RESERVE BANK OF NEW ZEALAND

DP2010/08

Intertemporal Choice: A Nash Bargaining Approach

David Baqaee

September 2010

JEL classification: D91, D01, C7

www.rbnz.govt.nz/research/discusspapers/

Discussion Paper Series

ISSN 1177-7567

DP2010/08

Intertemporal Choice: A Nash Bargaining Approach*

David Baqaee†

Abstract

A compelling, but highly tractable, axiomatic foundation for intertemporal decision making is established and discussed. This axiomatic foundation relies on methods employed in cooperative bargaining theory. Four simple axioms imply that the intertemporal objective function is a weighted geometric average of each period's utility function. This is in contrast to standard practice, which takes the objective function to be a weighted arithmetic average. The analysis covers both finite and infinite time.

* The Reserve Bank of New Zealand's discussion paper series is externally refereed. The views expressed in this paper are those of the author(s) and do not necessarily reflect the views of the Reserve Bank of New Zealand. The author thanks Thomas Steinke and Christie Smith for their helpful comments.

† Address: Economics Department, Reserve Bank of New Zealand, 2 The Terrace, PO Box 2498, Wellington, New Zealand. *email address:* david.baqaee@gmail.com

ISSN 1177-7567

1 Introduction

Dynamic economic problems are increasingly at the forefront of economic research. Just as the axiomatic development of expected utility aided the study of stochastic economics, a tractable, axiomatic model of intertemporal decision-making can help to provide a framework for studying dynamic economics. Standard preference relations (total, transitive, continuous preorders) do not take advantage of the special properties of time. In this paper, I propose an axiomatic model of intertemporal choice that exploits the structure of time, much in the same way that expected utility theory exploits the nature of probability.

Mas-Colell et al (1995, page 733) note that it is customary in economics to assume that an individual's preferences can be represented as the discounted sum of some *cardinal* per-period payoff function:

$$\sum_{t=0}^{n} \beta^t u(x_t). \tag{1}$$

Various other formulations have been proposed including nonstationary utility functions, hyperbolic discounting and recursive utility functions. A review of the various proposals can be found in Fredrick et al (2002).

However, despite the existence of a large literature on nonstandard intertemporal choice, uptake by the rest of the profession has been slow and uneven. This has possibly been due to the complexity these nonstandard formulations impose, but also due to the absence of a compelling axiomatic foundation. In short, for whatever reason, axiomatic developments of intertemporal choice have not been as successful as axiomatic developments of choice under uncertainty. According to Fredrick et al (2002), the popularity of (1) has "largely [been due] to its simplicity and its resemblance to the familiar compound interest formula". Indeed, when Samuelson (1937) first proposed the geometric discounting model, he used it as a rough descriptive estimate of reality noting that "any connection between utility as discussed here and any welfare concept is disavowed. The idea that the results of such a statistical investigation could have any influence upon ethical judgments of policy is one which deserves the impatience of modern economists."[1] In this paper, a tractable,

[1] Despite Samuelson's reservations, functions like (1) dominate welfare analyses, in everything from monetary policy to environmental economics.

and axiomatically-founded, objective function is proposed, for use in place of (1).

The outline of this paper is as follows. A finite time model of intertemporal choice is set out and studied in section 2. A note on this paper's place within the wider literature is given in section 2.1. In section 3, the results of section 2 are extended to the case where time is countably infinite. The paper concludes in section 4.

2 A model of intertemporal choice

Let $T \subseteq \mathbb{N}$ represent time. For the time being, assume that time is finite:

$$T = \{1, 2, 3, \ldots, n-1, n\}.$$

We will weaken this assumption in section 3.

Suppose that a decision maker (DM) has von Neumann-Morgenstern preferences corresponding to each period of time. The DM needs to decide on an optimal allocation g from a utility possibility set $S \subset \mathbb{R}^T$. Note that this set S is *not* the consumption set, but the utility payoffs associated with the consumption set. Implicitly, S is derived from $(X, \precsim_1, \ldots, \precsim_T)$, where X is a consumption set and $\precsim_i$ is a preference relation corresponding to preferences in period i.

The DM chooses $g \in S$, where $g = (u_1, \ldots, u_n)$, with the ith coordinate representing the *expected* utility in period i. Let d be an outcome in the possibility set that is strictly dominated by at least one other feasible point. We call d the **do-nothing outcome**. The interpretation of d will depend on the specific application, but in most cases we can take this to be the utility associated with consuming nothing with probability 1 in every period.

