

An Extension of Quasi-Hyperbolic Discounting to Continuous Time

Jinrui Pan^{*} Craig S. Webb[†] Horst Zank[‡]

December 3, 2014

Abstract

Two-Stage Exponential (TSE) discounting, the model developed here, generalises exponential discounting in a parsimonious way. It can be seen as an extension of Quasi-Hyperbolic discounting to continuous time. A TSE discounter has a constant rate of time preference before and after some threshold time; the switch point. If the switch point is expressed in calendar time, TSE discounting captures time consistent behaviour. If it is expressed in waiting time, TSE discounting captures time invariant behaviour. We provide preference foundations for all cases, showing how the switch point is derived endogenously from behaviour. We apply each case to Rubinstein's infinite-horizon, alternating-offers bargaining model.

JEL Codes: D74, D90

Keywords: Discounting, Present Bias, Decreasing Impatience, Bargaining.

^{*}Contact: jinrui.pan@manchester.ac.uk. This author's work was funded by a School of Social Sciences, Economics Discipline Area Studentship from the University of Manchester.

[†]Corresponding author. Contact: Economics, School of Social Sciences, University of Manchester, Manchester, M13 9PL, UK. Email: craig.webb@manchester.ac.uk. Tel: + 44 161 275 4805.

[‡]Economics, University of Manchester. Contact: horst.zank@manchester.ac.uk.

A decision maker’s rate of time preference is reflected in a variety of behaviours: smoothing consumption through saving, consuming or abstaining from tobacco, drugs or unhealthy food, investing in education, and so on. A constant rate of time preference, the predominant model of intertemporal choice in economics, rules out sudden changes in behaviour. Yet, people commonly resolve to start saving, quit smoking, eat better, and start exercising at some predetermined date. If utility is unchanged, it seems that their discount rate abruptly changes. This paper studies a model capturing this sudden change.

The *Quasi-Hyperbolic* (QH) discounting model elegantly captures a changing discount rate (Phelps and Pollak, 1968; Laibson, 1997; Hayashi, 2003; Attema, Bleichrodt, Rohde and Wakker, 2010; Olea and Strzalecki, 2013). Developed in discrete time, QH discounting involves weighting utility for outcomes using discount factors $\{1, \gamma\beta, \gamma\beta^2, \dots\}$. QH discounting has been applied extensively in economic theory (Asheim, 1997; Laibson, 1997; Barro, 1999; O’Donoghue and Rabin, 2001; Luttmer and Mariotti, 2003). Extending QH discounting to continuous time is important for economic applications. One approach has been considered by Harris and Laibson (2013). This paper studies an extension of QH discounting called *Two-Stage Exponential* (TSE) discounting. TSE discounting provides a more robust, in a way we will describe, approach to modelling present-biased preferences.

As with exponential, and many other nonexponential discounting models (see Abdellaoui, Attema and Bleichrodt, 2010: 849), TSE discounting retains a stationary instant utility for outcomes. This utility is discounted by a constant rate of time preference up to a *switch point*. After this point, the discount rate may change, but then remains constant. Violations of constant discounting occur only when comparing the near and distant future. This parametric form of discounting was first presented by Jamison and Jamison (2011). We provide a preference foundation for TSE discounting over timed outcomes.

We extend TSE discounting to dynamic choice by developing *time consistent* and *time invariant* (Halevy, 2012) versions of the model. It turns out that whether the

model captures time consistent or time invariant behaviour depends only on the interpretation of one parameter, the switch point. If the switch point is expressed in *calendar time*, then the model is time consistent. If it is expressed in *waiting time*, then the model is time invariant. We apply each dynamic version of TSE discounting to the infinite-horizon, alternating-offers bargaining model of Rubinstein (1982). There are few previous studies of non-exponential discounting preferences in sequential bargaining (Akin, 2007; Ok and Masatlioglu, 2007; Noor, 2011; Kodritsch, 2012), all of which have assumed time invariance.

The outline of this paper is as follows: Section 1 contains the notation and definitions. Section 2 presents the exponential discounting model, 2.1 as applied to static choice and 2.2 as applied to dynamic decision making. In Section 3.1 we present the TSE discounting model and in Section 3.2 we give a preference foundation for the model. Section 4 then considers extensions of the TSE discounting model to dynamic choice. The time consistent version of the model is presented in Section 4.1 and preference foundations are given. The time invariant version of the model is presented in Section 4.2 and, again, preference foundations are given. Then, the models are applied to infinite-horizon, alternating-offers bargaining in Sections 5.1 and 5.2. All proofs are in the Appendices.

1 Definitions

Let $[0, X]$, with $X > 0$, denote the set of *outcomes* and $[0, T]$, with $T > 0$, be the set of times at which an outcome can occur. The set of *timed outcomes* is $[0, X] \times [0, T]$. A typical element of $[0, X] \times [0, T]$ is (x, t) , which denotes the outcome x being received at time t . Such timed outcomes are the objects of choice.

A *static preference* relation $\succsim_{\underline{t}}$ is a binary relation defined over $[0, X] \times [\underline{t}, T]$; the set of timed outcomes occurring no sooner than time $\underline{t}$. A static preference relation characterises the preferences of our decision maker at time $\underline{t}$, as if they were making

decisions at that time.¹ An *initial preference* relation $\succsim_0$ is a static preference relation for $\underline{t} = 0$. For a set of *decision times* $\mathcal{D} \subseteq [0, T]$ with $0 \in \mathcal{D}$, a *dynamic preference structure* $\mathcal{R} := \{\succsim_{\underline{t}}\}_{\underline{t} \in \mathcal{D}}$ is a set of static preference relations indexed by $\mathcal{D}$.

Given a static preference $\succsim_{\underline{t}}$, the relations $\succ_{\underline{t}}$, $\preccurlyeq_{\underline{t}}$, $\prec_{\underline{t}}$ and $\sim_{\underline{t}}$ are defined in the usual way. A static preference $\succsim_{\underline{t}}$ is *complete* if, for all $(x, t), (x', t') \in [0, X] \times [\underline{t}, T]$, at least one of $(x, t) \succsim_{\underline{t}} (x', t')$ or $(x', t') \succsim_{\underline{t}} (x, t)$ holds. It is *transitive* if, for all $(x, t), (x', t'), (x'', t'') \in [0, X] \times [\underline{t}, T]$, $(x, t) \succsim_{\underline{t}} (x', t')$ and $(x', t') \succsim_{\underline{t}} (x'', t'')$ jointly imply $(x, t) \succsim_{\underline{t}} (x'', t'')$. It is a *weak order* if it is complete and transitive. It is *monotonic* if, for all $(x, t), (x', t) \in [0, X] \times [\underline{t}, T]$, $(x, t) \succsim_{\underline{t}} (x', t)$ iff $x \geq x'$. It is *impatient* if, for all $(x, t), (x, t') \in [0, X] \times [\underline{t}, T]$, $(x, t) \succsim_{\underline{t}} (x, t')$ iff $t' \geq t$. We will always assume that $(0, t) \sim_{\underline{t}} (0, t')$, for all $t, t' \in [0, T]$, and include this condition in the definition of impatience. A static preference relation $\succsim_{\underline{t}}$ is *continuous* if, for all $(x, t) \in [0, X] \times [\underline{t}, T]$, the sets $\{(x', t') : (x, t) \succsim_{\underline{t}} (x', t')\}$ and $\{(x', t') : (x, t) \preccurlyeq_{\underline{t}} (x', t')\}$ are closed subsets of $[0, X] \times [\underline{t}, T]$.

A static preference relation $\succsim_{\underline{t}}$ is *represented* by a real-valued function $V_{\underline{t}} : [0, X] \times [\underline{t}, T] \rightarrow \mathbb{R}$ if, for all $(x, t), (x', t') \in [0, X] \times [\underline{t}, T]$, the following holds:

$$(x, t) \succsim_{\underline{t}} (x', t') \Leftrightarrow V_{\underline{t}}(x, t) \geq V_{\underline{t}}(x', t').$$

A necessary condition for $\succsim_{\underline{t}}$ to admit such a representation is that $\succsim_{\underline{t}}$ is a weak order. By Debreu (1964), weak ordering and continuity of $\succsim_{\underline{t}}$ are necessary and sufficient for the existence of a *continuous* utility representation. Monotonicity and impatience ensure that such a representation is non-decreasing in x and non-increasing in t .

We call a set of functions $\mathcal{V} := \{V_{\underline{t}}\}_{\underline{t} \in \mathcal{D}}$, where $V_{\underline{t}} : [0, X] \times [\underline{t}, T] \rightarrow \mathbb{R}$ for each $\underline{t} \in \mathcal{D}$, a *dynamic model*. Finally, we say that a dynamic preference structure $\mathcal{R}$ is *represented* by a dynamic model $\mathcal{V}$ if, for all $\underline{t} \in \mathcal{D}$, the preference relation $\succsim_{\underline{t}} \in \mathcal{R}$ is represented by $V_{\underline{t}} \in \mathcal{V}$.

