

Additive valuations of streams of payoffs that satisfy the time value of money principle: A characterization and robust optimization

ABRAHAM NEYMAN

Institute of Mathematics, and Federmann Center for the Study of Rationality,
Hebrew University of Jerusalem

This paper characterizes those preferences over bounded infinite utility streams that satisfy the time value of money principle and an additivity property, and the subset of these preferences that, in addition, are either impatient or patient. Based on this characterization, the paper introduces a concept of optimization that is robust to a small imprecision in the specification of the preference, and proves that the set of feasible streams of payoffs of a finite Markov decision process admits such a robust optimization.

KEYWORDS. Valuations, utility, robust optimization, infinite streams of payoffs.

JEL CLASSIFICATION. C72, E43.

1. INTRODUCTION

It is well documented that economic agents' behavior is often incompatible with exact optimization. We believe that economic theory has overemphasized exact optimization, especially its implications for economic agents' behavior. The reason is that in most real-life applications, there is some imprecision in the specification of the model, in particular, in the specification of the preference over the possible outcomes. Therefore, a rational economic agent may, or even should, forgo exact optimization in a single economic model with well-defined parameters, and prefer behavior that is approximately optimal in a whole class of economic models whose differences reflect imprecision in specifying the parameters of the model. Moreover, it is desirable that the economic agent's behavior exhibits gradual change as the model changes.

We call this type of behavior robust optimization and believe that robust optimization may explain many of the gaps between the observed behavior of economic agents and the behavior that is implied by exact optimization.

The present paper focuses on robust optimization in an economic model where the decision maker must choose between different feasible bounded infinite streams of payoffs and the imprecision is in the specification of the preference.

Abraham Neyman: aneyman@math.huji.ac.il

This research was supported in part by Israel Science Foundation Grant 1596/10. The author wishes to thank Elon Kohlberg for many valuable comments throughout the writing of this paper, and Elchanan Ben-Porath, Sergiu Hart, John Hillas, Eilon Solan, and Yair Tauman for comments on an earlier draft.

The decision maker can be an individual, a firm, or a community of individuals. The stream of payoffs can be a stream of equal payoffs (called a *perpetuity*) or of payoffs that vary over time.

The first part of the paper characterizes the preferences over bounded infinite streams of payoffs that satisfy a few plausible assumptions. The characterization shows that each one of these preferences is represented¹ by a unique cardinal utility, called hereafter a *valuation*.

The essential difference between the characterization of valuations in the present paper and earlier studies (e.g., Koopmans (1960), Koopmans, Diamond, and Williamson (1964), Diamond (1965), Fishburn (1966), Koopmans (1972), Lauwers (1995), Chichilnisky (1996, 1997), Chambers and Echenique (2018), Dugeon and Huy (2021)) is the adoption of the *time value of money principle*. The time value of money principle reflects the preference for expediting the receipt of positive payoffs: the faster the accumulation of payoffs, the better. In other words, this principle states that a unit of payoff in a given period is weakly preferable to a unit of payoff that is spread out over later periods. This principle is natural when saving is costless.

The time value of money principle is the natural generalization of the time-preference principle (see, e.g., Olson and Bailey (1981), von Böehm-Bawerk (1912)) and of the overtaking criterion (see, e.g., von Weizsäcker (1967), Brock (1970), Brown and Lewis (1981) and the references therein). While the time value of money principle is meaningful for all preferences or ordinal utilities over streams of payoffs, the time-preference principle is meaningless for preferences or ordinal utilities for which only the payoffs in the distant future matter and the overtaking criterion is meaningless for preferences or ordinal utilities for which the payoffs in the distant future are negligible.

The other assumptions that are used in our characterization are variations of additivity, nontriviality, and Wold's condition (Wold (1943–1944), Beardon and Mehta (1994)).

The additivity property states that if the streams A and B are equivalent to the perpetuities C and D , respectively, then the sum of the streams A and B is equivalent to the sum of the perpetuities C and D .

The other assumption is that any stream of payoffs is equivalent to a perpetuity (Wold's condition) and the higher the perpetuity's (constant) payoff, the better.

A valuation is *impatient* if the contribution of payoffs in the distant future is negligible. It is *patient* if only the payoffs in the distant future matter.

The main result of the first part is Theorem 3, which characterizes the set of valuations. It shows that the set of valuations is the set of mixtures of impatient valuations, which are characterized in Theorem 1, and patient valuations, which are characterized in Theorem 2.

Our characterization of impatient valuations (respectively, patient valuations) uses the same properties that are used in the characterization of (general) valuations with the addition of an impatience property (respectively, a patience property).

¹A preference relation $\succsim$ over a set X is represented by a cardinal utility u if for all $x, y \in X$, it holds that $x \succsim y$ if and only if $u(x) \geq u(y)$.