Suppose that the utility possibility set S is bounded in each dimension.[2] Further, assume that it is also convex – this can be justified by presuming that the DM has the ability to select a probability mixture of any two possible

[2] This assumption is relatively innocuous and follows either from the St. Petersburg Paradox or compactness of the underlying consumption set. The economic implication of this assumption is that the utility attainable in each period is bounded above, and this seems reasonable for the applications envisaged.

allocations. In other words, if x and y are both feasible allocations, then the allocation whereby x is chosen with probability p and y is chosen with probability $1 - p$ is also feasible.

Definition 2.1. Let S be a bounded, closed, convex subset of $\mathbb{R}^T$. Let d be an element of S and suppose that there exists $x \in S$ with $x_i > d_i$ for every $i \in T$. Then the pair (S, d) is called an **intertemporal allocation problem**.

Let $\mathfrak{T}$ denote the set of all intertemporal allocation problems. For each $(S, d) \in \mathfrak{T}$ suppose that the DM chooses an allocation $g(S, d) \in S$. This allocation is taken to be DM's intertemporal choice. More formally,

$$g : \mathfrak{T} \to \mathbb{R}^T \text{ such that } g(S, d) \in S \text{ for each } (S, d) \in \mathfrak{T}.$$

We can make some assumptions about the properties of g.

First recall that if u is an expected utility function representing a preference relation $\preccurlyeq$, then for any $a \in \mathbb{R}_{++}$ and $b \in \mathbb{R}$, the function $au + b$ also represents the same preferences. Since the DM's payoffs, in each period, are expected utilities, we should expect the intertemporal choice to be invaraint to equivalent representations of these utilities.

Axiom 1 (EU)**.** For each $(S, d) \in \mathfrak{T}$, if we multiply a utility function by a positive scalar and translate it, then the solution g must similarly be scaled and translated.

This axiom is natural. However, a word of warning about its interpretation is necessary. We could have stated this axiom another way — we could have said:

For any positive affine transformation T, we have $g(T(S), T(d)) = T(g(S, d))$.

In terms of what happens to (S, d), this is equivalent to (EU). However, recall that, implicitly, (S, d) is given by $(X, \precsim_1, \ldots, \precsim_T)$. There are two changes that could result in an affine transformation to (S, d), one is scaling and translating individual utility functions, the other is changing the underlying consumption set X. Axiom (EU) only applies to the case where utility functions are being scaled, and not to the case where the prize set is being changed. This is a point stressed by Rubinstein et al (1992).

Also note that we place no restrictions on the per-period preference relation other than the fact that it satisfies the axioms of expected utility. In particular, in this general framework, we do not rule out the possibility of habit formation or path dependent utility functions.

The next axiom dictates that all periods are given "some" value by the DM.

Axiom 2 (TR)**.** For each $(S, d) \in \mathfrak{T}$, the point $g(S, d)$ dominates d in every coordinate.

This axiom implies that the DM does not resort to the do-nothing outcome in any period. Recall that this is always possible because, by assumption, there exists at least one point in S that strictly dominates d in every dimension. If the DM satisfies (TR), we say that the DM is **temporally rational**. We can go further and require the following efficiency axiom:

Axiom 3 (E)**.** Let $u, v \in S$. If $u_i \geq v_i$ for each $i \in T$ and $u_j > v_j$ for some $j \in T$, then the DM will not choose v over u.

This efficiency axiom (E) is obviously desirable. The next axiom is also a normatively appealing rationality condition.

Axiom 4 (IIA)**.** If (T, d) and (S, d) are elements of $\mathfrak{T}$ with $T \subseteq S$, then

$$g(S, d) \in T \text{ implies that } g(T, d) = g(S, d).$$

In other words, the DM's choice is independent of irrelevant alternatives. From a normative point of view, (IIA) is compelling.

For each strictly positive $\alpha \in \mathbb{R}^T$, define the **weighted Nash product** to be

$$F_\alpha(u) := \prod_{i \in T} (u_i - d_i)^{\alpha_i}. \tag{2}$$

Theorem 2.1. *For every $(S, d) \in \mathfrak{T}$, if $g : \mathfrak{T} \to \mathbb{R}^T$ satisfies (EU), (TR), (E), and (IIA), then*

$$g(S, d) \in \text{argmax}\left\{F_\alpha(u) : (u_1, \ldots, u_n) \in S\right\},$$

for some strictly positive vector α.

Remark. A simple convexity argument shows that the maximiser of the weighted Nash product is unique.

Theorem 2.1 implies that we can treat (2) as though it were an intertemporal utility function. Let us refer to this allocation as the **Nash temporal solution**. This is because the intertemporal axioms stated here closely parallel the axioms of a Nash bargaining game (see Nash 1950). In fact, the axioms for this problem correspond *exactly* to the assumptions of a nonsymmetric n-player Nash bargaining game (see Roth 1979). The standard proof therefore holds. Unfortunately, if the set T is countably infinite, the proof no longer goes through. In that case, we need to be more careful. The infinite-horizon time problem is treated in section 3.