¹We underline the decision time, $\underline{t}$, as it becomes useful in presenting what follows.

2 Exponential Discounting

2.1 Static Choice and Exponential Discounting

This section reviews the classical exponential discounting model, as applied to initial or static choice over timed outcomes. Initial preferences conform to exponential discounting if they can be represented as follows:

$$V_0(x, t) = \delta^t u(x)$$

for all $(x, t) \in [0, X] \times [0, T]$, with $\delta \in [0, 1]$ and $u : [0, X] \rightarrow \mathbb{R}$ a continuous, strictly increasing function with $u(0) = 0$. The uniqueness properties pertaining to this representation are discussed later. The key property of exponential discounting is *stationarity*:

Definition (Stationarity): A static preference relation $\succsim_t$ satisfies *stationarity* if for all $(x, t), (y, t + \tau), (x, s), (y, s + \tau) \in [0, X] \times [t, T]$ the following holds:

$$(x, t) \succsim_t (y, t + \tau) \Leftrightarrow (x, s) \succsim_t (y, s + \tau).$$

The stationarity axiom asserts that a decision maker's preferences are unaffected by translations that preserve the time distance between two timed outcomes. This formulation of stationarity is due to Fishburn and Rubinstein (1982). Koopmans (1960, 1972) formulated the condition for sequences of outcomes. The following result, characterising exponential discounting preferences for timed outcomes, is due to Fishburn and Rubinstein (1982):

Theorem 2.1.1 (Fishburn & Rubinstein, 1982). *The following statements are equivalent:*

- (i) *The initial preference relation $\succsim_0$ over $[0, X] \times [0, T]$ is a continuous, monotonic and impatient weak order that satisfies stationarity.*

(ii) The initial preference relation $\succsim_0$ over $[0, X] \times [0, T]$ is represented by a function V_0 such that,

$$V_0(x, t) = \delta^t u(x)$$

for some $\delta \in [0, 1]$ and $u : [0, X] \rightarrow \mathbb{R}$ a continuous, strictly increasing function with $u(0) = 0$.

The following uniqueness results hold for Theorem 2.1.1:

Proposition 2.1.2 (Uniqueness Results). *Let the representation obtained in Theorem 2.1.1 hold for some $\delta \in (0, 1)$ and u . Then, δ and u are unique up to a joint power $\sigma > 0$ and factor $\tau > 0$ for u . That is, the map $(x, t) \mapsto \gamma^t v(x)$ represents $\succsim_0$, if and only if $\gamma = \delta^\sigma$ and $v = \tau u^\sigma$ with $\sigma, \tau > 0$.*

Note two important points regarding the above proposition. First, since the discount factor δ is not unique, one cannot assign behavioural content to its magnitude. For a discussion of the interpretation of the model's parameters, see Benoit and Ok (2007). Second, this lack of uniqueness is due to the timed outcomes framework we are using. For sequences of more than one non-zero outcome, one obtains uniqueness.

2.2 Dynamic Exponential Discounting

We now present the exponential discounting model as applied to dynamic choice. A dynamic preference structure $\mathcal{R}$ conforms to *dynamic exponential discounting* if it is represented by a dynamic model $\mathcal{V}$ where,

$$V_{\underline{t}}(x, t) = \delta^t u(x)$$

for all $V_{\underline{t}} \in \mathcal{V}$ and $(x, t) \in [0, X] \times [\underline{t}, T]$, with $\delta \in [0, 1]$ and $u : [0, X] \rightarrow \mathbb{R}$ a continuous, strictly increasing function with $u(0) = 0$. That is, every static preference relation $\succsim_{\underline{t}} \in \mathcal{R}$ is represented by exponential discounting. Further, they are all represented by the same exponential discounting function. As all static preference relations $\succsim_{\underline{t}} \in \mathcal{R}$ in

such a dynamic preference structure satisfy stationarity, we say the structure satisfies stationarity.

Dynamic exponential discounting assumes that the discount factor δ and utility function u can be taken to be the same at every decision time. Stationarity does not constrain the relationship between static preferences across different times. Something further is required. We consider two properties, *time consistency* and *time invariance*. These are such that, adding either of them to stationarity is sufficient for dynamic exponential discounting.

Time consistency requires that the decision maker's static preference between two timed outcomes does not depend on the decision time.

Axiom (Time Consistency): A dynamic preference structure $\mathcal{R}$ satisfies *time consistency* if for all $\succsim_{\underline{t}}, \succsim_{\underline{t}'} \in \mathcal{R}$, $(x, t), (y, s) \in [0, X] \times [0, T]$ such that $\underline{t}, \underline{t}' \leq \min\{t, s\}$:

$$(x, t) \succsim_{\underline{t}} (y, s) \Leftrightarrow (x, t) \succsim_{\underline{t}'} (y, s).$$

Halevy (2012) used the term *time invariance* for the following condition:

Axiom (Time Invariance): A dynamic preference structure $\mathcal{R}$ satisfies *time invariance* if, for all $\succsim_{\underline{t}}, \succsim_{\underline{t}+\tau} \in \mathcal{R}$, $(x, t), (y, s), (x, t+\tau), (y, s+\tau) \in [0, X] \times [0, T]$ such that $\underline{t} \leq \min\{t, s\}$ and $\tau \geq 0$:

$$(x, t) \succsim_{\underline{t}} (y, s) \Leftrightarrow (x, t+\tau) \succsim_{\underline{t}+\tau} (y, s+\tau).$$

Time invariance captures the behaviour of a decision maker who evaluates timed outcomes in waiting time. That is, only the delay between the decision time and the outcome time matters; not the calendar time of the outcome. Time invariant sets of preference relations may fail to exhibit time consistency. The following theorem is due to Halevy (2012).

Theorem 2.2.1 (Halevy, 2012). *Let $\mathcal{R} := \{\succsim_t\}_{t \in \mathcal{D}}$ be a dynamic preference structure.*

Let $\mathcal{D} = [0, T]$. If $\mathcal{R}$ satisfies any two of stationarity, time consistency and time invariance, then it satisfies the remaining condition.

Theorem 2.2.1 can be used to give various preference foundations for exponential discounting in the dynamic framework. This is achieved by combining the usual conditions with any two of the three conditions for $\mathcal{R}$.

3 Two-Stage Exponential Discounting

3.1 The Two-Stage Exponential Discounting Model

Quasi-hyperbolic (QH), in its original, discrete-time formulation, involves weighting utility for outcomes using discount factors $\{1, \gamma\beta, \gamma\beta^2, \dots\}$. This paper is concerned with extending QH discounting to continuous time. Two possible ways are *Split-Function Exponential* (SFE) discounting and *Two-Stage Exponential* (TSE) discounting. SFE discounting incorporates a discontinuity in the discount *function*. TSE discounting incorporates a discontinuity in the discount *factor*. Harris and Laibson (2013) proposed SFE discounting and Jamison and Jamison (2011, p.40) advocated SFE over TSE discounting. We will show in this Section that the behavioural implications of TSE discounting are more natural than of SFE discounting.

A SFE discounter has initial preferences $\succsim_0$ over $[0, X] \times [0, T]$ represented by $V_0 : [0, X] \times [0, T] \rightarrow \mathbb{R}$ where,

$$V_0(x, t) = \begin{cases} \beta^t u(x) & \text{if } t \leq \lambda \\ \gamma \beta^t u(x) & \text{if } t > \lambda \end{cases}$$

for all $(x, t) \in [0, X] \times [0, T]$, with $\lambda \in [0, T]$, $\beta \in [0, 1]$, $\gamma \geq 0$, and $u : [0, X] \rightarrow \mathbb{R}$ a continuous, strictly increasing function with $u(0) = 0$. Under SFE discounting, outcomes occurring after a subjective “present,” i.e. after $[0, \lambda]$, are affected by a fixed factor γ .

Initial preferences are said to exhibit *immediacy bias* if there are $x, y \in [0, X]$, $\tau \in [0, T]$ and $\Delta > 0$ such that $(x, 0) \succ_0 (y, \Delta)$ and $(x, \tau) \prec_0 (y, \tau + \Delta)$. The delay Δ , for outcome y over x , becomes acceptable when the timed outcomes are translated τ units into the future. Immediacy bias, often called present bias, formalises the example of Thaler (1981), which suggests that one who prefers (one apple, 0) to (two apples, 1 day), will often prefer (two apples, 366 days) to (one apple, 365 days). Such preferences are incompatible with exponential discounting.

SFE discounting can explain Thaler's (1981) immediacy bias example. To do so, however, forces either $\lambda \in [0, 1]$, or $\lambda \in [365, 366]$ to hold. These are restrictive requirements. As such, SFE discounting cannot simultaneously explain minor adaptations of immediacy bias. If $\lambda \in [0, 1]$, then immediacy bias with a front-end delay of anything larger than one day cannot be explained. If $\lambda \in [365, 366]$, then immediacy bias with translation τ less than 364 days, or greater than 366 days cannot be explained by SFE discounting.

