whatever is passed by decision time b under one timed risk, is also passed under the other. In such cases, the requirement that previously expressed preferences are not reversed is reasonable, and we state this as an axiom:

Axiom 4 (Strong Time Consistency)**.** *For all* $a \leqslant b$ *and* $p, q \in \mathcal{L}_b$*, if* $p(x,t) = q(x,t)$ *for all* $(x,t) \in \mathcal{X} \times [a,b)$ *then* $p \succcurlyeq_a q$ *if and only if* $p \succcurlyeq_b q$.

Stationarity is the requirement that preferences, at a fixed decision time, are not reversed if timed outcomes are subject to a common delay. For timed risks, De Jarnette et al (2020) proposed a risk stationarity axiom. This requires that preferences between timed risks, at a fixed decision time, are not reversed if all timed outcomes in the support of those risks are subject to a common delay. We state here a stronger version of this axiom:

Axiom 5 (Strong Stationarity)**.** *For all $a \in \mathcal{T}$, $[p, \mathcal{S}, r], [q, \mathcal{S}, r] \in \mathcal{L}_a$, and $\Delta \geqslant 0$ such that $t \in S$ implies $t \geqslant a$ and $t \in S(p) \cup S(q)$ implies $t + \Delta \in \mathcal{S}$, we have $[p, \mathcal{S}, r] \succcurlyeq_a [q, \mathcal{S}, r]$ if and only if $[p_\Delta, \mathcal{S}, r] \succcurlyeq_a [q_\Delta, \mathcal{S}, r]$.*

Strong stationarity requires that preferences are not reversed if all except a common set of possible timed outcomes are subject to a common delay. Restricting to the case where $\mathcal{S} = \mathcal{T}$ gives the risk stationarity condition. Restricting the condition to degenerate timed risks gives the stationarity axiom of Fishburn and Rubinstein (1982). Stationarity axioms do not seem to have the same normative appeal as time consistency axioms. However, they are necessary and testable implications of EEDU. Furthermore, as Halevy (2015) showed, these conditions are closely related. A similar result, stated next, summarises the relationships between these conditions:

Proposition 5. *Consider a dynamic preference $\{\succcurlyeq_a\}_{a \in \mathcal{T}}$ that satisfies the basic assumptions. Then, the following hold:*

1. *If $\{\succcurlyeq_a\}_{a \in \mathcal{T}}$ satisfies conditional consistency and strong stationarity, then it satisfies the independence axiom.*

2. *If $\{\succcurlyeq_a\}_{a \in \mathcal{T}}$ satisfies foregone-risk independence, then any two of conditional consistency, strong stationarity and strong time invariance imply that all three are satisfied.*

3. *If $\{\succcurlyeq_a\}_{a \in \mathcal{T}}$ satisfies foregone-risk independence and strong time consistency, then it satisfies conditional consistency.*

4. *Any two of strong time consistency, strong stationarity and strong time invariance imply that all three are satisfied.*

5 Summary Main Result.

The following result is the main theorem of the paper, which summarises and follows as a corollary of the Propositions given above:

Main Theorem. *For a dynamic preference $\{\succsim_a\}_{a\in\mathcal{T}}$ that satisfies the basic assumptions, the following statements are equivalent:*

1. $\{\succsim_a\}_{a\in\mathcal{T}}$ *satisfies conditional consistency and strong time invariance.*

2. $\{\succsim_a\}_{a\in\mathcal{T}}$ *satisfies conditional consistency and strong stationarity.*

3. $\{\succsim_a\}_{a\in\mathcal{T}}$ *satisfies foregone-risk independence, strong time consistency, and strong time invariance.*

4. $\{\succsim_a\}_{a\in\mathcal{T}}$ *satisfies foregone-risk independence, strong time consistency, and strong stationarity.*

5. $\{\succsim_a\}_{a\in\mathcal{T}}$ *satisfies foregone-risk independence, strong time invariance and strong stationarity.*

6. $\{\succsim_a\}_{a\in\mathcal{T}}$ *can be represented by expected exponentially-discounted utility.*

The equivalence of statements 1 and 2 follows from statement 2 of Proposition 5. Statements 3 and 4 of Proposition 5 ensure the equivalence of statements 2 and 3 of the Main Theorem. The equivalence of statements 3, 4 and 5 of the Main

Theorem follows from statement 4 of Proposition 5. The equivalence of statement 1 and 6 of the Main Theorem repeats Proposition 4. Hence, the six statements are equivalent, and so provide five different dynamic preference foundations for expected exponentially-discounted utility.

6 Closing Remarks.

This paper provides various preferences conditions for choice over timed risks. We focus on EEDU, as the central model for time and risk in economics. We incorporate in the dynamic framework simple assumptions regarding information. The decision maker knows the current date and knows that the timed risks being considered, as they are being considered, have yet to pay out. In this setting, when combined with basic preference assumptions, the axioms of foregone-risk independence and strong time invariance are necessary and sufficient to establish a time-invariant expected utility representation. This provides a new perspective on expected utility maximisation, as the independence axiom emerges from dynamic preference conditions. How risks are perceived through time, rather than a separate treatment risk and time. Strengthening foregone-risk independence to conditional consistency delivers an expected exponentially discounted utility representation.