There are other characterizations of patient and impatient valuations that use specific axioms for the characterization of each subclass of valuations. As will be evident from our study of robust optimization in the second part, we view the non-impatient valuations as “limit points” of impatient ones or as an imprecise description of an impatient valuation. Therefore, it is advantageous to use the same properties that are meaningful in characterizing impatient valuations in the characterization of (general) valuations and in the characterization of patient valuations.

Our characterizations show that (a) any impatient valuation is a weighted average of the periods’ payoffs with averaging weights that are non-increasing in time (Theorem 1), (b) any patient valuation is a linear function that assigns to each stream a value that is between the limit inferior and the limit superior of the averages of the first n payoffs in the stream (Theorem 2), and (c) any valuation is a weighted average of an impatient valuation and a patient one (Theorem 3).

Two classic examples of impatient valuations are the n th Cesàro average valuation, which is denoted by u_n , and the r -discounted valuation ($0 < r \leq 1$), which is denoted by u_r . For a stream $f = (f_1, f_2, \dots)$,

$$u_n(f) = \frac{f_1 + \dots + f_n}{n} \quad \text{and} \quad u_r(f) = \sum_{t=1}^{\infty} r(1-r)^{t-1} f_t.$$

The t th period’s averaging weight of u_n is $1/n$ if $t \leq n$ and is 0 if $t > n$, and the t th period’s averaging weight of u_r is $r(1-r)^{t-1}$.

A mixture of a patient valuation and an impatient valuation takes into account payoffs both in the near and in the distant future, and, therefore, it is useful for studying economic models like global warming where one must take into account both the foreseeable and the distant future.

The second part of the paper uses the characterization in the first part to define and study robust optimization in models where the decision maker chooses between feasible bounded infinite streams of payoffs.

As there is a one-to-one correspondence between the preferences (that satisfy our assumptions) and the valuations, it suffices to study optimization that is robust to a small imprecision in the specification of the valuation.

Optimization that is robust to small changes in the valuation is common in a bank’s selection of its portfolio. A few considerations in selecting the portfolio are discussed as an illustration of the importance of robust optimization in selecting a proper feasible stream of payoffs.

A bank’s portfolio consists of its assets, which are mainly a collection of loans, and its liabilities, which are mainly a collection of customers’ (including other banks’) deposits and bonds issued by the bank, where each asset or liability has a different maturity and a different payment schedule.

The economic value of the bank is the present value of the stream of its portfolio payoffs. It is a function of the yield curve, which specifies the interest rate as a function of time.

The bank’s set of feasible portfolios depends on market and competitive conditions, as well as on regulatory constraints. One of the regulatory constraints, as well as an

important consideration in the bank's selection of its portfolio, is the sensitivity of its economic value to changes in the yield curve.²

The objective of maximizing the value of the bank's portfolio while ensuring that the losses due to changes in the yield curve remain within prescribed limits is essentially an approximate optimization that is robust to a given imprecision in the specification of the valuation.

The yield curve, and hence also the valuation, changes over time. Therefore, an additional desired property of the bank's portfolio-selection strategy is that the selected portfolio is gradually modified as the yield curve changes.

The fact that each preference can be represented by a unique valuation (and each valuation is a cardinal utility and thus allows comparison of utility differences) enables us to quantify approximate optimization: optimization of the valuation within a small positive ε .

To define a small imprecision in the specification of the preference, we define a topology on the space of valuations. A minimal requirement for the topology on the space of valuations is that the map that assigns to each valuation the utility, according to this valuation, of a given unit vector³ be continuous. The smaller the topology on the space of valuations is, the less stringent is the notion of a small change in the valuation and hence the more demanding is the notion of robust optimization. Therefore, the topology on the space of valuations that we are studying is the minimal topology for which the map that assigns to each valuation its value on a unit vector is continuous. We show that this topology is compact (and pseudo-metrizable).

We define the concepts of a *robust ε -optimizer at a given valuation*. It is a feasible stream of payoffs whose utility, according to any valuation in some neighborhood of the given valuation, is at least the utility that this valuation, and, moreover,⁴ any valuation in this neighborhood, assigns to any feasible stream of payoffs minus ε .

Theorem 4 proves that the existence of a robust ε -optimizer at any valuation implies the existence of finitely many streams of payoffs such that for each valuation, one of them is a robust ε -optimizer, and the existence of a robust ε -optimizer that depends continuously on the valuation whenever the set of feasible infinite streams of payoffs is convex.

The third part of the paper illustrates a nontrivial application of the theory of robust optimization that is developed in the first two parts. The application is to the classical model of a finite Markov decision process (MDP). It proves the existence of robust optimization in the model of a finite MDP.

A policy in a MDP defines a probability over infinite streams of payoffs rather than a deterministic infinite stream of payoffs. Our first result, Theorem 5, proves the existence of robust optimization in a finite MDP under the assumption that the preference over

²Obviously, there are other important sensitivity issues. We mention the sensitivity to the yield curve as the yield curve specifies the valuation.