Two-Stage Exponential (TSE) discounting also incorporates subjectivity to the time period called "present". A TSE discounter has initial preferences $\succsim_0$ over $[0, X] \times [0, T]$ represented by $V_0 : [0, X] \times [0, T] \rightarrow \mathbb{R}$ where,

$$V_0(x, t) = \begin{cases} \alpha^t u(x) & \text{if } t \leq \lambda \\ (\alpha/\beta)^\lambda \beta^t u(x) & \text{if } t > \lambda \end{cases}$$

for all $(x, t) \in [0, X] \times [0, T]$, with $\lambda \in [0, T]$, $\alpha, \beta \in [0, 1]$ and $u : [0, X] \rightarrow \mathbb{R}$ a continuous, strictly increasing function with $u(0) = 0$. We call λ the *switch point*. It is a threshold, a demarcation point that separates periods before and after a change in attitudes. When evaluating timed outcomes that occur before λ , the decision maker maximises an exponential discounting function with discount factor α . For timed outcomes occurring later than λ , the decision maker exponentially discounts the same utility function, but uses discount factor β . The weight $(\alpha/\beta)^\lambda$ ensures that the evaluation function is continuous everywhere.

We say that a TSE discounting function, with $\lambda \in (0, T)$, exhibits *decreasing impatience* if $\alpha < \beta$ and *increasing impatience* if $\alpha > \beta$. For TSE discounting to exhibit immediacy bias, decreasing impatience is necessary. The switch point λ , however, can be *any* time between today and Δ -from-today. Front-end delays of the immediacy bias condition are simultaneously accommodated, as are changes in τ to shorter or longer translations. If immediacy bias is adapted, keeping the earlier timed outcomes before λ and the later timed outcomes are after λ , then TSE discounting with decreasing impatience will not be contradicted.

3.2 A Preference Foundation

This section provides a preference foundation for the TSE discounting model. We present our result in the framework of initial choice over timed-outcomes. Our approach allows the switch point to be detected from observed behaviour.² First, two concepts are introduced: *stationarity-after-t* and *stationarity-before-t*.

Definition (Stationarity-after-t): A preference relation $\succsim_0$ satisfies *stationarity-after-t* if for all $(x, t), (y, t + \tau), (x, s), (y, s + \tau) \in [0, X] \times [0, T]$ with $\tau > 0$ and $s > t$, the following holds:

$$(x, t) \succsim_0 (y, t + \tau) \Rightarrow (x, s) \succsim_0 (y, s + \tau).$$

Stationarity-after- t demands that, when comparing two timed outcomes with the *soonest* outcome occurring at time t , preferences are invariant under translations that put each outcome backward in time by the same amount. Note that it is a one way implication; the preference regarding the earlier timed outcomes implying the preference regarding the later timed outcomes. TSE discounting preferences satisfy stationarity-after- t when $t \geq \lambda$.

²We avoid using “there exists” in our axiom. A similar technique was employed by Diecidue, Schmidt and Zank (2009) to provide a preference foundation for an inverse-S shaped probability weighting function.

Definition (Stationarity-before- t): A preference relation $\succsim_0$ satisfies *stationarity-before- t* if for all $(x, t), (y, t - \tau), (x, s), (y, s - \tau) \in [0, X] \times [0, T]$ with $0 < \tau < s < t$, the following holds:

$$(x, t) \succsim_0 (y, t - \tau) \Rightarrow (x, s) \succsim_0 (y, s - \tau).$$

When comparing two timed outcomes, with the *latest* outcome occurring at time t , these preferences are invariant under translations that bring each outcome forward in time by the same amount. One may also verify, by substitution of the preference functional, that TSE discounting preferences satisfy stationarity-before- t when $t \leq \lambda$.

Suppose we observe a violation of stationarity-after- t for some t . For TSE discounters, *this can only happen because their switch point is later than t* . Such an observation tells us, when trying to detect the switch point, *we do not need to look before t* as stationarity-before- t must hold. Analogously, a violation of stationarity-before- t' , for some t' , confirms that the switch point is not later than t' . Then, stationarity-after- t' must hold. Indeed, for *all* times $t \in [0, T]$, if one of these conditions does not hold, the other necessarily holds. This is the content of the *two-stage stationarity* axiom:

Axiom (Two-Stage Stationarity): For all times $t \in [0, T]$, preferences $\succsim_0$ are stationary-before- t , or stationary-after- t , or both.

Two-stage stationarity is not strong enough to ensure a time-independent utility function for outcomes.³ We use the following condition, *midpoint consistency*:

³To see this, consider an initial preference represented by,

$$V_0(x, t) = \begin{cases} \alpha^t u(x) & \text{if } t \leq \lambda \\ \phi(\beta^t v(x)) & \text{if } t > \lambda \end{cases}$$

with $\alpha, \beta \in (0, 1)$, $\lambda \in [0, T]$, $u : [0, X] \rightarrow \mathbb{R}$ continuous and strictly increasing, and $\phi : \mathbb{R} \rightarrow \mathbb{R}$ a continuous, strictly increasing function. Such preferences necessarily satisfy two-stage stationarity, but need not be TSE discounting preferences. The stationarity axiom would ensure that we could take ϕ as the identity function and that *both* $u = v$ and $\alpha = \beta$ hold. Two-stage stationarity implies *none* of this. Continuity of V_0 does ensure that $\alpha^\lambda u(x) = \phi(\beta^\lambda v(x))$ for all $x \in [0, X]$, therefore the utilities are ordinally equivalent. Beyond this, no more can be said.

Axiom (Midpoint Consistency): An initial preference relation $\succsim_0$ satisfies *midpoint consistency* (or “ $\succsim_0$ is *midpoint consistent*”) if, for all $x, y, z \in [0, X]$ and $t, t', s, s' \in [0, T]$, any three of the following indifferences imply the fourth:

$$\begin{aligned} (x, t) \sim_0 (y, t') & \quad (x, s) \sim_0 (y, s') \\ (y, t) \sim_0 (z, t') & \quad (y, s) \sim_0 (z, s') \end{aligned}$$

Midpoint consistency captures the idea that we may consistently measure utility ratios. Further, these utility ratios should not depend on which points in time are used. Köbberling and Wakker (2003) presented such a condition. Similar techniques are discussed in Baillon, Driesen and Wakker (2012). The following theorem provides the preference foundation for the TSE discounting model:

Theorem 3.2.1. *The following statements are equivalent:*

- (i) *The initial preference relation $\succsim_0$ over $[0, X] \times [0, T]$ is a continuous, monotonic, impatient and midpoint consistent weak order that satisfies two-stage stationarity.*
- (ii) *The initial preference relation $\succsim_0$ over $[0, X] \times [0, T]$ is represented by a function V_0 such that,*

$$V_0(x, t) = \begin{cases} \alpha^t u(x) & \text{if } t \leq \lambda \\ (\alpha/\beta)^\lambda \beta^t u(x) & \text{if } t > \lambda \end{cases}$$

for some $\alpha, \beta \in [0, 1]$, $\lambda \in [0, T]$ and a continuous, strictly increasing $u : [0, X] \rightarrow \mathbb{R}$ with $u(0) = 0$.

The following Proposition outlines the uniqueness results pertaining to Theorem 3.2.1. A switch point is *meaningful* if it is in $(0, T)$ and $\alpha \neq \beta$.

Proposition 3.2.2 (Uniqueness Results). *Let the representation obtained in Theorem 3.2.1 hold for some $\alpha, \beta \in (0, 1)$, $\lambda \in [0, T]$ and $u : [0, X] \rightarrow \mathbb{R}$. If a switch point*

is not meaningful, then $\lambda \in \{0, T\}$ or $\alpha = \beta$, so TSE discounting collapses to exponential discounting and the uniqueness results expressed in Proposition 2.1.2 hold. Now consider a meaningful switch point. In this case, the switch point is uniquely determined. Then, α , β and u are unique up to a joint power $\sigma > 0$ and factor $\tau > 0$ for u .

The proofs for Theorem 3.2.1 and Proposition 3.2.2 are in Appendices A.1 and A.2 respectively.

4 Dynamic TSE Discounting Models

When we consider dynamic preference structures, the distinction between calendar and waiting time is important. If the switch point is a point in calendar time λ then, eventually, it will pass. When this happens, TSE discounting collapses to exponential discounting. Suppose, on the other hand, that the switch point is expressed in waiting time. In this case, it remains λ units in the future at all decision times. Such a TSE discounter will never become an exponential discounter. These two interpretations lead to time consistent or time inconsistent models of dynamic choice. We will consider these possibilities, calendar time and waiting time, in Sections 4.1 and 4.2 respectively. In Sections 5.1 and 5.2 we study how each of these interpretations affects the application of TSE discounting to sequential bargaining.