The new axioms in this paper are falsifiable and experimental testing, in particular of those axioms relating to updating timed risks, is warranted.[1] Our consideration is more normative. Foregone-risk independence, in particular, seems sufficiently self evident that one would hope violations of the principle are corrected through learning or experience. Time invariance, implied by strong time invariance, has been

[1] We have not discussed empirical challenges to expected utility and exponential discounting. These are significant and well-known. For surveys, see Machina (1987) and Frederick et al (2002).

defended by an arbitrariness principle. If the condition holds, there is no special, arbitrary, time zero as all delays in all time periods are treated the same way. Conditional consistency is more demanding than foregone-risk independence. The extent to which it is compelling, and we believe it is compelling, provides the rationale for exponential discounting over the many other possible forms of discounting. Evaluating and choosing between such principles is important because, in some cases, one cannot have it all. In the context of riskless social choice, for example, where heterogeneous evaluations are aggregated, it has been shown that time consistency and time invariance are incompatible (Jackson and Yariv 2015, Millner and Heal 2018). The formulation of related concepts for risky decision making over time, as provided here, could be used to extend such results.

Appendix: Proofs.

Throughout these proofs we will assume that the basic assumptions hold.

Proof of Proposition 1: Suppose that $\{\succcurlyeq_a\}_{a\in\mathcal{T}}$ is represented by ETIU and consider $a \in \mathcal{T}$, $[p, \mathcal{S}, r], [q, \mathcal{S}, r] \in \mathcal{L}_a$, and $\Delta \geqslant 0$ such that $t \in S$ implies $t \geqslant a$ and $t \in S(p) \cup S(q)$ implies $t + \Delta \in \mathcal{S}$. Then $[p, \mathcal{S}, r] \succcurlyeq_a [q, \mathcal{S}, r]$ if and only if:

$$\sum_{(x,t)\in\mathcal{X}\times\mathcal{S}} p|a(x,t)U(x,t-a)) \geqslant \sum_{(x,t)\in\mathcal{X}\times\mathcal{S}} q|a(x,t)U(x,t-a),$$

where common terms have been cancelled. Timed risks are simple probability measures here and so we can enumerate the timed outcomes of $p|a$ and $q|a$ in $\mathcal{X}\times\mathcal{S}$ that occur with strictly positive probability. We will write that, in $\mathcal{X}\times\mathcal{S}$, $p|a$ has possible timed outcomes $(x_{p1}, t_{p1}), \ldots, (x_{pm}, t_{pm})$ and that $q|a$ has possible timed outcomes

$(x_{q1}, t_{q1}), \ldots, (x_{qn}, t_{qn})$. The above inequality can therefore be written as:

$$p|a(x_{p1}, t_{p1})U(x_{p1}, t_{p1} - a) + \ldots + p|a(x_{pm}, t_{pm})U(x_{pm}, t_{pm} - a)$$
$$\geqslant q|a(x_{q1}, t_{q1})U(x_{q1}, t_{q1} - a) + \ldots + q|A(x_{qn}, t_{qn})U(x_{qn}, t_{qn} - a).$$

Define $p_\Delta|a$ and $q_\Delta|a$ such that $p_\Delta|a(x, t+\Delta) = p|a(x, t)$ and $q_\Delta|a(x, t+\Delta) = q|a(x, t)$ for all $(x, t) \in \mathcal{X} \times \mathcal{S}$. Then, the above is equivalent to:

$$p_\Delta|a(x_{p1}, t_{p1} + \Delta)U(x_{p1}, t_{p1} - a) + \ldots + p_\Delta|a(x_{pm}, t_{pm} + \Delta)U(x_{pm}, t_{pm} - a)$$
$$\geqslant q_\Delta|a(x_{q1}, t_{q1} + \Delta)U(x_{q1}, t_{q1} - a) + \ldots + q_\Delta|a(x_{qn}, t_{qn} + \Delta)U(x_{qn}, t_{qn} - a).$$

As we have done with $p|a$ and $q|a$, we can enumerate the timed outcomes of $p_\Delta|a$ and $q_\Delta|a$ in $\mathcal{X}\times\mathcal{S}$. We denote these as $(x_{p\Delta 1}, t_{p\Delta 1}), \ldots, (x_{p\Delta m}, t_{p\Delta m})$ and $(x_{q\Delta 1}, t_{q\Delta 1}), \ldots, (x_{q\Delta n}, t_{q\Delta n})$. Furthermore, we can enumerate these so that $(x_{p\Delta i}, t_{p\Delta i}) = (x_{pi}, t_{pi} + \Delta)$, for all $i = 1, \ldots, n$ and $(x_{q\Delta i}, t_{q\Delta i}) = (x_{qi}, t_{qi} + \Delta)$ for all $i = 1, \ldots, m$. By replacing these terms in the previous inequality, we have:

$$p_\Delta|a(x_{p\Delta 1}, t_{p\Delta 1})U(x_{p\Delta 1}, t_{p\Delta 1} - (a + \Delta)) + \ldots + p_\Delta|a(x_{p\Delta m}, t_{p\Delta m})U(x_{p\Delta m}, t_{p\Delta m} - (a + \Delta))$$
$$\geqslant q_\Delta|a(x_{q\Delta 1}|a, t_{q\Delta 1})U(x_{q\Delta 1}, t_{q\Delta 1} - (a + \Delta)) + \ldots + q_\Delta|a(x_{q\Delta n}, t_{q\Delta n})U(x_{q\Delta n}, t_{q\Delta n} - (a + \Delta)).$$