Theorem 2.1 says that the optimal allocation must be a weighted geometric mean of the expected utilities in each period. The weight vector $\alpha = (\alpha_1, \ldots, \alpha_T)$ represents the DM's impatience. Theorem 2.1 provides us with a tractable, axiomatically-founded objective function for use in dynamic decision problems. This representation gels nicely with our existing theory of choice under uncertainty. Here, the per-period payoff functions are standard axiomatically-founded expected utility functions (determined only up to positive affine transformations) and not vaguely defined cardinal felicity functions. Observe that if $\alpha_i > 0$ for period i, then the DM will *not* assign reservation (no resources) to that period. This is unlike the discounted sum representation of intertemporal choice which can easily give corner solutions (see example in section 2.2).

Why is this theorem true? Let us sketch an informal argument, since it is instructive. First, consider a symmetric, linear utility possibility set A where each period is indistinguishable from any other – in other words

$$A := \left\{ x \in \mathbb{R}^T : \sum_{i \in T} x_i \leq T \text{ and } x \geq 0 \right\}.$$

Observe that $(A, 0) \in \mathfrak{T}$ – that is, $(A, 0)$ is an intertemporal allocation problem as in definition 2.1. Axioms (E) and (TR) imply that $g(A, 0)$ must be a non-corner point on the "Pareto"-frontier of this possibility set. Let $\alpha = g(A, 0)$. The point α will be our weight vector. Now consider a general intertemporal allocation problem (S, d). Let z denote the maximiser of F_α over S. We can use a positive affine transformation $T : \mathbb{R}^T \to \mathbb{R}^T$ to map d to 0 and z to α. Then it can be shown that $T(S) \subset A$. Then by (IIA), the solution to $(T(S), T(d))$ must also be $\alpha = g(A, 0)$. By using (EU), we can conclude that z is the solution to (S, d).

The key insight comes from extracting information about the DM's impatience by using a symmetric utility possibility set like A. In other words, we ask, what point does the DM select when, in terms of utility payoffs, every period is identical? The point selected by the DM in this case gives us information about the impatience of the DM. This information can then be used, in conjunction with (IIA) and (EU), to derive the solution to any intertemporal allocation problem.

Note that this logic will not work if the intertemporal objective function is linear in utilities, as in (1), because for a symmetric set, this objective function would either have no unique maximiser or it would give a corner solution. This is another way of saying that a linear objective function cannot satisfy the axioms.

Proposition 2.2. *For each $(S, d) \in \mathfrak{T}$ there exists $\beta \in \mathbb{R}^T$ such that $\beta > 0$ and*

$$g(S, d) \in \text{argmax} \left\{ \sum_{t=1}^{T} \beta_t (u_t - d_t) : (u_1, \ldots, u_n) \in S \right\}.$$

The proof of this proposition is omitted (the partial derivatives of the weighted Nash product at the Nash temporal solution produce the requisite weights).[3] This proposition tells us that the temporal Nash solution can be reached via a discounted sum where the discount weights are, in general, time varying. The time discount vector β depends on the relative concavities of the utility functions, the possibility set, the do-nothing outcome, and the weight vector.

Remark. Note that in the deterministic case the Nash temporal solution is the maximiser of a discounted sum like (1). This is seen by taking logs of the utility functions:

$$\underset{u \in S}{\text{argmax}} \left\{ \prod_i u_i(x_t)^{\beta_t} \right\} = \underset{u \in S}{\text{argmax}} \left\{ \sum_i \beta_t \log(u_i(x_t)) \right\}.$$

[3] For a more detailed analysis of the links between the weighted Nash product and discounted sum utility functions, see the graphical analysis on page 180 of the book by Binmore (1992). Although Binmore's exposition is geared towards bargaining, his mathematical results are unchanged when applied to an intertemporal problem.

Here, in the discounted sum, the per-period utility function is $\log(u_i(\cdot))$ rather than $u_i(\cdot)$. However in the stochastic case, this is no longer possible because, in general,

$$\beta_t \log\left(\int u_t \mathrm{d}P_t\right) \neq \beta_t \int \log(u_t)\mathrm{d}P_t.$$

Proposition 2.3. *The only per period utility function that results in the Nash solution being the same as the optimisation of*

$$\sum_{t=0}^{n} \beta^t u(x_t), \quad \beta \in (0,1)$$

is $u \equiv \text{constant}$. *That is, for all intents and purposes, the time-homogenous, constant discounting model of choice is inconsistent with the axioms.*

2.1 Bargaining and intertemporal choice

In this subsection, I briefly outline my take on this approach to intertemporal choice, its relationship to bargaining theory, and its place within the wider literature of time preference.