4.1 Consistent TSE Discounting

This section presents a model we call *consistent TSE discounting* (CTSE). Suppose that our decision maker has TSE discounting initial preferences, as in Theorem 3.2.1. For a CTSE preference structure, we ask that the switch point be expressed in calendar time. Formally, our decision maker is a CTSE maximiser if their dynamic

preference structure $\mathcal{R}$ is represented by a dynamic model $\mathcal{V}$ where,

$$V_{\underline{t}}(x, t) = \begin{cases} \alpha^t u(x) & \text{if } t \leq \lambda \\ (\alpha/\beta)^\lambda \beta^t u(x) & \text{if } t > \lambda \end{cases}$$

for all $V_{\underline{t}} \in \mathcal{V}$ and $(x, t) \in [0, X] \times [\underline{t}, T]$, with $\lambda \in [0, T]$, $\alpha, \beta \in [0, 1]$ and $u : [0, X] \rightarrow \mathbb{R}$ a continuous, strictly increasing function with $u(0) = 0$. For CTSE representations, α, β, λ and u do not vary with the decision time. Each $V_{\underline{t}}$ is the restriction of the initial representing function V_0 to timed outcomes occurring later than $\underline{t}$. Since λ is expressed in calendar time units, CTSE preference structures exhibit time consistency.

Theorem 2.2.1 tells us that this model cannot satisfy time invariance without collapsing to exponential discounting. We will now dispense with two-stage stationarity and consider an axiom that applies to the dynamic preference structure. Our new condition here is *two-stage time invariance*. We must first introduce the notions of *time-invariance-before-t* and *time-invariance-after-t*. The former is as follows:

Definition (Time-Invariance-before-t): A dynamic preference structure $\mathcal{R}$ satisfies *time-invariance-before-t* if, for all $(x, t'), (y, t''), (x, t' - \tau), (y, t'' - \tau) \in [0, X] \times [0, T]$ with $0 \leq \tau \leq \underline{t} < t', t'' \leq t$, the following holds:

$$(x, t') \succsim_{\underline{t}} (y, t'') \Rightarrow (x, t' - \tau) \succsim_{\underline{t} - \tau} (y, t'' - \tau).$$

Time-invariance-before- t of a dynamic preference structure asks that, if we observe a preference between two outcomes with the latest occurring before time t , then the preference does not change if we bring both outcomes and the decision time forward in time by the same amount. Notice that this is a one-way implication; the later preference implying the earlier. The analogous condition, *time-invariance-after-t*, is defined as follows:

Definition (Time-Invariance-after-t): A dynamic preference structure $\mathcal{R}$ satis-

fies *time-invariance-after- t* if, for all $(x, t'), (y, t''), (x, t' + \tau), (y, t'' + \tau) \in [0, X] \times [0, T]$ with $0 \leq \tau$ and $\underline{t} \leq t \leq t', t''$, the following holds:

$$(x, t') \succsim_{\underline{t}} (y, t'') \Rightarrow (x, t' + \tau) \succsim_{\underline{t} + \tau} (y, t'' + \tau).$$

Time-invariance-after- t of a dynamic preference structure asks that, if we observe a preference between two outcomes with the earliest occurring after time t , then the preference does not change if we delay both outcomes and the decision time by the same amount.

Under CTSE discounting, preferences must satisfy time-invariance-before- λ and time-invariance-after- λ . We need not know λ *a priori*. Our axiom is formulated such that the existence of λ and the appropriate representation are implied. We offer the following, *two-stage time invariance*:

Axiom (Two-Stage Time Invariance): For all times $t \in [0, T]$, the set of preference relations $\mathcal{R}$ satisfies time-invariance-before- t , or time-invariance-after- t , or both.

Two-stage time invariance provides a testable condition that must hold for *all* t in $[0, T]$. The following theorem asserts the equivalence of our previous axiom set, using two-stage stationarity, and an axiom set using two-stage time invariance.

Theorem 4.1.1. *Let the set of decision times $\mathcal{D} = [0, T]$. Then, the following statements are equivalent:*

- (i) *The initial preference relation $\succsim_0$ over $[0, X] \times [0, T]$ is a continuous, monotonic, impatient and midpoint consistent weak order that satisfies two-stage stationarity and the set of preference relations $\mathcal{R}$ satisfies time consistency.*
- (ii) *The initial preference relation $\succsim_0$ over $[0, X] \times [0, T]$ is a continuous, monotonic, impatient and midpoint consistent weak order and the set of preference relations $\mathcal{R}$ satisfies time consistency and two-stage time invariance.*

(iii) Each preference relation $\succsim_{\underline{t}} \in \mathcal{R}$ can be represented by a function $V_{\underline{t}}$ where:

$$V_{\underline{t}}(x, t) = \begin{cases} \alpha^t u(x) & \text{if } t \leq \lambda \\ (\alpha/\beta)^\lambda \beta^t u(x) & \text{if } t > \lambda \end{cases}$$

for some $\alpha, \beta \in [0, 1]$, $\lambda \in [0, T]$ and a continuous, strictly increasing $u : [0, X] \rightarrow \mathbb{R}$ with $u(0) = 0$.

The uniqueness results pertaining to Theorem 4.1.1 are the same as those in Proposition 3.2.2, applied at each decision time. The proof of Theorem 4.1.1 is in Appendix A.3.

4.2 Invariant TSE Discounting

This section presents our second dynamic model. We call this model *invariant two-stage discounting* (ITSE). In this case, we think of the switch point as being expressed in waiting time. Formally, our decision maker is an ITSE discounter if their dynamic preference structure $\mathcal{R}$ is represented by a dynamic model $\mathcal{V}$ where,

$$V_{\underline{t}}(x, t) = \begin{cases} \alpha^t u(x) & \text{if } t \leq \underline{t} + \lambda \\ (\alpha/\beta)^{\underline{t} + \lambda} \beta^t u(x) & \text{if } t > \underline{t} + \lambda \end{cases}$$

for all $V_{\underline{t}} \in \mathcal{V}$ and $(x, t) \in [0, X] \times [\underline{t}, T]$, with $\lambda \in [0, T]$, $\alpha, \beta \in [0, 1]$ and $u : [0, X] \rightarrow \mathbb{R}$ a continuous, strictly increasing function with $u(0) = 0$.

The main difference between ITSE and CTSE discounting is that the decision time $\underline{t}$ now plays an important role. At each decision time, an ITSE discounter will evaluate timed outcome (x, t) using discount factor α if t is less than λ units from the current decision time ($t \leq \underline{t} + \lambda$). They will use discount factor β if t is more than λ units in front of the current decision time ($t > \underline{t} + \lambda$). Such preferences are not time consistent (unless $\alpha = \beta$) because the evaluation of a timed outcome (x, t) changes with the decision time. They are, however, time invariant.

We may use time invariance, combined with our initial conditions for TSE discounting, to characterise ITSE discounting. We will provide a foundation for ITSE discounting, with the key axioms applying to the dynamic preference structure. We introduce a relaxation of time consistency, *two-stage time consistency*, which permits violations of time consistency only when comparing timed outcomes that are subjectively far enough apart in time.

Before presenting our two-stage time consistency axiom, it is necessary to develop two conditions: *Time-consistency-within-t-from-now* and *time-consistency-beyond-t-from-now*. The former is as follows:

Definition (Time-consistency-within-t-from-now): A set of preference relations $\mathcal{R}$ satisfies *time-consistency-within-t-from-now* if for all $\succsim_{\underline{t}} \in \mathcal{R}$ and all $(x, t), (y, t') \in [0, X] \times [0, T]$, with $0 \leq \underline{t} \leq t' \leq t$, the following holds:

$$(x, t) \succsim_0 (y, t') \Rightarrow (x, t) \succsim_{\underline{t}} (y, t').$$

Time-consistency-within-t-from-now of a dynamic preference structure demands that, for timed outcomes occurring before t , initial preferences are not later reversed. The initial preference relation $\succsim_0$ plays a key role in this axiom. Preference relations at later points in time, $\succsim_{\underline{t}}$ with $\underline{t} \leq t$, are forced to agree with the initial preference relation. This is a one-way implication, with initially expressed preferences implying the later ones.

ITSE discounting preference structures must satisfy time-consistency-within-t-from-now whenever $t \leq \lambda$. This condition need not hold for ITSE discounters when $t > \lambda$, but in this case there is an analogous condition:

Definition (Time-consistency-beyond-t-from-now): A set of preference relations $\mathcal{R}$ satisfies *time-consistency-beyond-t-from-now* if for all $\succsim_{\underline{t}} \in \mathcal{R}$ and all $(x, t + \underline{t}), (y, t' + \underline{t}) \in [0, X] \times [0, T]$ with $t \leq t'$ the following holds:

$$(x, t + \underline{t}) \succsim_{\underline{t}} (y, t' + \underline{t}) \Rightarrow (x, t + \underline{t}) \succsim_0 (y, t' + \underline{t}).$$

Time-consistency-beyond- t -from-now of a dynamic preference structure demands that, for timed outcomes occurring no sooner than t after decision time $\underline{t}$, preferences expressed are respected by initial preferences. Again, it is a one-way implication with the later preference implying the earlier, initial preference. Time-consistency-beyond- t -from-now, when combined with time invariance, will be equivalent to stationarity-after- t . The distance between decision time and the timed outcomes increases when comparing the first and second preference expressions. That is, at time decision $\underline{t}$ a timed outcome occurring at time $t + \underline{t}$ is “ t -from-now”. As “now” becomes earlier, back to time zero, the same timed outcome is beyond- t -from-now. Under time invariance, the former preference implies $(x, t) \succsim_0 (y, t')$. We see that the implied preference above completes the requirement for stationarity-after- t . We are now set to define the characteristic axiom for ITSE discounting.