Notice that the earliest time in S that $p_\Delta|a$ possibly pays out is no sooner than time $a + \Delta$. Then, for all $a \leqslant b \leqslant a + \Delta$, we must have $p_\Delta|a(x, t) = p_\Delta|b(x, t)$ for all $(x, t) \in \mathcal{X} \times \mathcal{S}$. For all such b, the previous inequality is equivalent to:

$$\sum_{(x,t)\in\mathcal{X}\times\mathcal{S}} p_\Delta|b(x, t)U(x, t - (a + \Delta)) \geqslant \sum_{(x,t)\in\mathcal{X}\times\mathcal{S}} q_\Delta|b(x, t)U(x, t - (a + \Delta)).$$

Choosing $b = a + \Delta$, this is easily seen to be equivalent to $[p_\Delta, \mathcal{S}, r] \succcurlyeq_{a+\Delta} [q_\Delta, \mathcal{S}, r]$. ETUI preferences therefore necessarily satisfy strong time invariance. ■

Proof of Proposition 2: Assume that $\{\succcurlyeq_a\}_{a\in\mathcal{T}}$ satisfies foregone-risk independence and strong time invariance. We will show that $\succcurlyeq_0$ satisfies independence. Showing that each $\succcurlyeq_a$ satisfies independence is entirely similar. Let $p, q \in \mathcal{L}$ and consider $r \in \mathcal{L}$ such that timed outcomes of r occur at different times to those of p and q, $\mathcal{T}(r) \cap \mathcal{T}(p) = \emptyset$ and $\mathcal{T}(r) \cap \mathcal{T}(q) = \emptyset$. The remaining case is considered later. Let $\alpha \in [0, 1]$ and denote $\tilde{p} = \alpha p + (1-\alpha)r$ and $\tilde{q} = \alpha q + (1-\alpha)r$. Assume that $\tilde{p} \succcurlyeq_0 \tilde{q}$. Let $\tilde{\mathcal{S}} = \{t : r(x,t) = 0\}$. Defined in this way, we have $\tilde{p} = \left[\tilde{p}, \tilde{\mathcal{S}}, \tilde{p}\right]$ and $\tilde{q} = \left[\tilde{q}, \tilde{\mathcal{S}}, \tilde{q}\right]$, and for all (x,t) we have:

$$\left[\tilde{p}, \tilde{\mathcal{S}}, \tilde{p}\right](x,t) = \begin{cases} \alpha p(x,t) & \text{if } t \in \tilde{\mathcal{S}}, \\ (1-\alpha)r(x,t) & \text{if } t \notin \tilde{\mathcal{S}}, \end{cases}$$

and:

$$\left[\tilde{q}, \tilde{\mathcal{S}}, \tilde{q}\right](x,t) = \begin{cases} \alpha q(x,t) & \text{if } t \in \tilde{\mathcal{S}}, \\ (1-\alpha)r(x,t) & \text{if } t \notin \tilde{\mathcal{S}}. \end{cases}$$

Let t^* solve $\max t$ subject to $t \in \mathcal{T}(r)$. Choosing $\Delta > t^*$, we have $t \in \tilde{\mathcal{S}}(p) \cup \tilde{\mathcal{S}}(q)$ implies $t + \Delta \in \tilde{\mathcal{S}}$. By strong time invariance, taking $\mathcal{S} = \tilde{\mathcal{S}}$, we have $\tilde{p} \succcurlyeq_0 \tilde{q}$ if and only if $\left[\tilde{p}_\Delta, \tilde{\mathcal{S}}, \tilde{p}\right] \succcurlyeq_\Delta \left[\tilde{q}_\Delta, \tilde{\mathcal{S}}, \tilde{q}\right]$. By foregone-risk independence, this holds if and only if $\left[\tilde{p}_\Delta, \tilde{\mathcal{S}}, \tilde{p}\right]|\Delta \succcurlyeq_\Delta \left[\tilde{q}_\Delta, \tilde{\mathcal{S}}, \tilde{q}\right]|\Delta$. For all (x,t) with $t \geqslant \Delta > t^*$, we have:

$$\left[\tilde{p}_\Delta, \tilde{\mathcal{S}}, \tilde{p}\right]|\Delta(x,t) = \frac{\alpha p_\Delta(x,t)}{\sum_{y\in\mathcal{X}, s\geqslant\Delta} \alpha p_\Delta(y,s)} = p_\Delta(x,t)$$

because $\sum_{y\in\mathcal{X}, s\geqslant\Delta} p_\Delta(y,s) = 1$. That is, $\left[\tilde{p}_\Delta, \tilde{\mathcal{S}}, \tilde{p}\right]|\Delta = p_\Delta$. By the same reasoning,

$\left[\tilde{q}_{\Delta}, \tilde{\mathcal{S}}, \tilde{q}\right] = q_{\Delta}$. So, we have shown that $\alpha p + (1-\alpha) r \succcurlyeq_0 \alpha q + (1-\alpha) r$ if and only if $p_{\Delta} \succcurlyeq_{\Delta} q_{\Delta}$. By strong time invariance, but this time taking $\mathcal{S} = \mathcal{T}$, this holds if and only if $p \succcurlyeq_0 q$, hence $\succcurlyeq_0$ satisfies independence for the case where r is such that timed outcomes of r occur at different times to those of p and q.