In the set up above, the axioms have been stated abstractly with no reference whatever to bargaining. This is because I find the axioms more normatively compelling if viewed as stand-alone assumptions about intertemporal decision making than if interpreted as the consequence of a metaphorical bargaining game between multiple selves. However, with that said, there is nothing to disallow such an interpretation. Indeed, there is an influential school of theorists who have viewed intertemporal decision making as the outcome of intrapersonal conflict.

These models, like the ones proposed by Schelling (1984) or Winston (1980), are used to model addiction and temptation. In such models, there is a myopic self and a farsighted self who alternately take control of decision making. Continuing this train of thought, researchers like Thaler and Shefrin (1981) frame this conflict in the context of a principal-agent problem. Other models, like the one proposed by Elster (1985), draw further connections between intertemporal choice and interpersonal strategic interactions.

Fredrick et al (2002), in their review of this literature, report that

> "Few of these multiple-self models have been expressed formally, and even fewer have been used to derive testable implications that go much beyond the intuitions that inspired them in the first place. However, perhaps it is unfair to criticise the models for these shortcomings. These models are probably best viewed as metaphors intended to highlight specific aspects of intertemporal choice."

One can easily reinterpret my development of intertemporal choice as the Nash bargaining solution to a cooperative bargaining game (or, in some cases, the Nash equilibrium to an appropriately set up noncooperative game). In this sense, the axioms mentioned above can be taken to establish a normative foundation for the intuitive multiple-self approach of other researchers. Furthermore, the intuitive multiple-self approach, to the extent that it is a true description of reality, can help lend descriptive credence to the axiomatic approach.

It should be noted that this multiple-self view of intertemporal choice runs into conceptual problems when time is taken to be infinite. But the abstract axiomatic approach continues to make sense.

2.2 An example

Suppose there is a quantity α of a non-depreciable good that can be consumed in period 1 or period 2. Suppose that the DM is risk-neutral and has time discount factor $\beta \in (0, 1)$. Then the "classical" intertemporal choice will be given by the following optimisation problem:

$$\max_{x \in [0,\alpha]} x + \beta(\alpha - x).$$

This gives a corner solution of total consumption in period 1 and no consumption in period 2 regardless of the discount factor β. On the other hand, the Nash intertemporal solution is given by

$$\max_{x \in [0,\alpha]} x(\alpha - x)^{\beta},$$

which gives $\frac{1}{1+\beta}$ as the consumption in period 1 and $\frac{\beta}{1+\beta}$ as consumption in period 2. The difference between the Nash solution and the classical solution becomes even starker if we introduce uncertainty. If α is a random variable, then to work out the Nash solution we would need to

$$\max_x \mathrm{E}(x)(\mathrm{E}(\alpha - x))^\beta,$$

or equivalently,

$$\max_x \log\left(\mathrm{E}(x)\right) + \beta \log\left((\mathrm{E}(\alpha - x))\right),$$

where x is now a random variable (see Baqaee and Watt (2010) for more information).

3 Infinite period case

The applicability of the Nash temporal allocation is greatly enhanced if we can take time to be infinite. This is because in many intertemporal models, there is no natural "end-date". It is trivial to show that the Nash temporal allocation is consistent with our axioms (EU), (TR), (E), and (IIA) even when time is infinite. So, in most modeling applications, one can use the weighted Nash product as an agent's intertemporal objective function knowing that this choice function will satisfy our postulated axioms.[4]

However, it would be nice to know that these axioms *compel* one to use a weighted Nash product to represent preferences. Unfortunately, the standard proof of the nonsymmetric Nash bargaining game does not automatically apply to the case where T is countably infinite, so we have to tread more cautiously.

In the infinite-dimensional case, the utility possibility set S is a subset of $\mathbb{R}^{\mathbb{N}}$ and the do-nothing outcome d is an element of $\mathbb{R}^{\mathbb{N}}$. Note that for simplicity, we can take $d = 0$ without losing generality. This is because both $g : \mathfrak{T} \to \mathbb{R}^{\mathbb{N}}$ and the Nash temporal solution are invariant to translations.

The first problem that confronts us in the infinite-dimensional context is our choice of ambient space. The set of all real-valued sequences is too large and unruly to work with, but at the same time, we wish to keep the space as large as possible. Observe that due to its the finite-dimensional context, theorem 2.1 is purely an algebraic statement. Its infinite-dimensional analogue,

[4] It is equally trivial to observe that a standard discounted sum utility function will violate (E) and (EU) in the infinite period case.

however, requires additional topological structure because we need to keep track of convergence.