Axiom (Two-Stage Time Consistency): For all times $t \in [0, T]$, the set of preference relations $\mathcal{R}$ satisfies time-consistency-within- t -from-now, or time-consistency-beyond- t -from-now, or both.

We have explained how time-consistency-beyond- t -from-now, under the assumption of time invariance of $\mathcal{R}$, is equivalent to stationarity-after- t . Under the same assumptions, time-consistency-within- t -from-now of $\mathcal{R}$ can be shown to be equivalent stationarity-before- t of $\mathcal{R}$. Two-stage time consistency, although by itself completely distinct from two-stage stationarity, must be equivalent to two-stage stationarity for time invariant structures. This is summarised in the following theorem:

Theorem 4.2.1. *Let the set of decision times $\mathcal{D} = [0, T]$. Then, the following statements are equivalent:*

- (i) *The initial preference relation $\succsim_0$ over $[0, X] \times [0, T]$ is a continuous, monotonic, impatient and midpoint consistent weak order that satisfies two-stage stationarity and the set of preference relations $\mathcal{R}$ satisfies time invariance.*
- (ii) *The initial preference relation $\succsim_0$ over $[0, X] \times [0, T]$ is a continuous, monotonic,*

impatient and midpoint consistent weak order and the set of preference relations $\mathcal{R}$ satisfies time invariance and two-stage time consistency.

(iii) *Each preference relation $\succsim_{\underline{t}} \in \mathcal{R}$ can be represented by a function $V_{\underline{t}}$ where:*

$$V_{\underline{t}}(x, t) = \begin{cases} \alpha^t u(x) & \text{if } t \leq \underline{t} + \lambda \\ (\alpha/\beta)^{t+\lambda} \beta^t u(x) & \text{if } t > \underline{t} + \lambda \end{cases}$$

for some $\alpha, \beta \in [0, 1]$, $\lambda \in [0, T]$ and a continuous, strictly increasing $u : [0, X] \rightarrow \mathbb{R}$ with $u(0) = 0$.

The uniqueness results pertaining to Theorem 4.2.1 are the same as those in Proposition 3.2.2, applied at each decision time. The proof of Theorem 4.2.1 is in Appendix A.4.

5 Bargaining

This section gives an application of the CTSE and ITSE discounting models. For this purpose, we consider the infinite-horizon alternating-offers bargaining model of Rubinstein (1982). Rubinstein's model provides clear predictions under exponential discounting. The basic framework is highly adaptable and has been used widely in economics.

The game $\mathcal{G}$ is as follows. There are two players, 1 and 2, and a surplus of (normalised) size 1. The players have exponential discounting (ED) preferences with linear utility, that is, an outcome of x at time t gives utility at time zero of $\delta_i^t x$, where $\delta_i \in (0, 1)$ is player i 's discount factor. The players alternate in proposing and considering offers regarding how the surplus should be divided. Player 1 proposes first at $t = 0$ and player 2 may accept or reject the proposal. If player 2 accepts the proposal, the game ends at that point and the payoffs are those specified in player 1's offer. If player 2 rejects the proposal, the game continues to $t = \Delta$ at which time

the players' previous roles are exchanged. The Δ parameter is called the *bargaining delay*. The friction caused by the bargaining delay, in particular the players' aversion to delayed payoffs, is central in Rubinstein's analysis. The game continues, perhaps indefinitely, with players 1 and 2 making offers at time $k\Delta$, for even and odd k respectively. Indefinite disagreement yields zero payoffs for both players.

Rubinstein (1982) showed that $\mathcal{G}$ has a unique subgame perfect equilibrium (SPE). The SPE prescribes an immediate agreement. At time 0, player 1 suggests $x = \frac{1-\delta_2^\Delta}{1-\delta_1^\Delta\delta_2^\Delta}$ for himself and $1-x$ for player 2, which player 2 accepts and the game ends.

5.1 CTSE Discounting in Bargaining

Suppose we take any infinite-horizon game where the players are CTSE discounters. Then *eventually* (that is, after the last of the players' switch points) the game will become an infinite-horizon game where each player has exponential discounting preferences. One may then use existing results to find the subgame perfect Nash equilibrium (SPE) of this subgame. Then, one may find the SPE of the overall game by: truncating the game at the latest switch point, assigning each player their (sub-) SPE payoff instead, and then solving the truncated game by backward induction. We apply this technique to bargaining.

Consider a game $\mathcal{G}'$ that is the same as $\mathcal{G}$, except that player 2 is a CTSE discounter with $\lambda = 1$ and discount factors $\alpha, \beta \in (0, 1)$. For simplicity, we let $\Delta = 1$. If the game were to continue to $t = 2$, then the remaining game is a standard Rubinstein bargaining model with player 1 as first mover. The game starting at $t = 2$ has a unique SPE with payoffs $(\frac{1-\beta}{1-\delta_1\beta}, \frac{\beta(1-\delta_1)}{1-\delta_1\beta})$. To solve the game $\mathcal{G}'$ we construct a truncated game $\mathcal{G}'(2)$ with the following structure. $\mathcal{G}'(2)$ is the same as $\mathcal{G}'$ at times 0 and 1. If the game reaches $t = 1$ and player 1 rejects player 2's offer, instead of continuing to $t = 2$ the game ends and players 1 and 2 are assigned payoffs $\frac{1-\beta}{1-\delta_1\beta}$ and $\frac{\beta(1-\delta_1)}{1-\delta_1\beta}$, respectively.

A payoff profile $(x, 1-x)$, paid immediately, occurs as a SPE of $\mathcal{G}'$ if and only

if it occurs as a SPE of $\mathcal{G}'(2)$. The truncated game $\mathcal{G}'(2)$ is solved by backward induction. At $t = 1$, by standard arguments, player 2 should offer player 1 a share of $x^1 = \delta_1 \frac{1-\beta}{1-\delta_1\beta}$. At $t = 0$ player 1 should offer player 2 a share of $1 - x^0 = \alpha(1 - x^1)$ implying x^0 for himself. This is immediately accepted by player 2. The unique SPE involves immediate agreement and payoffs $(x^0, 1 - x^0)$ where:

$$x^0 = \frac{1 - [\delta_1\beta + (1 - \delta_1)\alpha]}{1 - \delta_1\beta}.$$

Notice that $x^0 = \frac{1-\delta_2}{1-\delta_1\delta_2}$ if and only if $\alpha = \beta = \delta_2$. The agreement is still reached immediately, however the SPE is different as the incentives to delay agreement are different. The change in discount rate that occurs after player 2's switch point affects the SPE payoffs in a predictable way.

A well-known result in bargaining under exponential discounting is that, other things being equal, greater impatience leads to worse equilibrium payoffs. One can construct examples where the player who is more impatient *at the time of agreement* does better. In such cases the player in question will, at some point, become less impatient than his opponent. This fact must, therefore, be integrated into the determination of the equilibrium payoffs, even though the equilibrium prescribes agreement without delay.

5.2 ITSE Discounting in Bargaining

A time inconsistent decision maker has different preferences at different decision times. Effectively, they are a different decision maker at each decision time. Although the decision maker is a collection of selves, the “decision time self” decides how to treat the other selves. Three strategies advanced in the literature are *naive*, *resolute* and *sophisticated* choice (Blackorby, Nissen, Primont and Russell, 1973; Hammond, 1976; Machina, 1989; McClennen, 1990; O'Donogue and Rabin, 2001; Hey and Lotito, 2009; Hey and Panaccione, 2011). We address each of these as

applied to ITSE discounting and sequential bargaining.

The naive approach to address time inconsistency is, essentially, to ignore it. That is, the decision time self acts as if they are ignorant of the fact that their preferences will change. In general, this leads to different outcomes to time consistent choice. In the case of the Rubinstein bargaining game, however, the outcome will be as the time consistent case. Since the naive subgame perfect equilibrium prescribes immediate agreement, the players will never learn of their preferences changing.

A *resolute* decision maker (McClennen, 1990) achieves time consistent behaviour by sticking to their initial plan, regardless of their evolving preferences. As with the naive case, agreement occurs immediately, hence we need not appeal to any form of commitment device. In this application, the difference between naive and resolute choice is only the underlying reasoning. A naive decision maker does not know his preferences will change; a resolute decision maker does not care.

The final case we consider is *sophisticated* choice. In this case, the time inconsistent decision makers are fully aware of their future preferences, fully anticipate the optimal choices of their future selves, and integrate these into their current strategy. We will assume, for ease of exposition, that player 1 has exponential discounting preferences and player 2 is a sophisticated, ITSE discounter.