We have considered the case where $p, q, r \in \mathcal{L}$ and the timed outcomes of r occur at different times to those of p and q, $\mathcal{T}(r) \cap \mathcal{T}(p) = \emptyset$ and $\mathcal{T}(r) \cap \mathcal{T}(q) = \emptyset$. We now consider with the remaining case where at least one of $\mathcal{T}(r) \cap \mathcal{T}(p)$ or $\mathcal{T}(r) \cap \mathcal{T}(q)$ is non-empty. The idea is to perturb r, apply the above, allow these perturbations to tend to zero, then appeal to continuity. For the details, let $\mathcal{R} \subset \mathcal{T}$ denote the set of times where the timing of an r outcome coincides with the timing of a p or q outcome, $\mathcal{R} = \mathcal{T}(r) \cap (T(p) \cup T(q))$. Given $m > 0$ let $r_{\mathcal{R},m}$ be defined so that $r_{\mathcal{R},m}(x, t + \frac{1}{m}) = r(x,t)$ for all $(x,t) \in \mathcal{X} \times \mathcal{R}$ and $r_{S,m}(x,t) = r(x,t)$ for all $(x,t) \in \mathcal{X} \times \mathcal{T} \setminus \mathcal{R}$. By choosing a sufficiently large m, specifically m such that $\frac{1}{m} < \min\{|t-s| : t \in \mathcal{R}, s \in T(p) \cup T(q)\}$, the timing of outcomes of $r_{\mathcal{R},m}$ do not coincide with the timings of p or q, that is $\mathcal{T}(r_{\mathcal{R},m}) \cap \mathcal{T}(p) = \emptyset$ and $\mathcal{T}(r_{\mathcal{R},m}) \cap \mathcal{T}(q) = \emptyset$. Therefore, by the arguments given above, we have $p \succcurlyeq_0 q$ if and only if $\alpha p + (1-\alpha) r_{\mathcal{R},m} \succcurlyeq_0 \alpha q + (1-\alpha) r_{\mathcal{R},m}$ for all m that are sufficiently large. Notice that $\lim_{m\to\infty} r_{\mathcal{R},m} = r$ and so, by continuity, we have $\alpha p + (1-\alpha) r \succcurlyeq_0 \alpha q + (1-\alpha) r$, as required. ■

Proof of Proposition 3: We show that statement 1 implies statement 2. The converse has been shown in the main text and Proposition 1. If statement 1 holds then, by Proposition 2, $\succcurlyeq_0$ satisfies independence. Combined with the basic assumptions, this guarantees that $\succcurlyeq_0$ admits an expected utility representation, $p \succcurlyeq_0 q$ if and only if $\mathbb{E}_p[U(X,T)] \geqslant \mathbb{E}_q[U(X,T)]$ for a continuous function $U : \mathcal{X} \times \mathcal{T} \to \mathbb{R}$ (Grandmont, 1972). For all x such that $(x, 0) \succ_0 (x_0, 0)$, impatience implies that $U(x,t)$ decreases strictly with t. The utility function U is a cardinal scale, and so $\succcurlyeq_0$ is represented by $\mathbb{E}_p V$ if and only if there exists $\lambda > 0$ and $\kappa \in \mathbb{R}$ such that $V = \lambda U + \kappa$.

Given time $a > 0$, using foregone-risk independence, for all $p, q \in \mathcal{L}_a$ we have $p \succcurlyeq_a q$ if and only if $p|a \succcurlyeq_a q|a$. Notice that all $p \in \mathcal{L}_a$ implies $(p|a)_{-a} \in \mathcal{L}$. Then, by strong time invariance, $p|a \succcurlyeq_a q|a$ if and only if $(p|a)_{-a} \succcurlyeq_0 (q|a)_{-a}$, which holds if and only if $\mathbb{E}_{(p|a)_{-a}}[U(X,T)] \geqslant \mathbb{E}_{(q|a)_{-a}}[U(X,T)]$. Notice that $\mathbb{E}_{(p|a)_{-a}}[U(X,T)]$ can be equivalently expressed as $\mathbb{E}_{(p|a)}[U(X, T-a)]$. Therefore, for all $p, q \in \mathcal{L}_a$, we have $p \succcurlyeq_a q$ if and only if $\mathbb{E}_{(p|a)}[U(X, T-a)] \geqslant \mathbb{E}_{(q|a)}[U(X, T-a)]$. This holds for all $a \in \mathcal{T}$ so an ETIU representation of $\{\succcurlyeq_a\}_{a \in \mathcal{T}}$ exists. ■

Proof of Proposition 4: By Proposition 3, $\succcurlyeq_0$ is represented by $\mathbb{E}_p[U(X,T)]$. Given $a \geqslant 0$, conditional consistency requires that $p \succcurlyeq_a q$ if and only if $p|a \succcurlyeq_0 q|a$, and so $\succcurlyeq_a$ is represented by $\mathbb{E}_{p|a}[U(X,T)]$. There exists a null outcome $x_0 \in X$ and we set $U(x_0, 0) = 0$. Recall that, for all $s, t \in T$, we have $\delta_{(x_0,s)} \sim_0 \delta_{(x_0,t)}$ and so $U(x_0, \cdot) \equiv 0$. As X is non-negative, $U \geqslant 0$. Define U_a such that $U_a(x,t) = U(x, t-a)$ for all (x,t). Proposition 3 also implies that $\succcurlyeq_a$ is represented by $\mathbb{E}_{p|a}[U_a(X,T)]$. Since U and U_a are both von Neumann-Morgenstern utilities for $\succcurlyeq_a$ and the locations are fixed so that $U(x_0, \cdot) \equiv U_a(x_0, \cdot) \equiv 0$, we must have $U_a = \lambda_a U$ for some $\lambda_a > 0$. Then, we have $U(x,t) = U_a(x, t+a) = \lambda_a U(x, t+a)$ for all $(x,t) \in \mathcal{X} \times \mathcal{T}$. Fixing a $(z,r) \in \mathcal{X} \times \mathcal{T}$ with $\delta_{(z,r)} \succ_0 \delta_{(x_0,r)}$, we can define a function $D : T \to \mathbb{R}_{++}$ such that $D(t) = \frac{1}{\lambda_t} = \frac{U(z,r+t)}{U(z,r)}$ for all $t \in T$ and be assured that D does not depend on the choice of (z,r), because $U(z, r+t) = D(t)U(z,r)$ holds for all $(z,r) \in \mathcal{X} \times \mathcal{T}$. Defined as such, D is a continuous and strictly decreasing function and $D(0) = 1$. Given $s, t \in T$, we can choose $r = 0$ followed by $r = s$ to get $D(t) = \frac{U(z,t)}{U(z,0)} = \frac{U(z,s+t)}{U(z,s)}$.