For each $1 \leq p < \infty$, define l_p to be the space of real-valued sequences equipped with the norm

$$\|x\|_p = (\sum_i |x_i|^p)^{1/p}.$$

We define

$$l_\infty = \left\{ x \in \mathbb{R}^{\mathbb{N}} : \sup_{i \in \mathbb{N}}(x_i) < \infty \right\},$$

and endow it with the supremum norm.

The infinite-period intertemporal choice problem is framed in l_∞.

For each strictly positive sequence $\beta \in l_1$, define the weighted Nash product to be

$$F_\beta(x) = \prod_{i \in \mathbb{N}} (x_i - d_i)^{\beta_i}.$$

To control for the possibility of divergence, we restrict ourselves to those utility possibility sets for which the following holds:

$$S \subset l_\infty, \text{ and for each } p \in l_1, \text{ the set } \{F_p(x) : x \in S\} \text{ is bounded above in } \mathbb{R}. \tag{3}$$

This restriction is a technical one and its economic implications are not immediately obvious. One might fear that it is far too strong and that it rules out economically interesting utility possibility sets. These fears can be allayed by observing that if the set S is taken to be bounded in l_∞, then it will satisfy condition (3). Boundedness in l_∞ simply means that there exists $M \in \mathbb{N}$ such that, for every $x \in S$, $\|x\|_\infty = \sup(x_i) < M$.

The main result in this section is the following:

Theorem 3.1. *Let $T = \mathbb{N}$. If condition (3) is fulfilled, and $g : \mathfrak{T} \to l_\infty$ satisfies (EU), (IIA), (E), and (TR), then there exists a strictly positive sequence $(\alpha_i)_{i=1}^\infty$ with $\sum_{i=1}^\infty \alpha_i = 1$, such that*

$$g(S, d) \in \operatorname{argmax} \left\{ \prod_{i \in T} (u_i - d_i)^{\alpha_i} : (u_i)_{i \in T} \in S \right\}.$$

In order to prove this result, we first need to do some spadework.

Lemma 3.2. *For strictly positive $p \in l_1$ with $\|p_i\| = 1$,*

$$p \in \operatorname{argmax} \left\{ \prod_i x_i^{p_i} : \|x\|_1 \leq 1 \wedge x_i \geq 0 \right\}$$

Proof. It is clear that if the maximiser exists, then it must satisfy the constraint with equality. Also note that the maximiser of $\prod x_i^{p_i}$ subject to the constraint $\|x\|_1 = 1$ will also maximise $\sum p_i \log(x_i)$ subject to the same constraint. So, we can use the method of Lagrange multipliers to derive the solution.

The Lagrangian is

$$\mathfrak{L}(x, \lambda) = \sum_i p_i \log(x_i) - \lambda(1 - \sum_i x_i),$$

and the first order conditions are

$$\frac{p_i}{x_i} = -\lambda \quad \text{(for every } i\text{)}. \tag{4}$$

Hence

$$\begin{aligned} \sum_i x_i = -\sum_i \frac{p_i}{\lambda} &= 1 \\ \sum_i p_i &= -\lambda. \end{aligned}$$

We know that $\sum_i p_i = 1$ and this implies that $\lambda = -1$. Therefore, (4) implies that

$$p_i = x_i \quad \text{(for every } i\text{)}$$

■

Lemma 3.3. *For each positive p with $\|p\|_1 = 1$, function $F_p = \prod_{i \in \mathbb{N}} x_i^{p_i}$ is concave over strictly positive l_∞ sequences.*

Proof. Let $x \in l_\infty$ and $x > 0$. Then for each positive p with $\|p\|_1 = 1$, the concavity of the log function implies that

$$\log \left(\prod_{i=1}^n x_i^{p_i} \right) = \sum_{i=1}^n p_i \log(x_i) \leq \log(\sum_{i=1}^n p_i x_i).$$

By exponentiation,

$$\prod_{i=1}^{n} x_i^{p_i} \leq \sum_{i=1}^{n} p_i x_i \quad (n \geq 1).$$

By taking limits, first of the right hand side and then the left hand side, deduce that

$$\prod_{i=1}^{\infty} x_i^{p_i} \leq \sum_{i=1}^{\infty} p_i x_i. \tag{5}$$

Now let $x, y \in l_\infty$ with $x, y > 0$ and $\lambda \in [0, 1]$. From (5),

$$\begin{aligned} \prod_i \left(\frac{\lambda x_i}{\lambda x_i + (1-\lambda) y_i} \right)^{p_i} &\leq \sum_i p_i \frac{\lambda x_i}{\lambda x_i + (1-\lambda) y_i} \\ \prod_i \left(\frac{(1-\lambda) y_i}{\lambda x_i + (1-\lambda) y_i} \right)^{p_i} &\leq \sum_i p_i \frac{(1-\lambda) y_i}{\lambda x_i + (1-\lambda) y_i}. \end{aligned}$$