We call a strategy profile a *myopic equilibrium* if there are no preferable single period deviations. This solution concept is sophisticated, players recognise their changing preferences. But, only single periods matter for the solution and, under time invariance, attitudes to single period delays never change.

Suppose that the bargaining delay Δ is larger than player 2's λ . Then, every time player 2 has to make a decision of accept or reject now, time invariance ensures they must be using discount factor β . It must be that the myopic equilibrium is the same if we construct a new game, replacing player 2 with an exponential discounter with discount factor β . Alternatively, if the delay Δ is smaller than λ , then the myopic equilibrium of the game will be the same if we replace player 2 with an exponen-

tial discounter with discount factor α . The myopic equilibrium strategies are those prescribed by Rubinstein's analysis, using one of the two discount factors, depending on the bargaining delay. The myopic equilibrium outcome involves immediate agreement and payoffs $(x^0, 1 - x^0)$ where:

$$x^0 = \begin{cases} \frac{1-\alpha^\Delta}{1-\delta^\Delta\alpha^\Delta} & \text{if } \Delta \leq \lambda \\ \frac{1-\beta^\Delta}{1-\delta^\Delta\beta^\Delta} & \text{if } \Delta > \lambda \end{cases}$$

The outcome in this case is sensitive to the specification of the game. In particular, the payoffs do not vary continuously with the bargaining delay. There is a pronounced jump in the equilibrium partition as Δ becomes longer or shorter than the period before the switch point.

It remains to justify the use of this solution concept for the bargaining game under consideration. One can consider the bargaining game with sophisticated, time inconsistent players as a game with an infinite number of players. *Strotz-Pollak equilibrium* is the subgame perfect Nash equilibrium of that game. As such, Strotz-Pollak equilibrium is the appropriate solution concept for sophisticated choice. For more on Strotz-Pollak equilibrium, see Shefrin (1998). The myopic equilibrium strategies form a Strotz-Pollak equilibrium here. Adhering to the strategy always guarantees immediate agreement. Deviating in any single round leads to agreement in either one or two rounds which is, at best, no better.

Although the myopic equilibrium is unique, and is a Strotz-Pollak equilibrium, there may be other Strotz-Pollak equilibria. Theorem 4 of Kodritsch (2012) provides a sufficient condition for the unique myopic equilibrium we derived to be the unique Strotz-Pollak equilibrium of the bargaining game. A sufficient condition is a weak present-bias condition, equivalent here to constant or decreasing impatience, $\alpha \leq \beta$.

Increasing impatience allows for delayed agreement as an equilibrium. Intuitively, if one expects one's opponent to "lose his cool" after a disagreement, the strategic advantage for delay is apparent. Although it is not the predominant finding, increasing

impatience has been observed in individual choice experiments (Attema, Bleichrodt, Rohde and Wakker, 2010; Epper, Fehr-Duda and Bruhin, 2011; Takeuchi, 2011; Abdellaoui, Bleichrodt and l’Haridon, 2013).

6 Concluding Comments

We have presented two-stage exponential (TSE) discounting, provided an axiomatic foundation, extended it to the dynamic framework in two different ways, and in each dynamic interpretation provided two different axiomatic foundations. These characterisations provide simple, testable conditions (two-stage stationarity, two-stage time consistency and two-stage time invariance) that merit empirical study. We have demonstrated how CTSE and ITSE discounting can be applied to bargaining theory, in particular to the case with time between bargaining rounds being a continuous variable. We expect that TSE discounting will be useful for further applications of intertemporal choice.

Acknowledgements

We are grateful to Geir Asheim, Manell Baucells, Glenn Harrison, Sebastian Korditsch, Mark Machina, Kirsten Rohde, Alejandro Saporiti, Peter Wakker and to those who attended presentations in Manchester, Exeter Business School, FUR XV in Atlanta, and the SAET 2013 conference in Paris. Three reviewers made extensive, constructive comments, for which we are grateful. The usual disclaimer applies.

References

- [1] Asheim, Geir B. (1997). Individual and collective time-consistency. *Rev Econ Stud* 64 (1997) 427-443.

- [2] Abdellaoui, Mohammed, Arthur E. Attema, and Han Bleichrodt (2010). Intertemporal tradeoffs for gains and losses: An experimental measurement of discounted utility, *Econ J* 120, 845-866.
- [3] Abdellaoui, Mohammed, Han Bleichrodt and Olivier l'Haridon (2013). Sign-dependence in intertemporal choice. Forthcoming in *J Risk Uncertain*.
- [4] Aczel, Janos and Fulvia Skof (2007). Local pexider and cauchy equations. *Aequationes Mathematicae* 73, 311-320.
- [5] Akin, Zafer (2007). Time inconsistency and learning in bargaining games. *Int J Game Theory* 36, 275-299.
- [6] Attema, Arthur E., Han Bleichrodt, Kirsten I. M. Rohde and Peter P. Wakker (2010). Time trade-off sequences for analyzing discounting and time inconsistency. *Manag Sci* 56:11, 2015-2030.
- [7] Baillon, Aurelien, Bram Driesen and Peter P. Wakker (2012). Relative concave utility for risk and ambiguity. *Games Econ Behav* 75, 481-489.
- [8] Barro, Robert (1999). Laibson meets Ramsey in the neoclassical growth model. *QJE* 114, 1125-1152.
- [9] Benoit, Jean-Pierre and Efe A. Ok (2007). Delay aversion. *Theoretical Econ* 2, 71-113.
- [10] Blackorby, Charles, David Nissen, Daniel Primont and R. Robert Russell (1973). Intertemporal decision making. *Rev Econ Stud* 40:2, 239-248.
- [11] Bleichrodt, Han, Amit Kothiyal, Drazen Prelec and Peter Wakker (2012). Compound invariance implies prospect theory for simple prospects. Forthcoming in *J Math Psych*.
- [12] Diecidue, Enrico, Ulrich Schmidt and Horst Zank (2009). Parametric weighting functions. *J Econ Theory* 144:3, 1102-1118.

- [13] Epper, T., H. Fehr-Duda and A. Bruhin (2011). Viewing the future through a warped lens: Why uncertainty generates hyperbolic discounting. *J Risk Uncertain* 43, 169-203.
- [14] Fishburn, Peter C. and Ariel Rubinstein (1982). Time preference. *Int Econ Rev* 23, 677-694
- [15] Halevy, Yoram (2012). Time consistency: stationarity and time invariance. Working paper, The University of British Columbia.
- [16] Hammond, Peter J. (1976). Changing tastes and coherent dynamic choice. *Rev Econ Stud* 43, 159-173.
- [17] Harris, Christopher and David Laibson (2013). Instantaneous gratification. *QJE*, 128:1, 205-248.
- [18] Hayashi, T. (2003). Quasi-stationary cardinal utility and present bias. *J Econ Theory* 112, 343-352.
- [19] Hey, John D. and Gianna Lotito (2009). Naive, resolute or sophisticated? A study of dynamic decision making. *J Risk Uncertain* 38, 1-25.
- [20] Hey, John D. and Luca Panaccione (2011). Dynamic decision making: What do people do? *J Risk Uncertain* 42, 85-123
- [21] Jamison, Dean T. and Julian Jamison (2011). Characterizing the amount and speed of discounting procedures. *J Benefit Cost Anal* 2:2, Article 1.
- [22] Köbberling, Veronika and Peter P. Wakker (2003). Preference foundations for nonexpected utility: A generalized and simplified technique. *Math Oper Res* 28, 395-423.
- [23] Kodritsch, Sebastian (2012). On time-inconsistency in bargaining. Working paper, London School of Economics.

- [24] Koopmans, Tjalling C. (1960). Stationary ordinal utility and impatience. *Econometrica* 28, 287-309.
- [25] Koopmans, Tjalling C. (1972). Representation of preference orderings with independent components of consumption, and representations of preference orderings over time. In C. B. McGuire and Roy Radner (Eds.), *Decision and organization*: Vol. 57-100. Amsterdam: North-Holland.
- [26] Laibson, David (1997). Golden eggs and hyperbolic discounting. *QJE* 112:2, 443-477.
- [27] Luttmer, Erzo, and Thomas Mariotti (2003). Subjective discounting in an exchange economy. *J Polit Econ* 111:5, 959-989.
- [28] Machina, Mark J. (1989). Dynamic consistency and non-expected utility models of choice under uncertainty. *J Econ Lit* 27, 1622-1668.
- [29] McClennen, Edward F. (1990). *Rationality and dynamic choice: Foundational explorations*. Cambridge University Press, Cambridge.
- [30] Noor, Jawwad (2011). Intertemporal choice and the magnitude effect. *Games Econ Behav* 72, 255-270.
- [31] O'Donoghue, Ted and Matthew Rabin (2001). Choice and procrastination. *QJE* 116, 121-160.
- [32] Ok, Efe A. and Yusufcan Masatlioglu (2007). A theory of (relative) discounting. *J Econ Theory* 137, 214-245.
- [33] Olea, Jose Luis Monteval and Tomasz Strzalecki (2011). A simple axiomatization of quasi-hyperbolic discounting. Working paper, Harvard University.
- [34] Phelps, E. and R. Pollak (1968). On second-best national saving and game-equilibrium growth. *Rev Econ Stud* 35:2, 185-199.