³That is, a stream where in one period the payoff is 1 and in all other periods the payoff is 0.

⁴This additional stronger requirement guarantees that the oscillation, in this neighborhood, of the optimal value at a valuation, namely, the supremum of the utility of the valuation over all feasible streams of payoffs, is no more than ε .

distributions coincides with the preference over the deterministic stream of expected payoffs in each period.

The time value of money principle plays a crucial role in the existence of robust optimization in a finite MDP. Without it, and even with the addition of the time-preference principle and the strong Pareto optimality assumptions, the existence of a robust ε -optimizer at a preference for which the payoffs in the distant future are not negligible does not hold in a finite MDP.⁵

Theorem 6, which generalizes Theorem 5, demonstrates the existence of robust optimization in a finite MDP where the valuations are viewed as von Neumann–Morgenstern utilities.

2. CHARACTERIZATION OF VALUATIONS

This section defines formally the concepts of valuation, impatient valuation, and patient valuation, and states the theorems that characterize each of them in turn.

2.1 Streams of payoffs

A stream of payoffs is a sequence $g = (g_1, g_2, \dots)$ of real numbers. It is bounded if $\|g\| := \sup_t |g_t| < \infty$. The linear space of all bounded streams of payoffs is denoted by ℓ_∞ .

For $g, h \in \ell_\infty$ and $a \in \mathbb{R}$, $g + h$ is the element $(g_1 + h_1, g_2 + h_2, \dots)$ of ℓ_∞ , i.e., the t th coordinate of $g + h$ is $g_t + h_t$, and ag is the element $(ag_1, ag_2, \dots)$ of ℓ_∞ , i.e., the t th coordinate of ag is ag_t .

2.2 Linearity

The t th coordinate, g_t , of the stream g is often interpreted as the utility of consumption at stage t , and several classic sets of axioms (see [Debreu \(1959\)](#), [Fishburn \(1966\)](#)) lead to a presentation of a utility over infinite streams of consumption that is a linear function of the stream g .

A real-valued function u that is defined on ℓ_∞ is *additive* if for every $g, h \in \ell_\infty$, we have that $u(g + h) = u(g) + u(h)$. As $\mathbf{0} + \mathbf{0} = \mathbf{0}$, where $\mathbf{0} = (0, 0, \dots)$, an additive u satisfies $u(\mathbf{0}) = 0$.

A real-valued function u that is defined on ℓ_∞ is *linear* if it is additive and $u(ag) = au(g)$ for every $g \in \ell_\infty$ and $a \in \mathbb{R}$.

2.3 The time value of money principle

This principle captures two desirable properties of a function $u : \ell_\infty \rightarrow \mathbb{R}$ that represents a preference over streams of payoffs.

⁵Explicitly, one can (i) characterize all linear and monotonic real-valued functions that are defined on the space of bounded streams of payoffs and satisfy the time-preference principle and the strong Pareto optimality assumptions, (ii) define a natural topology on the set of these functions, (iii) define (analogously) the concept of a robust ε -optimizer, and (iv) show that there is a finite MDP for which a robust ε -optimizer does not exist at any preference for which the payoffs in the distant future are not negligible.

The first is monotonicity: the higher the stage payoffs are the better. For an additive u , monotonicity is equivalent to the property that a stream of nonnegative payoffs is at least as desirable as the stream of zero payoffs.

The second desirable property of u expresses the fact that the earlier the payments are, the better: a unit payoff in a given period is at least as desirable as its spread over later periods. This implies the *time-preference*⁶ property: $u(\mathbf{e}_t) \geq u(\mathbf{e}_{t+1})$ for all t , where $\mathbf{e}_t$ is the t th unit vector in ℓ_∞ .

An additive u satisfies the time-preference property (i.e., $u(\mathbf{e}_t) \geq u(\mathbf{e}_{t+1})$ for all t) if and only if for any two streams g and h that differ only in finitely many periods of nonzero payoffs and satisfy $\sum_{t=1}^s g_t \geq \sum_{t=1}^s h_t \forall s$, we have $u(g) \geq u(h)$.

The time value of money principle, which is defined formally below, is a generalization of the time-preference property and is a key principle in the characterization of valuations.

DEFINITION 1. A real-valued function u that is defined on ℓ_∞ satisfies the *time value of money principle* if, for every two streams g and h such that $\sum_{t=1}^s g_t \geq \sum_{t=1}^s h_t \forall s$, we have $u(g) \geq u(h)$.

REMARK 1. A function $u : \ell_\infty \rightarrow \mathbb{R}$ that satisfies the time value of money principle is monotonic, i.e., $u(g) \geq u(h)$ whenever $g_t \geq h_t \forall t$, and satisfies $u(\mathbf{e}_t) \geq u(\mathbf{e}_{t+1})$ for all t .