Sum the last two lines to get

$$\begin{aligned} \prod_i \left(\frac{\lambda x_i}{\lambda x_i + (1-\lambda) y_i} \right)^{p_i} + \prod_i \left(\frac{(1-\lambda) y_i}{\lambda x_i + (1-\lambda) y_i} \right)^{p_i} &\leq \sum_i p_i \frac{\lambda x_i + (1-\lambda) y_i}{\lambda x_i + (1-\lambda) y_i} \\ &= \sum_i p_i \\ &= 1. \end{aligned}$$

Hence,

$$\prod_i (\lambda x_i)^{p_i} + \prod_i ((1-\lambda) y_i)^{p_i} \leq \prod_i (\lambda x_i + (1-\lambda) y_i)^{p_i},$$

which implies that

$$\lambda F_p(x) + (1-\lambda) F_p(y) \leq F_p(\lambda x + (1-\lambda) y).$$

■

Lemma 3.4. *Let U and S be two convex sets in l_∞. Suppose that a strictly positive, continuous linear functional $f \in l_1$ separates U and S via the hyperplane $H = \{x : f \cdot x = c\}$. Consider a linear bijection $T : X \to X$ defined by*

$$(T(x))_i = \beta_i x_i,$$

where $\beta \in l_1$ and $\beta > 0$. Then $T(H)$ separates $T(U)$ from $T(S)$.

Proof. Let $(x, z, y) \in T(S) \times T(H) \times T(U)$. Observe then that, without loss of generality,

$$f \cdot T^{-1}(x) \leq f \cdot T^{-1}(z) \leq f \cdot T^{-1}(y),$$

or in other words

$$\sum_i \frac{f_i}{\beta_i} x_i \leq \sum_i \frac{f_i}{\beta_i} z_i \leq \sum_i \frac{f_i}{\beta_i} y_i,$$

for every $(x, z, y) \in T(S) \times T(H) \times T(U)$. So,

$$T(H) = \left\{ x : \sum_i \frac{f_i}{\beta_i} x_i = c \right\},$$

is a hyperplane separating $T(U)$ from $T(S)$. ■

Lemma 3.5. *For each $r \in \mathbb{R}$, the set*

$$F_p^{-1}(r, \infty) := \left\{ x \in \mathbb{R}^{\mathbb{N}} : F_p(x) > r \right\}$$

has nonempty interior.[5]

Proof. Let $x_i = d_i + e^{|r|}$ for each i. Hence $F_p(x) = e^{|r|} > r$ implying that $x \in F_p^{-1}(r, \infty)$. Let

$$0 < \delta < \frac{1}{2}.$$

Suppose that $\|x - y\| < \delta$. It suffices to show that $F_p(y) > r$. For each i,

$$\left| \frac{y_i - d_i}{x_i - d_i} - 1 \right| = \frac{|x_i - y_i|}{|x_i - d_i|} < \delta e^{-|r|} < 1.$$

If $|x - 1| < 1$, then $|\log(x)| \leq |x - 1|$. Therefore,

$$\left| \log\left(\frac{y_i - d_i}{x_i - d_i} \right) \right| \leq \left| \frac{y_i - d_i}{x_i - d_i} - 1 \right| < \delta e^{-|r|}.$$

Hence,

$$\begin{aligned} |\log(F_p(y)) - \log(F_p(x))| &= |\sum_i p_i(\log(y_i - d_i) - \log(x_i - d_i))| \\ &\leq \left| \sum_i p_i \log\left(\frac{y_i - d_i}{x_i - d_i} \right) \right| \\ &\leq \sum_i p_i \delta e^{-|r|} = \delta e^{-|r|} < \frac{1}{2}. \end{aligned}$$

[5] The proof of this proposition was communicated to me by Thomas Steinke.

For $|a-b| < 1/2$,

$$|e^a - e^b| = e^a|1-e^{b-a}| \leq e^a \left|1 - \sum_{i\geq 0} \frac{|b-a|^i}{i!}\right| \leq e^a \sum_{i\geq 1} |b-a|^i = e^a \frac{|b-a|}{1-|b-a|} \leq 2e^a|b-a|.$$

Hence, because $|\log(F_p(y)) - \log(F_p(x))| < 1/2$, we conclude that

$$\begin{aligned} |F_p(x) - F_p(y)| &\leq 2F_p(x)|\log(F_p(y)) - \log(F_p(x))| \\ &< 2e^{|r|}\delta e^{-|r|} = 2\delta. \end{aligned}$$

This implies that

$$F_p(y) > F_p(x) - 2\delta = e^{|r|} - 2\delta \geq 1 + |r| - 2\delta > r.$$

So $F_p(y) \in F_p^{-1}(r, \infty)$ whenever $\|x - y\| < \delta$. This completes the proof. ■

For any set U, denote the interior of U by U°.