Similarly, we have $D(s) = \frac{U(z,s)}{U(z,0)}$ and so:

$$D(s+t) = \frac{U(z,s+t)}{U(z,0)} = \frac{U(z,s)}{U(z,0)}\frac{U(z,s+t)}{U(z,s)} = D(s)D(t).$$

The Cauchy functional equation $D(s+t) = D(s)D(t)$ therefore holds for all $s, t \in T$. The only continuous solution, not equal to zero everywhere, is $D(t) = \beta^t$ for some $\beta > 0$ (see, for example, Corollary 1.36 in Kannappan, 2009). Defining $u(x) := U(x,0)$ for all $x \in X$, noting that $U(x,t) = D(t)U(x,0)$ for all $(x,t) \in \mathcal{X} \times \mathcal{T}$, we have $U(x,t) = \beta^t u(x)$ for all $(x,t) \in \mathcal{X} \times \mathcal{T}$. Because impatience requires that $\beta^t u(x)$ is strictly decreasing in t for all non-null x, we have $\beta \in (0,1)$. Under the requirement that $U(x_0, \cdot) \equiv 0$, U is a ratio scale and therefore the utility for outcomes u satisfies $u(x_0) = 0$ and is a ratio scale. The discount factor β is uniquely determined because $\beta^t = \frac{U(x,t)}{U(x,0)}$ for all (x,t). ■

Proof of Proposition 5:

Statement 1: Assume that $\{\succcurlyeq_a\}_{a \in \mathcal{T}}$ satisfies conditional consistency and strong stationarity. We show that $\succcurlyeq_0$ satisfies independence. That each $\succcurlyeq_a$ satisfies independence is entirely similar. Let $p, q \in \mathcal{L}$ and consider $r \in \mathcal{L}$ such that $\mathcal{T}(r) \cap \mathcal{T}(p) = \emptyset$ and $\mathcal{T}(r) \cap \mathcal{T}(q) = \emptyset$. The remaining case can be shown, by perturbation and continuity, exactly as in the proof of Proposition 2. Let $\tilde{\mathcal{S}} = \{t : r(x,t) = 0\}$. Let t^* solve $\max t$ subject to $t \in \mathcal{T}(r)$. Choosing $\Delta > t^*$, we have $t \in \tilde{\mathcal{S}}(p) \cup \tilde{\mathcal{S}}(q)$ implies $t + \Delta \in \tilde{\mathcal{S}}$. Let $\alpha \in [0,1]$ and denote $\tilde{p} = \alpha p + (1-\alpha)r$ and $\tilde{q} = \alpha q + (1-\alpha)r$. To develop a contradiction, suppose a violation of independence: $\tilde{p} \succcurlyeq_0 \tilde{q}$ and $p \prec_0 q$.

The assumed $\tilde{p} \succcurlyeq_0 \tilde{q}$ preference, by strong stationarity with $\mathcal{S} = \mathcal{T}$, is equivalent to $\tilde{p}_\Delta \succcurlyeq_0 \tilde{q}_\Delta$. This can be rewritten as $[\tilde{p}_\Delta, \tilde{\mathcal{S}}, \tilde{p}_\Delta] \succcurlyeq_0 [\tilde{q}_\Delta, \tilde{\mathcal{S}}, \tilde{q}_\Delta]$. By strong stationarity again, but now taking $\mathcal{S} = \tilde{\mathcal{S}}$, this preference is equivalent to $[\tilde{p}_{2\Delta}, \tilde{\mathcal{S}}, \tilde{p}_\Delta] \succcurlyeq_0$

$[\tilde{q}_{2\Delta}, \tilde{\mathcal{S}}, \tilde{q}_\Delta]$. Notice that $[\tilde{p}_{2\Delta}, \tilde{\mathcal{S}}, \tilde{p}_\Delta] = [\tilde{p}_{2\Delta}, \tilde{\mathcal{S}}, \tilde{p}_\Delta]|\Delta$ and $[\tilde{q}_{2\Delta}, \tilde{\mathcal{S}}, \tilde{q}_\Delta] = [\tilde{q}_{2\Delta}, \tilde{\mathcal{S}}, \tilde{q}_\Delta]|\Delta$, because the timed outcomes that are possible under these timed risks cannot occur before time Δ. Then, by conditional consistency, this preference is equivalent to $[\tilde{p}_{2\Delta}, \tilde{\mathcal{S}}, \tilde{p}_\Delta] \succcurlyeq_\Delta [\tilde{q}_{2\Delta}, \tilde{\mathcal{S}}, \tilde{q}_\Delta]$.