REMARK 2. An additive and monotonic function $u : \ell_\infty \rightarrow \mathbb{R}$ satisfies $u(\mathbf{e}_t) \geq 0$ and $\sum_{t=1}^\infty u(\mathbf{e}_t) < \infty$, and, therefore, $u(\mathbf{e}_t)$ goes to zero as t goes to infinity.

2.4 Valuations

DEFINITION 2. A real-valued function u that is defined on ℓ_∞ is *normalized* if $u(\mathbf{1}) = 1$, where $\mathbf{1} = (1, 1, \dots)$.

DEFINITION 3. A normalized additive real-valued function that is defined on ℓ_∞ and satisfies the time value of money principle is called a *valuation*.

Recall that two classic examples of valuations are the n th Cesàro average valuation u_n and the r -discounted valuation ($0 < r \leq 1$) u_r . For a stream $g = (g_1, g_2, \dots)$,

$$u_n(g) = \frac{g_1 + \dots + g_n}{n} \quad \text{and} \quad u_r(g) = \sum_{t=1}^{\infty} r(1-r)^{t-1} g_t.$$

Another example of a valuation is a linear function L on ℓ_∞ such that for every $g \in \ell_\infty$, $L(g) \geq \liminf_{n \rightarrow \infty} \bar{g}_n$, where $\bar{g}_n = (g_1 + g_2 + \dots + g_n)/n$. (Note that as L is linear, the inequality $L(g) \geq \liminf_{n \rightarrow \infty} \bar{g}_n \forall g$ implies that for every $g \in \ell_\infty$, $L(g) \leq \limsup_{n \rightarrow \infty} \bar{g}_n$.)

⁶For theoretical, empirical, or historical accounts of time-preference properties, see, e.g., von Böhm-Bawerk (1912), Olson and Bailey (1981), Rothbard (1990), Fishburn and Edwards (1997), Doyle (2013), Galor and Özak (2016), Falk, Becker, Dohmen, Enke, Huffman, and Sunde (2018) and the references therein.

2.5 Preferences and valuations

Many writers, e.g., Debreu (1959, 1960), Koopmans (1960), Koopmans, Diamond, and Williamson (1964), Diamond (1965), Fishburn (1966), Koopmans (1972), Olson and Bailey (1981), Lauwers (1995), Chichilnisky (1996, 1997), Fishburn and Edwards (1997), Chambers and Echenique (2018), Dugeon and Huy (2021), have studied the implications of various axioms on preferences over product sets, e.g., on sequences of consumptions or on streams of payoffs, and the representation of the preferences by ordinal utilities.

In this section, we present a list of axioms (on preferences over bounded streams of payoffs) such that a preference over bounded streams of payoffs satisfies the axioms if and only if it is represented by a valuation.

A preference relation $\succsim$ on ℓ_∞ satisfies the *time value of money principle* if $g \succsim h$ whenever g and h are two streams in ℓ_∞ such that $\sum_{t=1}^s g_t \geq \sum_{t=1}^s h_t \forall s$; it is *additive* if for every $\alpha, \beta \in \mathbb{R}$, we have that $(g + h) \succsim (\alpha + \beta)\mathbf{1} \equiv (\alpha + \beta, \alpha + \beta, \dots)$ whenever $g \succsim \alpha\mathbf{1}$ and $h \succsim \beta\mathbf{1}$; it is *nontrivial* if there are $g, h \in \ell_\infty$ such that $g \succ h$, i.e., $g \succsim h$ and not $h \succsim g$; it is *complete* if for every g and h , either $g \succsim h$ or $h \succsim g$; it is *transitive* if $f \succsim h$ whenever $f \succsim g$ and $g \succsim h$.

The next result states properties of a preference relation that are sufficient for it to be represented by a valuation.

PROPOSITION 1. *For every nontrivial preference relation $\succsim$ on ℓ_∞ that is complete (alternatively, transitive), additive, and satisfies the time value of money principle, and such that for every stream g there is $\alpha \in \mathbb{R}$ such that $g \sim \alpha\mathbf{1}$, i.e., $g \succsim \alpha\mathbf{1}$ and $\alpha\mathbf{1} \succsim g$, there exists a unique valuation v such that v represents $\succsim$ as an ordinal utility, i.e., $g \succsim h$ if and only if $v(g) \geq v(h)$.*

Obviously, if a valuation v represents the preference relation $\succsim$ on ℓ_∞ , then $\succsim$ is complete, transitive, satisfies the time value of money principle, and for every stream g there is $\alpha \in \mathbb{R}$ such that $g \sim \alpha\mathbf{1}$.

2.6 Impatient valuations

This section defines an impatient valuation (Definition 4 below), remarks that an impatient valuation is continuous on norm-bounded subsets of ℓ_∞ when ℓ_∞ is equipped with the product discrete topology and the range $\mathbb{R}$ is equipped with the standard topology (Remark 3 below), and notes that in the characterization of impatient valuations, the time value of money principle can be replaced by the time-preference assumption.