Theorem 3.6 (Interior Separating Hyperplane Theorem)**.** *In any topological vector space X, if the interior of a convex set A is nonempty and disjoint from another nonempty convex set B, then there exists a continuous linear functional f and $p \in \mathbb{R}$ such that $A \subset \{x \in X : f(x) \geq p\}$ and $B \subset \{x \in X : f(x) \leq p\}$ with $(A \cup B) \not\subset \{x \in X : f(x) = p\}$.*

See Aliprantis and Border (2007, pg. 202) for a proof.

Definition 3.1. Let X be a Banach space and U an open set in X. We say a function $f : U \to \mathbb{R}$ has **support** at $x_0 \in U$ if there is an affine function $A : X \to \mathbb{R}$ such that $A(x_0) = f(x_0)$ and $A(x) \geq f(x)$ for every $x \in U$

Theorem 3.7. *Let f be concave on an open set $U \subseteq X$. Then if the Gateaux derivative $f'(x_0)$ exists, f has $A(x) = f(x_0) + f'(x_0)(x - x_0)$ as its unique support at x_0.*

See Roberts and Varberg (1973, pg. 115) for a proof.

Given lemmas 3.2 to 3.5 and theorems 3.6 and 3.7, we can now prove our target.

Proof of theorem 3.1. Define

$$A = \left\{ x \in \mathbb{R}^T : \sum_{i \in T} x_i \leq 1 \wedge x_i \geq 0 \right\}.$$

Observe that $(A, 0) \in \mathfrak{T}$. Denote $g(A, 0) = \alpha$ and observe that, by the efficiency axiom (E), $\alpha > 0$ and $\|\alpha\|_1 = 1$. Now consider a general intertemporal allocation problem (S, d). As described before, without loss of generality, take $d = 0$. Let

$$z \in \operatorname{argmax}\{F_\alpha(x) : x \in S\}.$$

This point exists by closedness of S and the boundedness condition (3).

Let

$$U = \left\{ x \in \mathbb{R}^T : F_\alpha(x) \geq F_\alpha(z) \right\}.$$

By lemma 3.3, the weighted Nash product F_α is concave, and therefore U is a convex set. By lemma 3.5, U has nonempty interior. Furthermore $S \cap U^\circ = \varnothing$. Hence, by theorem 3.6, there exists a continuous linear functional f, and a corresponding hyperplane $H = \{x : f(x) = p\}$, separating U and S. Observe that $z \in H$. In other words, f supports F_α at z. But by the concavity of F_α and theorem 3.7,

$$f(x) = F_\alpha(z) + F'_\alpha(z) \cdot (x - z).$$

We can compute $F'_\alpha(z)$ explicitly by differentiation:

$$F'_\alpha(z) \cdot (y) = \frac{\mathrm{d}}{\mathrm{d}t} \left(\prod_i (z_i + t y_i)^{\alpha_i} \right) \Bigg|_{t=0}.$$

This gives

$$f(x) = F_\alpha(z) + F_\alpha(z) \sum \frac{\alpha_i}{z_i}(x_i - z_i). \tag{6}$$

So, $x \in U$ implies that $f(x) \geq f(z)$ and $x \in S$ implies that $f(x) \leq f(z)$. Hence,

$$H = \left\{ x \in \mathbb{R}^T : f(x) = f(z) = F_\alpha(z) \right\}$$

is the separating hyperplane.

Now, let T denote the invertible linear transformation that scales the ith coordinate by

$$(T(x))_i = \frac{\alpha_i}{z_i} x_i. \tag{7}$$

and observe that

$$\text{if } x \in H, \text{ then } \sum_i (T(x))_i = 1.$$

This is because if $x \in H$, then $f(x) = F_\alpha(z)$, meaning that

$$F_\alpha(z) + F_\alpha(z) \sum \frac{\alpha_i}{z_i}(x_i - z_i) \quad = \quad F_\alpha(z).$$

This can be simplified to show that

$$\sum \frac{\alpha_i}{z_i} x_i = 1,$$

but (7) implies that

$$\sum (T(x))_i = 1.$$

By lemma 3.4, observe that $T(S)$ is now separated from $T(U)$ via $T(H)$ and that $T(z) \in T(H)$. Since f supports U at z,

$$F_\alpha(z) \geq F_\alpha(x) \text{ for all } x \in \{x : f(x) \leq f(z)\}.$$

Hence,

$$F_\alpha(T(z)) \geq F_\alpha(x) \text{ for all } x \in \{x : \sum x_i \leq 1\}.$$