- [35] Rubinstein, Ariel (1982). Perfect equilibrium in a bargaining model. *Econometrica* 50, 97-110.
- [36] Shefrin, Hersh (1998). Changing utility functions. In Barbera, Hammond and Shell (eds.) *Handbook of utility theory volume 1: Principles*. Kluwer academic publishers, 1998.
- [37] Takeuchi, K. (2011). Non-parametric test of time consistency: Present bias and future bias. *Games Econ Behav* 71, 456-478.
- [38] Thaler, Richard H. (1981). Some empirical evidence of dynamic inconsistency. *Econ Let* 8, 201-207.

A Appendices

A.1 Proof of Theorem 3.2.1

First suppose that the initial preference relation is represented as in statement (ii) of the theorem. That this implies statement (i) is straightforward. Weak ordering, continuity, monotonicity and impatience are immediate. Define a function $D : [0, T] \rightarrow \mathbb{R}$ such that $D(t) = \alpha^t$ if $t \leq \lambda$ and $D(t) = (\alpha/\beta)^\lambda \beta^t$ if $t > \lambda$. For midpoint consistency, suppose (for instance) that the first three indifferences of the condition hold. Under statement (ii) this is equivalent to:

$$D(t)u(x) = D(t')u(y) \ \& \ D(t)u(y) = D(t')u(z) \ \& \ D(s)u(x) = D(s')u(y)$$

First notice that, if one of x, y or z are zero, then they are all zero. In that case, the condition holds, given our extended definition of impatience. Suppose one of x, y, z is non-zero. As $D(\cdot)$ is positive and u strictly increasing, it follows that all x, y, z must be positive. Therefore, $u(x), u(y), u(z) > 0$ holds. The first two equalities jointly imply, $u(x)/u(y) = u(y)/u(z) = D(t')/D(t) := \mu$. Given this, the third implies

$D(s')/D(s) = \mu$. The equality $D(s)u(y) = D(s')u(z)$ follows immediately, as does the equivalent, required, fourth indifference. The necessity of two-stage stationarity, given statement (ii) is explained next. If λ is zero or T then all of the conditions of two-stage stationarity hold at all times. Suppose $\lambda \in (0, T)$. Taking any time $t \in [0, T]$, Stationarity-before- t holds whenever $t \leq \lambda$ and stationarity-after- t holds when $t > \lambda$. Both conditions hold at λ . This covers all cases, establishing two-stage stationarity.

For the remaining part of the proof we assume statement (i) of the theorem and derive statement (ii). We first outline some of the implications of two-stage stationarity.

Under weak ordering, and using the definitions of two-stage stationarity, if $\succsim_0$ satisfies stationarity-after- t then it satisfies stationarity-after- t' for any $t' > t$. Similarly, if $\succsim_0$ satisfies stationarity-before- t then it satisfies stationarity-before- t' for any $t' < t$. We will now show that if stationarity-before- t and stationarity-after- t' hold with $t > t'$, then stationarity holds everywhere. To see this, suppose that the conditions of the claim are true. The restriction of preferences to $[0, X] \times [0, t]$ satisfies all the conditions of Theorem 2.1.1 and therefore admits an exponential discounting representation. The same holds for preferences restricted to $[0, X] \times [t', T]$. By the uniqueness results attached to Fishburn and Rubinstein's theorem, we can choose the same δ for each case. Then, there will be a u such that (x, t) mapped to $\delta^t u(x)$ represents preferences on $[0, X] \times [0, t]$, and a $\tilde{u}$ such that (x, t) mapped to $\delta^t \tilde{u}(x)$ represents on $[0, X] \times [t', T]$. By assumption, there is a set $[0, X] \times [t', t]$ where both functions must represent preferences, hence they can be chosen to be equal. Then, preferences over the whole set of timed outcomes admit one exponential discounting representation, and stationarity must necessarily hold everywhere.

We now show the existence of λ . Firstly, if the initial preference relation is stationary, then we may choose either $\lambda = 0$ or $\lambda = T$. Suppose now that the conditions of (ii)

hold, but stationarity does not hold. We use the following definitions:

$$\begin{aligned} t_* &= \sup\{t \in [0, T] : \succsim_0 \text{ satisfies stationarity-before-}t\} \\ t^* &= \inf\{t \in [0, T] : \succsim_0 \text{ satisfies stationarity-after-}t\} \end{aligned}$$

Two-stage stationarity demands that $[0, T] = [0, t_*] \cup [t^*, T]$. By connectedness, if the union of $[0, t_*]$ and $[t^*, T]$ cover $[0, T]$, they must have a non-empty intersection. We cannot have $t_* > t^*$, or else stationarity would hold everywhere as argued above. Then, there is a unique point in this intersection, $t_* = t^* := \lambda$ as required.

To complete the theorem when $\lambda = 0$ or $\lambda = T$, simply notice that the conditions coincide with Theorem 2.1.1, except midpoint consistency being redundant. For the remainder of the proof, we consider the $\lambda \in (0, T)$ case. We apply Observation 4.1 of Bleichrodt, Kothiyal, Prelec and Wakker (2013) to derive a separable representation: preferences are represented by a function that maps (x, t) to $D(t)u(x)$, with D continuous, strictly decreasing and positive with $D(0) = 1$, and u continuous and strictly increasing.⁴ To determine the structure of $D : [0, T] \rightarrow \mathbb{R}_{++}$ we consider its behaviour on $[0, \lambda]$ and $[\lambda, T]$ separately. As shown, preferences satisfy stationarity-before- λ . Then, for $t, s, t + s \leq \lambda$ and $x, x' \in [0, X]$ the following equivalence holds: $(0, x) \sim (t, x')$ if and only if $(s, x) \sim (t + s, x')$. The existence of suitable x and x' is straightforward. Substituting the separable representation we obtain: $u(x) = D(t)u(x')$ if and only if $D(s)u(x) = D(t + s)u(x')$. Equivalently, D satisfies the following local functional equation:

$$D(t + s) = D(t)D(s) \quad t, s, t + s \in [0, \lambda].$$

This is the second of Cauchy's functional equations, restricted to a connected subset of the reals. The classic approach to solving this applies to the case where the equation holds on all of $\mathbb{R}$. One must show that there is an extension of D that

⁴What we call midpoint consistency is referred to as the hexagon condition, and as 1-unit invariance in Bleichrodt, Kothiyal, Prelec and Wakker (2013).

preserves the functional equation. This has been addressed by Aczel and Skof (2007), whose results apply here as D is strictly positive. The general, continuous solution gives $D(t) = \pi\alpha^t$ for all $t \in [0, \lambda]$ for non-zero α and π . The initial condition, $D(0) = 1$, gives $\pi = 1$.

The existence of the separable representation of preferences, when combined with stationarity-after- λ , will lead to a local functional equation on $[\lambda, T]$. There is no $t \in [\lambda, T]$ with $D(t) = 1$. Define a function $\tilde{D}$ such that $\tilde{D}(t) = D(t)/D(\lambda)$ for all $t \in [\lambda, T]$. Notice that $\tilde{D}(\cdot)u(\cdot)$ still represents preferences and that $\tilde{D}(\lambda) = 1$. Stationarity-after- λ guarantees that, for $t, s, t+s \geq \lambda$ and $x, x' \in [0, X]$, $(\lambda, x) \sim (t, x)$ if and only if $(\lambda + s, x) \sim (t + s, x')$. Substituting the rescaled representation gives:

$$\tilde{D}(t+s) = \tilde{D}(t)\tilde{D}(s) \quad t, s, t+s \in [\lambda, T]$$

The general, continuous solution is of the form $\tilde{D}(t) = \tilde{\pi}\beta^t$ for all $t \in [\lambda, T]$, for non-zero $\tilde{\pi}$ and β . That $\tilde{\pi} = \beta^{-\lambda}$ follows immediately from the initial condition. Recall that $D = D(\lambda)\tilde{D}$ on $[\lambda, T]$. Summing up, we have shown that:

$$V_0(x, t) = D(t)u(x) = \begin{cases} \alpha^t u(x) & \text{if } t \leq \lambda \\ (\alpha/\beta)^\lambda \beta^t u(x) & \text{if } t > \lambda \end{cases}$$

as required. ■

A.2 Proof of Proposition 3.2.2

Assume preferences admit a TSE discounting representation $V_0 : [0, X] \times [0, T] \rightarrow \mathbb{R}$ for some parameters $\alpha, \beta \in (0, 1)$, $\lambda \in [0, T]$ and utility function $u : [0, X] \rightarrow \mathbb{R}$. The uniqueness of λ , when $\lambda \notin \{0, T\}$, has been explained in the proof of Theorem 3.2.1 in Appendix A.1. Either λ is unique, or else stationarity must hold everywhere.