Let $\mathbf{1}_{>n}$ be the stream of payoffs $g = (g_1, g_2, \dots)$ with $g_t = 1 \forall t > n$ and $g_t = 0 \forall t \leq n$.

DEFINITION 4. An *impatient valuation* is a valuation u such that

$$u(\mathbf{1}_{>n}) \rightarrow_{n \rightarrow \infty} 0.$$

REMARK 3. If u is an impatient valuation, then for every $g \in \ell_\infty$, $u(g_1, g_2, \dots, g_n, 0, 0, \dots)$ converges to $u(g)$ as n goes to infinity, where $(g_1, g_2, \dots, g_n, 0, 0, \dots)$ stands for the stream whose t -coordinate equals g_t if $t \leq n$ and equals 0 if $t > n$.

Moreover, if u is an impatient valuation, then $u(g_1, \dots, g_n, h_{n+1}, \dots)$, where $g, h \in \ell_\infty$ and $(g_1, \dots, g_n, h_{n+1}, \dots)$ stands for the stream whose t -coordinate equals g_t if $t \leq n$ and equals h_t if $t > n$, converges to $u(g)$ as n goes to infinity.⁷

Moreover, the above convergence is, for each $K > 0$, uniform in $g \in \ell_\infty$ and h with $\|h\| \leq K$.

The first result characterizes all impatient valuations.

THEOREM 1. *A real-valued function u that is defined on ℓ_∞ is an impatient valuation if and only if there are weights ω_t , where $t \geq 1$ ranges over the positive integers, with $\omega_t \geq 0$ and $\sum_{t=1}^\infty \omega_t = 1$, such that*

$$u(g) = \sum_{t=1}^{\infty} \omega_t g_t.$$

The r -discounted valuation and the n th Cesàro average valuations are impatient valuations. The weights representing the r -discounted valuation u_r are $\omega_t = r(1-r)^t$, and those representing the k th Cesàro average valuation u_k are $\omega_t = 1/k$ if $t \leq k$ and $\omega_t = 0$ if $t > k$.

The following result shows that in the characterization of impatient valuations, the time value of money principle can be replaced by monotonicity and the time-preference property that $u(\mathbf{e}_t) \geq u(\mathbf{e}_{t+1})$.

LEMMA 1. *A monotonic, impatient, and additive function $u : \ell_\infty \rightarrow \mathbb{R}$ that satisfies $u(\mathbf{e}_t) \geq u(\mathbf{e}_{t+1})$ satisfies the time value of money principle.*

Therefore, a real-valued function that is defined on ℓ_∞ is an impatient valuation if and only if it is normalized, linear, $u(\mathbf{1}_{>n}) \rightarrow_{n \rightarrow \infty} 0$, and $u(\mathbf{e}_t) \geq u(\mathbf{e}_{t+1})$ for all t .

2.7 Convergence of impatient valuations

Next, we define convergence of a sequence of impatient valuations.

DEFINITION 5. A sequence u^k of impatient valuations *converges* if for every positive integer t , the sequence $u^k(\mathbf{e}_t)$ converges as $k \rightarrow \infty$.

The subspace of ℓ_∞ of all converging sequences $g \in \ell_\infty$, i.e., the limit of g_t exists as t goes to infinity, is denoted by c . An equivalent definition of convergence of a sequence of impatient valuations follows.

⁷This property of a function $u : \ell_\infty \rightarrow \mathbb{R}$ is Fishburn's convergence axiom (Fishburn (1966)).

REMARK 4. A sequence v^k of impatient valuations converges if and only if $v^k(g)$ converges for every $g \in c$.

It follows that the limit of a converging sequence of impatient valuations defines a real-valued function on c . On this restricted domain, the “limit” v satisfies the following properties of a valuation: linearity, $v(\mathbf{1}) = 1$, and the time value of money principle.

Examples of converging sequences of impatient valuations are the k th Cesàro average valuations, u_k , which converge as k goes to infinity, and the r -discounted valuations, u_r , which converge as $r > 0$ goes to zero.

The limit v of a sequence of impatient valuations need not coincide with the restriction of an impatient valuation to the domain c . For example, if v is the limit of the k th Cesàro average valuations u_k , then, for every fixed n , the sequence $u_k(\mathbf{1}_{>n})$ converges to 1 as k goes to infinity, and, therefore, $v(\mathbf{1}_{>n}) = 1$; hence, v is not impatient.

2.8 Patient valuations

DEFINITION 6. A *patient valuation* is a valuation u such that

$$u(\mathbf{1}_{>n}) = 1 \quad \forall n \geq 1.$$

Note that for any valuation u , if for some $n \geq 1$, we have $u(\mathbf{1}_{>n}) = 1$, then for all $n \geq 1$, we have $u(\mathbf{1}_{>n}) = 1$.