Now consider the other assumed preference, $p \prec_0 q$. This is equivalent, by strong stationarity with $\mathcal{S} = \mathcal{T}$, to $p_\Delta \prec_0 q_\Delta$. Note that $p_\Delta = [\tilde{p}_\Delta, \tilde{\mathcal{S}}, \tilde{p}]|\Delta$ and $q_\Delta = [\tilde{q}_\Delta, \tilde{\mathcal{S}}, \tilde{q}]|\Delta$, and so by conditional consistency this preference is equivalent to $[\tilde{p}_\Delta, \tilde{\mathcal{S}}, \tilde{p}] \prec_\Delta [\tilde{q}_\Delta, \tilde{\mathcal{S}}, \tilde{q}]$. By strong stationarity, now with $\mathcal{S} = \tilde{\mathcal{S}}$, this is equivalent to $[\tilde{p}_{2\Delta}, \tilde{\mathcal{S}}, \tilde{p}_\Delta] \prec_\Delta [\tilde{q}_{2\Delta}, \tilde{\mathcal{S}}, \tilde{q}_\Delta]$. Thus, we arrive at a contradiction, and so $\succcurlyeq_0$ must satisfy independence.

Statement 2: Assume strong stationarity and strong time invariance. Let $p, q \in \mathcal{L}_b$ and $a \leqslant b$. By strong time invariance, $p|b \succcurlyeq_a q|b$ if and only if $(p|b)_{b-a} \succcurlyeq_b (q|b)_{b-a}$. By strong stationarity, $(p|b)_{b-a} \succcurlyeq_b (q|b)_{b-a}$ if and only if $p|b \succcurlyeq_b q|b$. By foregone-risk independence, $p|b \succcurlyeq_b q|b$ if and only if $p \succcurlyeq_b q$, and so conditional consistency holds.

Next, suppose that $\{\succcurlyeq_a\}_{a\in\mathcal{T}}$ satisfies conditional consistency and strong time invariance. Consider $a \in \mathcal{T}$, $[p, \mathcal{S}, r], [q, \mathcal{S}, r] \in \mathcal{L}_a$, and $\Delta \geqslant 0$ such that $t \in S$ implies $t \geqslant a$ and $t \in S(p) \cup S(q)$ implies $t + \Delta \in \mathcal{S}$. By strong time invariance, $[p, \mathcal{S}, r] \succcurlyeq_a [q, \mathcal{S}, r]$ if and only if $[p_\Delta, \mathcal{S}, r] \succcurlyeq_{a+\Delta} [q_\Delta, \mathcal{S}, r]$. By Proposition 2, $\{\succcurlyeq_a\}_{a\in\mathcal{T}}$ satisfies independence, and so $[p_\Delta, \mathcal{S}, r] \succcurlyeq_{a+\Delta} [q_\Delta, \mathcal{S}, r]$ if and only if $[p_\Delta, \mathcal{S}, \tilde{r}] \succcurlyeq_{a+\Delta} [q_\Delta, \mathcal{S}, \tilde{r}]$, provided that $\tilde{r}$ is chosen so that these timed risks are well-defined. Choose such an $\tilde{r}$ with the additional property that $\tilde{r}(x, t) = 0$ for all $(x, t) \in \mathcal{X} \times [a, a + \Delta)$. Then, $[p_\Delta, \mathcal{S}, \tilde{r}]$ and $[q_\Delta, \mathcal{S}, \tilde{r}]$ must also both assign zero probability to all timed outcomes with times in the $[a, a + \Delta)$ interval. For these timed risks, the timed outcomes in $\mathcal{S}$ occur no earlier than $a + \Delta$, and for times not in S we chose $\tilde{r}$ accordingly. This means that $[p_\Delta, \mathcal{S}, \tilde{r}]\,|(a + \Delta) = [p_\Delta, \mathcal{S}, \tilde{r}]\,|a$ and $[q_\Delta, \mathcal{S}, \tilde{r}]\,|(a + \Delta) = [q_\Delta, \mathcal{S}, \tilde{r}]\,|a$. By conditional consistency, $[p_\Delta, \mathcal{S}, \tilde{r}] \succcurlyeq_{a+\Delta} [q_\Delta, \mathcal{S}, \tilde{r}]$ if and only if $[p_\Delta, \mathcal{S}, \tilde{r}]\,|(a + \Delta) \succcurlyeq_a$

$[q_\Delta, \mathcal{S}, \tilde{r}]\,|(a+\Delta)$, equivalent to $[p_\Delta, \mathcal{S}, \tilde{r}]\,|a \succcurlyeq_a [q_\Delta, \mathcal{S}, \tilde{r}]\,|a$. By foregone-risk independence, $[p_\Delta, \mathcal{S}, \tilde{r}]\,|a \succcurlyeq_a [q_\Delta, \mathcal{S}, \tilde{r}]\,|a$ if and only if $[p_\Delta, \mathcal{S}, \tilde{r}] \succcurlyeq_a [q_\Delta, \mathcal{S}, \tilde{r}]$. With a final application of independence, this is equivalent to $[p_\Delta, \mathcal{S}, r] \succcurlyeq_a [q_\Delta, \mathcal{S}, r]$, and so strong stationarity holds.