Since V_0 represents preferences, it can be replaced by $f \circ V_0$ whenever f is strictly increasing. In general, such transformations need not retain the separable form.

Suppose α is replaced with any $\tilde{\alpha} \in (0, 1)$ and utility u replaced with $\tilde{u} = u^k$ with $k = \ln(\tilde{\alpha})/\ln(\alpha)$. One can verify that $\ln(\alpha^t u(x)) = (1/k) \ln(\tilde{\alpha}^t \tilde{u}(x))$ for all $(x, t) \in [0, X] \times [0, \lambda]$, and because $\ln$ is strictly increasing and $k > 0$, it must be that preferences over $[0, X] \times [0, \lambda]$ are represented by $\tilde{\alpha}^t \tilde{u}(x)$. Similarly, one may verify that $\ln(\beta^t u(x)) = (1/k) \ln(\beta^{kt} \tilde{u}(x))$ for all $(x, t) \in [0, X] \times [\lambda, T]$, hence preferences over $[0, X] \times [\lambda, T]$ are represented by $\tilde{\beta}^t \tilde{u}(x)$ with $\tilde{\beta} = \beta^k$. By the same reasoning, one may choose any $\tilde{\beta} \in (0, 1)$, and proceed as above replacing u and α appropriately.

Once α and β are chosen, utility must be a ratio scale. This follows from well-known results on separable representations, given that the location of utility is fixed. To see this, recall that we included the condition $(0, t) \sim_{\underline{t}} (0, t')$, for any $t, t' \in [0, T]$, in the definition of impatience. Then $u(0) = 0$ holds, or else the representation would not exhibit impatience.

Having chosen either of α or β , however, the other is uniquely determined. To see this, one may take any $x < y$ and find a unique t such that $(x, 0) \sim (y, t)$. Choose x and y such that $t > \lambda$. Substituting the representation and rearranging gives:

$$\beta = \left[\frac{u(x)}{\alpha^\lambda u(y)} \right]^{\frac{1}{t-\lambda}}$$

Given that u is a ratio scale, the right hand side of the above equation is dimensionless. Therefore β , for given α (or vice versa), is uniquely determined. ■

A.3 Proof of Theorem 4.1.1

First assume statement (iii) of the theorem holds. At each decision time, the static preference relation admits a TSE discounting representation. By Theorem 3.2.1, the initial preference is a continuous, monotonic, impatient and midpoint consistent weak order that satisfies two-stage stationarity. The assumed dynamic model is such that, each decision time's representation is the restriction of the initial representation to timed outcomes occurring no sooner than that decision time. At no time, therefore,

can initial preferences be reversed. As such, the dynamic preference structure is time consistent and statement (i) is proved.

We now assume the conditions of statement (i) and derive statement (iii). For $\underline{t} \in \mathcal{D}$, time consistency allows us to identify $\succsim_{\underline{t}}$ with $\succsim_0|_Z$ where $Z = X \times [\underline{t}, T]$; the restriction of initial preferences $\succsim_0$ to the set of timed outcomes occurring no sooner than time $\underline{t}$. Hence, by Theorem 3.2.1, each $\succsim_{\underline{t}}$ may be represented by $V_{\underline{t}} := V_0|_Z$ as required. Each decision time's TSE discounting representation is the restriction of the initial representation, hence every associated parameter (α, β, λ and u) is independent of decision time.

To complete proof of Theorem 4.1.1, we now show the equivalence of statements (i) and (ii). Fix a set of weakly ordered, monotonic, impatient, midpoint consistent and continuous preference relations $\mathcal{R}$. We prove that if $\mathcal{R}$ is time consistent then, stationarity-before- t of $\succsim_0$ is equivalent to time-invariance-before- t of $\mathcal{R}$. Let $(x, t), (y, t'), (x, t - \tau), (y, t' - \tau) \in [0, X] \times [0, T]$ with $\tau \geq 0$ and $t \geq t'$. The following diagram aids the proof of the theorem:

$$\begin{array}{ccc}
 (x, t) \succsim_0 (y, t') & \Rightarrow^1 & (x, t - \tau) \succsim_0 (y, t' - \tau) \\
 \Downarrow^2 & & \Downarrow^3 \\
 (x, t) \succsim_{\underline{t}} (y, t') & \Rightarrow^4 & (x, t - \tau) \succsim_{\underline{t} - \tau} (y, t' - \tau)
 \end{array}$$

Note that implication 1 is stationarity-before- t , equivalences 2 and 3 are time consistency, and implication 4 is time-invariance-before- t . Notice that implication 1 may be deduced by starting at the top left preference, then 2, then 4, and then 3. Implication 4 may be deduced by starting with the bottom left preference, then 2, then 1, and then 3. The equivalence of time consistency with stationarity-after- t , and time consistency with time-invariance-after- t may be similarly shown. ■

A.4 Proof of Theorem 4.2.1

First assume statement (iii) of the theorem holds. At each decision time, the static preference relation admits a TSE discounting representation. By the Theorem 3.2.1, the initial preference is a continuous, monotonic, impatient and midpoint consistent weak order that satisfies two-stage stationarity. The assumed dynamic model is such that, each decision time $\underline{t}$ has a representation obtained by translating the initial representation, and appropriately restricting it, as follows:

$$V_{\underline{t}}(x, t) = V_0(x, t - \underline{t}) \quad \text{on } [0, X] \times [\underline{t}, T].$$

As such, initial preferences cannot be reversed when the timed outcomes and decision time are all translated by a fixed amount. That is, the dynamic preference structure is time invariant and statement (i) is proved.

We now assume the conditions of statement (i) and derive statement (iii). For $\underline{t} \in \mathcal{D}$, define $\succsim_{0, \underline{t}}$ according to:

$$(x, t) \succsim_{0, \underline{t}} (x', t') \Leftrightarrow (x, t - \underline{t}) \succsim_0 (x', t' - \underline{t})$$

for all $(x, t), (x', t'), (x, t - \underline{t}), (x', t' - \underline{t}) \in X \times [0, T]$. By Theorem 3.2.1, $\succsim_0$ is represented by a TSE discounting function V_0 . Construct $\tilde{V}_{\underline{t}} : X \times [\underline{t}, T] \rightarrow \mathbb{R}$ such that $\tilde{V}_{\underline{t}}(\cdot, t) \equiv V_0(\cdot, t - \underline{t})$. Clearly, $\tilde{V}_{\underline{t}}$ represents $\succsim_{0, \underline{t}}$. Time invariance allows us to identify $\succsim_{\underline{t}}$ with the restriction of $\succsim_{0, \underline{t}}$ to the set of timed outcomes occurring no sooner than time $\underline{t}$. Hence, $\succsim_{\underline{t}}$ may be represented by a function $V_{\underline{t}} := \tilde{V}_{\underline{t}}|_Z$ where $Z = X \times [\underline{t}, T]$ as required.

To complete the proof of Theorem 4.2.1, we show the equivalence of statements (i) and (ii). Fix a set of weakly ordered, monotonic, impatient, midpoint consistent and continuous preference relations $\mathcal{R}$. We prove that if $\mathcal{R}$ is time invariant then, stationarity-before- t of $\succsim_0$ is equivalent to time-consistency-within- t -from-now of $\mathcal{R}$. Let $(x, t), (y, t'), (x, t - \underline{t}), (y, t' - \underline{t}) \in [0, X] \times [0, T]$. The following diagram contains

the proof of this claim:

$$\begin{array}{ccc}
 & & (x, t) \succsim_{\underline{t}} (y, t') \\
 & \nearrow^1 & \\
 (x, t) \succsim_0 (y, t') & & \Downarrow^3 \\
 & \searrow^2 & \\
 & & (x, t - \underline{t}) \succsim_0 (y, t' - \underline{t})
 \end{array}$$

Note that implication 1 is time-consistency-within- t -from-now, implication 2 is stationarity-before- t and equivalence 3 is time invariance. We next show that, if $\mathcal{R}$ satisfies time invariance, then stationarity-after- t of $\succsim_0$ and time-consistency-beyond- t -from-now of $\mathcal{R}$ are equivalent. The following diagram contains the proof of this claim:

$$\begin{array}{ccc}
 (x, t) \succsim_0 (y, t') & & \\
 & \searrow^2 & \\
 \Downarrow^1 & & (x, t + \underline{t}) \succsim_0 (y, t' + \underline{t}) \\
 & \nearrow^3 & \\
 (x, t + \underline{t}) \succsim_{\underline{t}} (y, t' + \underline{t}) & &
 \end{array}$$

Note that equivalence 1 is time invariance, implication 2 is stationarity-after- t and implication 3 is time-consistency-beyond- t -from-now. We have, therefore, established that time invariance and two-stage stationarity are jointly equivalent, given the other conditions, to time invariance and two-stage time consistency. ■.