The second result characterizes the patient valuations.

THEOREM 2. A real-valued function u that is defined on ℓ_∞ is a patient valuation if and only if it is a linear function on the bounded streams of payoffs such that

$$\liminf_{n \rightarrow \infty} \bar{g}_n \leq u(g) \leq \limsup_{n \rightarrow \infty} \bar{g}_n. \quad (4)$$

The next result shows that the lower and upper bounds in Theorem 2 are tight.

LEMMA 2. For every bounded g there are patient valuations u and v such that $v(g) = \liminf_{n \rightarrow \infty} \bar{g}_n$ and $u(g) = \limsup_{n \rightarrow \infty} \bar{g}_n$.

The next result shows that in the characterization of patient valuations it is impossible to replace the time value of money principle with the condition that $u(\mathbf{e}_t) \geq u(\mathbf{e}_{t+1})$ for all t .

LEMMA 3. There is a normalized linear function $w : \ell_\infty \rightarrow \mathbb{R}$ that is monotonic and satisfies $w(\mathbf{e}_t) = 0 \forall t$; hence, w satisfies $w(\mathbf{e}_t) \geq w(\mathbf{e}_{t+1})$, but does not satisfy the time value of money principle.

A patient valuation can be viewed informally as a limit of the k th Cesàro average valuation as k goes to infinity and of the r -discounted valuations as $0 < r < 1$ goes to zero. This informal view will be made formal at a later stage.

2.9 Characterization of valuations

There are other possible informal limits of impatient valuations. For example, a weighted average $\beta v + (1 - \beta)w$, $0 \leq \beta < 1$, of an impatient valuation w and a patient one v is the informal limit, as k goes to infinity, of the impatient valuations $\beta u_k + (1 - \beta)w$.

The next result characterizes all valuations by showing that the weighted averages of an impatient valuation and a patient one are all the valuations.

THEOREM 3. *A real-valued function u that is defined on ℓ_∞ is a valuation if and only if it is a convex combination of an impatient valuation and a patient one.*

2.10 Ordinal and cardinal utilities on ℓ_∞

This section introduces properties of general real-valued functions that are defined on the bounded infinite streams of payoffs. It serves to relate the characterization of valuations to other known results.

An ordinal utility on a set X (e.g., ℓ_∞) is a real-valued function u that is defined on X . It represents a preference $\succsim$ on X if for all $x, y \in X$, $x \succsim y$ if and only if $u(x) \geq u(y)$.

A valuation is, in particular, an ordinal utility on ℓ_∞ . An ordinal utility u on ℓ_∞ is *strong Pareto optimal* if for all distinct $x, y \in \ell_\infty$, $x \geq y \implies u(x) > u(y)$.

DEFINITION 7. An ordinal utility u on ℓ_∞ is *impatient* if for all $g, h \in \ell_\infty$ with $u(g) > u(h)$, we have that for any positive constant C , there is a period $T(g, h, C)$ such that for all $T \geq T(g, h, C)$ and $g', h' \in \ell_\infty$ with $\|g'\|, \|h'\| \leq C$,

$$u(g_{\leq T}, g'_{> T}) > u(h_{\leq T}, h'_{> T}),$$

where $(g'_{\leq T}, g_{> T})$ is the sequence of payoffs whose payoff in period t equals g_t if $t > T$ and equals g'_t if $t \leq T$.

Impatience of an ordinal utility u on ℓ_∞ is called *dictatorship of the present* in Chichilnisky (1996).

Note that an impatient ordinal utility u satisfies $u(\mathbf{1}_{>n}) \rightarrow_{n \rightarrow \infty} u(\mathbf{0})$, where $\mathbf{0}$ is the perpetuity with constant payoff 0. However, an ordinal utility that satisfies $u(\mathbf{1}_{>n}) \rightarrow_{n \rightarrow \infty} u(\mathbf{0}) = 0$ need not be impatient. An impatient valuation is an impatient ordinal utility since in addition to its impatient property, $u(\mathbf{1}_{>n}) \rightarrow_{n \rightarrow \infty} 0$, it is linear and monotonic.

DEFINITION 8. A ordinal utility u on ℓ_∞ is *patient* if for all $g, h \in \ell_\infty$ with $u(g) > u(h)$, we have that for any positive constant C there is a period $T(g, h, C)$ such that for all $T \geq T(g, h, C)$ and $g', h' \in \ell_\infty$ with $\|g'\|, \|h'\| \leq C$,

$$u(g'_{\leq T}, g_{> T}) > u(h'_{\leq T}, h_{> T}).$$

Patience of an ordinal utility u on ℓ_∞ is called *dictatorship of the future* in Chichilnisky (1996).