Finally, suppose that $\{\succcurlyeq_a\}_{a\in\mathcal{T}}$ satisfies conditional consistency and strong stationarity. Consider $a \in \mathcal{T}$, $[p, \mathcal{S}, r], [q, \mathcal{S}, r] \in \mathcal{L}_a$, and $\Delta \geqslant 0$ such that $t \in \mathcal{S}$ implies $t \geqslant a$ and $t \in \mathcal{S}(p) \cup \mathcal{S}(q)$ implies $t + \Delta \in \mathcal{S}$. By strong stationarity, $[p, \mathcal{S}, r] \succcurlyeq_a [q, \mathcal{S}, r]$ if and only if $[p_\Delta, \mathcal{S}, r] \succcurlyeq_a [q_\Delta, \mathcal{S}, r]$. By foregone-risk independence, this is equivalent to $[p_\Delta, \mathcal{S}, r]\,|a \succcurlyeq_a [q_\Delta, \mathcal{S}, r]\,|a$. As above, we can choose $\tilde{r}$ such $[p_\Delta, \mathcal{S}, \tilde{r}]\,|(a+\Delta) = [p_\Delta, \mathcal{S}, \tilde{r}]\,|a$ and $[q_\Delta, \mathcal{S}, \tilde{r}]\,|(a+\Delta) = [q_\Delta, \mathcal{S}, \tilde{r}]\,|a$. By the first statement of this Proposition, $\{\succcurlyeq_a\}_{a\in\mathcal{T}}$ satisfies independence, and so this implies that $[p_\Delta, \mathcal{S}, r] \succcurlyeq_a [q_\Delta, \mathcal{S}, r]$ is equivalent to $[p_\Delta, \mathcal{S}, \tilde{r}]\,|(a+\Delta) \succcurlyeq_a [q_\Delta, \mathcal{S}, \tilde{r}]\,|(a+\Delta)$. By conditional consistency, this holds if and only if $[p_\Delta, \mathcal{S}, \tilde{r}] \succcurlyeq_{a+\Delta} [q_\Delta, \mathcal{S}, \tilde{r}]$. With a final application of independence, this is equivalent to $[p_\Delta, \mathcal{S}, r] \succcurlyeq_{a+\Delta} [q_\Delta, \mathcal{S}, r]$, and so strong time invariance holds, establishing statement 2.

Statement 3: Assume that $\{\succcurlyeq_a\}_{a\in\mathcal{T}}$ satisfies strong time consistency and foregone-risk independence and let $p, q \in \mathcal{L}_b$ and $a \leqslant b$. Because $p|b(x,t) = q|b(x,t) = 0$ for all $(x,t) \in \mathcal{X} \times [a,b)$, strong time consistency requires that $p|b \succcurlyeq_a q|b$ if and only if $p|b \succcurlyeq_b q|b$. By foregone-risk independence, $p|b \succcurlyeq_b q|b$ if and only if $p \succcurlyeq_b q$ and so conditional consistency holds. This establishes statement 3 of the Proposition.

Statement 4: Assume strong stationarity and strong time invariance. Let $a \in \mathcal{T}$ and $p, q \in \mathcal{L}_b$, with $p(x,t) = q(x,t)$ for all $(x,t) \in \mathcal{X} \times [a,b)$. Setting $\mathcal{S} = [b,\infty)$ can write $p = [p, \mathcal{S}, p]$ and $q = [q, \mathcal{S}, p]$. If $\mathcal{S} = [b,\infty)$, then $t \in \mathcal{S}(p) \cup \mathcal{S}(q)$ implies both $t \geqslant b$ and $t + b - a \in \mathcal{S}$. Then, by strong time invariance, $p \succcurlyeq_a q$ if and only $[p_{b-a}, \mathcal{S}, p] \succcurlyeq_b [q_{b-a}, \mathcal{S}, p]$ and, by strong stationarity, $[p_{b-a}, \mathcal{S}, p] \succcurlyeq_b [q_{b-a}, \mathcal{S}, p]$ if and only if $[p,\mathcal{S},p] \succcurlyeq_b [q,\mathcal{S},p]$. That is, strong time consistency holds.

Next, suppose that $\{\succcurlyeq_a\}_{a\in\mathcal{T}}$ satisfies strong time consistency and strong time invariance. Consider $a \in \mathcal{T}$, $[p,\mathcal{S},r], [q,\mathcal{S},r] \in \mathcal{L}_a$, and $\Delta \geqslant 0$ such that $t \in S$ implies $t \geqslant a$ and $t \in S(p) \cup S(q)$ implies $t + \Delta \in \mathcal{S}$. By strong time invariance, $[p,\mathcal{S},r] \succcurlyeq_a [q,\mathcal{S},r]$ if and only if $[p_\Delta,\mathcal{S},r] \succcurlyeq_{a+\Delta} [q_\Delta,\mathcal{S},r]$. We then have $[p_\Delta,\mathcal{S},r](x,t) = [q_\Delta,\mathcal{S},r](x,t) = 0$ for all (x,t) such that $t \in [a,a+\Delta) \cap S$ and also have $[p_\Delta,\mathcal{S},r](x,t) = [q_\Delta,\mathcal{S},r](x,t) = r(x,t)$ for all (x,t) such that $t \in [a,a+\Delta) \setminus S$. The prerequisites of strong time consistency hold, and so $[p_\Delta,\mathcal{S},r] \succcurlyeq_{a+\Delta} [q_\Delta,\mathcal{S},r]$ if and only if $[p_\Delta,\mathcal{S},r] \succcurlyeq_a [q_\Delta,\mathcal{S},r]$. That is, strong stationarity holds.