Therefore, we conclude that $T(z)$ maximises F_α over A. By lemma 3.2, $T(z) = \alpha$. But by definition, $\alpha = g(A, 0)$. Hence $g(A, 0)$ is the point that maximises the Nash product over A. By (E) and (IIA),

$$g(T(S), 0) = g(A, 0) = T(z).$$

By (EU), we infer that $g(S, 0) = z$. Since both the solution and the maximiser of the weighted Nash product are invariant to translations, we observe that the result holds in general, even when $d \neq 0$. ■

In other words, if the DM obeys (EU), (IIA), (E), and (TR), then the DM's preferences can be represented via *some* weighted Nash product.

Remark. Equation (6) allows an easy extension to proposition 2.2 to be stated in the infinite-time context.

This approach to intertemporal choice holds promising possibilities for use in intergenerational welfare analysis because, unlike a discounted sum, it does not rely on cardinal utilities. Furthermore, because the axioms stated here are intuitively appealing for making interpersonal comparisons, it is a more credible way of justifying concepts like the "dynastic" motive in fields like environmental economics and macroeconomics.

Remark. The choice of modeling time discretely is not critical. The logic in this section can be extended to cover the case where $S \subset \mathbb{R}^{\mathbb{R}}$. In this case, we would have to frame the problem in L_∞. We would also need to generalise the weighted Nash product so that it makes sense over an uncountable set. The simplest way to do this is to define the weighted Nash product to be the antilog of the Lebesgue integral of the log of each period's payoff function:

$$F_p(x) := \exp\left(\int_{t\in\mathbb{R}} p(t)\log(x(t))\,\mathrm{d}t\right).$$

Alternatively, one could use the notion of a "product integral" to define the weighted Nash product (see, for example, Grossman and Katz (1972) for more information).

3.1 Dynamic programming and Euler equations

Dynamic programming and Euler equations are bread and butter tools for studying dynamic problems. The applicability and tractability of the Nash temporal allocation is further enhanced if these tools can also be used to study the Nash allocation.

By logging the objective function, most of the dynamic programming results from Stokey, Lucas, and Prescott (1989) carry through. This is despite the fact that logarithms and expectations do not commute. Similarly, variational techniques, like Euler equations, are just as easily applicable to a discounted product as they are to a discounted sum.

4 Conclusion

In this paper, we have seen that four simple axioms, axioms that parallel those of a Nash bargaining game, imply that an individual's intertemporal choice can be represented as the maximiser of a weighted geometric average of his per-period expected utilities. These axioms cover countably infinite time, and they may also be generalised to cover continuous time. We observed that this intertemporal objective function, due to its reliance on expected utility, rather than cardinal felicity, also holds promises for intergenerational welfare analysis.

The weighted geometric average of utility functions proposed here is a tractable and axiomatically-founded way of modeling rational intertemporal choice.

References

Aliprantis, C D and K C Border (2007), *Infinite Dimensional Analysis: a Hitchhiker's Guide*, Springer, 3 edn.

Binmore, K (1992), *Fun and Games*, D.C. Heath and Company.

Elster, J (1985), "Weakness of will and the free-rider problem," *Mathematical Social Sciences*, 10(3), 293–294.

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

Grossman, M and R Katz (1972), *Non-Newtonian Calculus*, Lee Press.

Mas-Colell, A, M D Whinston, and J R Green (1995), *Microeconomic Theory*, Oxford University Press.

Nash, J (1950), "The bargaining problem," *Econometrica*, 18(2), 155–162.

Roberts, A W and D E Varberg (1973), *Convex Functions*, Academic Press.

Roth, A (1979), *Axiomatic Models of Bargaining*, vol 170 of *Lecture Notes in Economics and Mathematical Systems*, Springer Verlag.

Rubinstein, A, Z Safra, and W Thomson (1992), "On the interpretation of the nash bargaining solution and its extension to non-expected utility preferences," *Econometrica*, 60(5), 1171–1186.

Samuelson, P (1937), "A note on measurement of utility," *Review of Economic Studies*, 4, 155–161.

Schelling, T C (1984), "Self-command in practice, in policy, and in a theory of rational choice," *American Economic Review*, 74(2), 1–11.

Stokey, N, R Lucas, and E Prescott (1989), *Recursive Methods in Economic Dynamics*, Harvard University Press.

Thaler, R H and H M Shefrin (1981), "An economic theory of self-control," *Journal of Political Economy*, 89(2), 392–406.

Winston, G C (1980), "Addiction and backsliding : A theory of compulsive consumption," *Journal of Economic Behavior & Organization*, 1(4), 295–324.