Note that a patient ordinal utility u satisfies $u(\mathbf{1}_{>n}) \rightarrow_{n \rightarrow \infty} u(\mathbf{1})$. However, an ordinal utility that satisfies $u(\mathbf{1}_{>n}) \rightarrow_{n \rightarrow \infty} u(\mathbf{1}) = 1$ need not be patient. A patient valuation is a patient ordinal utility since in addition to its patient property, $u(\mathbf{1}_{>n}) \rightarrow_{n \rightarrow \infty} = 1$, it is linear and monotonic.

The monotonic linear functionals on ℓ_∞ have the form

$$u(g) = \sum_{t=1}^{\infty} w_t g_t + \phi(g),$$

where $w_t \geq 0$, $\sum_{t=1}^{\infty} w_t < \infty$, and ϕ is a monotonic linear function on ℓ_∞ with $\liminf_{t \rightarrow \infty} g_t \leq \phi(g) \leq \limsup_{t \rightarrow \infty} g_t$. It is strong Pareto optimal if and only if $w_t > 0 \forall t$. It is normalized if and only if $\phi(\mathbf{1}) + \sum_{t=1}^{\infty} w_t = 1$. It satisfies the time-preference property if and only if $w_t \geq w_{t+1}$. It is impatient if and only if $\phi = 0$, and it is patient if and only if $w_t = 0$ for all t .

A normalized and monotonic linear functional u on ℓ_∞ , e.g., a valuation, is a cardinal utility since the difference $u(g) - u(h)$ is the unique number c such that $g \sim h + c\mathbf{1}$.

Using the above characterization of monotonic linear functionals on ℓ_∞ , Chichilnisky (1996) characterized all linear ordinal utilities on ℓ_∞ that are neither impatient nor patient, and satisfy the strong Pareto optimality.

The theory of robust optimization that is developed in the next section applies also to robust optimization of normalized monotonic linear functionals on ℓ_∞ . However, the existence of robust optimization for a finite MDP, which is stated in Theorem 5 in the sequel, does not hold when the possible preferences are represented by normalized monotonic linear cardinal utilities on ℓ_∞ .

The above representation of monotonic linear functionals on ℓ_∞ can be used to provide the first step in an alternative proof to our elementary proof of the characterization of valuations. The outline of the alternative proof follows.

A valuation u is, in particular, a monotonic linear functional on ℓ_∞ . Hence it has the representation $u(g) = \sum_{t=1}^{\infty} w_t g_t + \phi(g)$ as above.

The time-preference property of a valuation (which follows from the time value of money principle) implies that in a valuation, $w_t \geq w_{t+1}$, which together with monotonicity implies that $\sum_{t=1}^{\infty} w_t < \infty$, and the normalization assumption implies that $\phi(\mathbf{1}) + \sum_{t=1}^{\infty} w_t = 1$.

The last step is to show that the time value of money principle implies that $\phi(g)$ is between $1 - \sum_{t=1}^{\infty} w_t$ times the limit inferior of $\bar{g}_n$ and $1 - \sum_{t=1}^{\infty} w_t$ times the limit superior of $\bar{g}_n$, as $n \rightarrow \infty$. This last step is essentially the proof of Theorem (4) that appears in Section 5.2 in the sequel.

2.11 Valuations that satisfy additional properties

In this section we state several easily derived results that identify the valuations that satisfy various additional properties/assumptions/postulates that were used in previous studies of ordinal or cardinal utilities on ℓ_∞ or on preferences over product sets.

2.11.1 Debreu's independent and essential factors properties Any preference $\succsim$ on ℓ_∞ that is defined by a valuation v satisfies Debreu's independent factor property (Debreu (1960, Definition 4)), and the i th factor is essential (Debreu (1960, Definition 4)) if and only if $v(\mathbf{e}_i) > 0$.

2.11.2 Continuity properties

FACT 1. Any preference $\succsim$ that is defined by a valuation v satisfies Diamond's PSC continuity axiom (Diamond (1965)), i.e., $\forall g \in \ell_\infty$, $\{g' : g' \succsim g\}$ and $\{g' : g \succsim g'\}$ are closed in the sup (norm) topology. It satisfies Diamond's PPC continuity axiom (Diamond (1965)), i.e., $\forall g \in \ell_\infty$ and $\forall C > 0$, the sets $\{g' : \|g'\| \leq C \text{ and } g' \geq g\}$ and $\{g' : \|g'\| \leq C \text{ and } g \succsim g'\}$ are closed in the product topology, if and only if v is an impatient valuation.

FACT 2. The valuation v satisfies Fishburn's convergence axiom (Fishburn (1966, (UC))), i.e., $\forall g, h \in \ell_\infty$, $\lim_{n \rightarrow \infty} v(g_1, \dots, g_n, h_{n+1}, h_{n+2}, \dots) = v(g)$, if and only if v is an impatient valuation.

2.11.3 Diamond's sensitivity properties Diamond's sensitivity properties are versions of monotonicity, which states that more is better. Recall that weak monotonicity of a valuation is implied by the time value of money principle.