Finally, suppose that $\{\succcurlyeq_a\}_{a\in\mathcal{T}}$ satisfies strong time consistency and strong stationarity. Consider $a \in \mathcal{T}$, $[p,\mathcal{S},r], [q,\mathcal{S},r] \in \mathcal{L}_a$, and $\Delta \geqslant 0$ such that $t \in S$ implies $t \geqslant a$ and $t \in S(p) \cup S(q)$ implies $t + \Delta \in \mathcal{S}$. By strong stationarity, $[p,\mathcal{S},r] \succcurlyeq_a [q,\mathcal{S},r]$ if and only if $[p_\Delta,\mathcal{S},r] \succcurlyeq_a [q_\Delta,\mathcal{S},r]$. As above, we have $[p_\Delta,\mathcal{S},r](x,t) = [q_\Delta,\mathcal{S},r](x,t)$ for all (x,t) such that $t \in [a,a+\Delta)$, as each are equal to zero where $[a,a+\Delta)$ intersects S and are equal to $r(x,t)$ otherwise. By strong time consistency, $[p_\Delta,\mathcal{S},r] \succcurlyeq_a [q_\Delta,\mathcal{S},r]$ if and only if $[p_\Delta,\mathcal{S},r] \succcurlyeq_{a+\Delta} [q_\Delta,\mathcal{S},r]$. That is, strong time invariance holds, establishing Statement 4 and the Proposition.
■

References

[1] Abdellaoui, M., E. Diecidue, and A. Öncüler (2011). Risk preferences at different time periods: An experimental investigation. Management Science 57, 975-897.

[2] Anchugina, N. (2017). A simple framework for the axiomatization of exponential and quasi-hyperbolic discounting. Theory and Decision 82, 185-210.

[3] Attema, A., H. Bleichrodt, K. Rohde, and P.P. Wakker (2010). Time-tradeoff sequences for analyzing discounting and time inconsistency. Management Science 56, 2015-2030.

[4] Bastianello, L. and J. H. Faro (2022). Choquet expected discounted utility. Economic Theory, forthcoming.

[5] Benoit, J-P. and E. Ok (2007). Delay aversion. Theoretical Economics 2, 71-113.

[6] Binmore, K., A. Rubinstein and A. Wolinsky (1986). The Nash bargaining solution in economic modelling. Rand Journal of Economics, 17: 176-188.

[7] Debreu, G. (1954). Representation of a preference ordering by a numerical function. In Robert M. Thrall, Clyde H. Coombs, and Robert L. Davis (Eds.) Decision Processes, 159-165, Wiley, New York.

[8] Debreu, G. (1964). Continuity properties of Paretian utility. International Economic Review 5, 285-293.

[9] DeJarnette, P., D.Dillenberger, D. Gottlieb and P. Ortoleva (2020). Time lotteries and stochastic impatience. Econometrica 88, 619-656.

[10] Ebert, S. (2020). Decision making when things are only a matter of time. Operations Research 68(5), 1564-1575.

[11] Epstein, L. (1983). Stationary cardinal utility and optimal growth under uncertainty Journal of Economic Theory 31(1), 133-152.

[12] Fishburn, P. C. (1970). Utility theory for decision making. Wiley, New York.

[13] Fishburn, P. C. and A. Rubinstein (1982). Time preference. International Economic Review 23, 677-695.

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

[15] Gilboa, I., A. Postlewaite and L. Samuelson (2016). Memorable consumption. Journal of Economic Theory 165, 414-455.

[16] Grandmont, J.-M. (1972). Continuity properties of a von Neumann-Morgenstern utility. Journal of Economic Theory 4, 45-57.

[17] Halevy, Y. (2015). Time consistency: Stationarity and time invariance. Econometrica 83, 335-352.

[18] Hammond, P. (1988). Consequentialist foundations for expected utility. Theory and Decision 25, 25-78.

[19] Hayashi, T. (2003). Quasi-stationary cardinal utility and present bias. Journal of Economic Theory 112(2), 343-352.

[20] Jackson, M. and L. Yariv (2015). Collective dynamic choice: the necessity of time inconsistency. American Economic Journal: Microeconomics 7, 150-178.

[21] Kannappan, P. (2009). Functional Equations and Inequalities with Applications. Springer Monographs in Mathematics. Springer, US.

[22] Karni, E. and D. Schmeidler (1991). Atemporal dynamic consistency and expected utility theory. Journal of Economic Theory 54, 401-408.

[23] Kochov, A. (2015). Time and no lotteries: An axiomatization of maxmin expected utility. Econometrica 83(1), 239-262.

[24] Machina, M. (1987). Choice under uncertainty: Problems solved and unsolved. Journal of Economic Perspectives 1, 121-154.

[25] Millner, A. and G. Heal (2018). Time consistency and time invariance in collective intertemporal choice. Journal of Economic Theory 176, 158-169.

[26] Muthoo, A. (1999). Bargaining theory with applications. Cambridge University Press.

[27] Nachman, D. C. (1975). Risk aversion, impatience, and optimal timing decisions. Journal of Economic Theory 11, 196-246.

[28] Prakash, P. (1977). On the consistency of a gambler with time preference. Journal of Economic Theory 15, 92-98.

[29] Rubinstein, A. (1982). Perfect equilibrium in a bargaining model. Econometrica 50, 97-110.