Diamond's S1 sensitivity property (Diamond (1965)) is composed of two properties: (S1.1), $g' \geq g \implies g' \succsim g$, and (S1.2), also called weak Pareto, $g'_t > g_t \forall t \implies g' > g$, and Diamond's S2 sensitivity property (Diamond (1965)), also called strong Pareto, is $(g' \geq g \text{ and } g \neq g') \implies g' > g$.

FACT 3. Any preference that is defined by a valuation v satisfies (S1.1). It satisfies (S1.2) if and only if v is not a patient valuation (equivalently, $v(\mathbf{e}_1) > 0$), and it satisfies S2 if and only if $v(\mathbf{e}_t) > 0 \forall t$.

2.11.4 Koopmans' postulates (Koopmans (1960, 1972), Koopmans, Diamond, and Williamson (1964))

FACT 4. A valuation v satisfies Koopmans' stationary recursiveness property; i.e., there is a function V that is defined on $\mathbb{R}^2$ such that

$$v(g_1, g_2, \dots) = V(g_1, v(g_2, g_3, \dots)) \quad \forall g = (g_1, g_2, \dots) \in \ell_\infty$$

if and only if v is either a patient valuation (and then $V(a, b) = b$) or v is the discounted valuation u_r (and then $V(a, b) = ra + (1 - r)b$).

FACT 5. A valuation v satisfies Koopmans' sensitivity postulate (Koopmans (1960, Postulate 2)); i.e., there exist $x, x' \in \mathbb{R}$ and $g \in \ell_\infty$ such that $(x, g) \succ (x', g)$ if and only if v is not a patient valuation.

Therefore, we can state the following facts.

FACT 6. *A valuation v satisfies Koopmans' stationary recursiveness postulate and Koopmans' sensitivity postulate if and only if it is a discounted valuation u_r , $0 < r \leq 1$.*

FACT 7. *Any valuation v satisfies Koopmans' aggregation by period postulates (Koopmans (1960, (P3a) and (P3b))), equivalently, the limited complementarity postulates (Koopmans, Diamond, and Williamson (1964, (P3a) and (P3b))), i.e., for all $x, x' \in \mathbb{R}$ and for all $g, g' \in \ell_\infty$, we have that $v(x, g) \geq v(x', g)$ implies $v(x, g') \geq v(x', g')$ and $v(x', g) \geq v(x', g')$ implies $v(x, g) \geq v(x, g')$.*

FACT 8. *A valuation v satisfies Koopmans' stationarity postulate (Koopmans (1960, Postulate 4)) (i.e., for some $x \in \mathbb{R}$, for all $g, g' \in \ell_\infty$, we have $v(x, g) \geq v(x, g')$ if and only if $v(g) \geq v(g')$) if and only if v is a mixture of a patient valuation and a discounted valuation u_r for some $0 < r < 1$.*

2.11.5 The equal-treatment and time-neutrality properties

FACT 9. *A valuation v satisfies Diamond's equal-treatment property (Diamond (1965, (C))), also called time neutrality or intergenerational equity or finite anonymity (i.e., $v(g) = v(\pi g)$ for every permutation π of the positive integers with only finitely many t with $\pi(t) \neq t$ (where πg is the stream of payoffs whose i -period payoff is $g_{\pi(i)}$)) if and only if v is a patient valuation.*

However, the next fact follows from Lemma 3.

FACT 10. *A normalized and monotonic linear functional $v : \ell_\infty \rightarrow \mathbb{R}$ that satisfies Diamond's equal-treatment property need not satisfy the time value of money principle and, therefore, need not be a patient valuation.*

Forges (1986) labels a linear functional v on ℓ_∞ as *time-neutral* if v satisfies (4), i.e., if and only if v is a patient valuation (Theorem 2), and Lauwers (1995) proves that a linear functional u on ℓ_∞ is time-neutral if and only if it is monotonic, $u(\mathbf{1}) = 1$, and $u(g) = u(\pi g)$ for every permutation π such that $\lim_n \pi(n)/n = 1$ (where (πg) is defined by $(\pi g)_t = g_{\pi(t)}$).

2.11.6 *The overtaking and catching-up criteria* Fix a real-valued function $v : \ell_\infty \rightarrow \mathbb{R}$. We say that v satisfies the overtaking criterion if $v(g) > v(h)$ whenever $\liminf_{T \rightarrow \infty} \sum_{t=1}^T (g_t - h_t) > 0$. We say that v satisfies the catching-up criterion if $v(g) \geq v(h)$ whenever $\liminf_{T \rightarrow \infty} \sum_{t=1}^T (g_t - h_t) \geq 0$. It satisfies the alternative catching-up criterion if $v(g) \geq v(h)$ whenever $\sum_{t=1}^T (g_t - h_t) \geq 0$ for all sufficiently large values of T . Note that if v satisfies the catching-up criterion, then it satisfies the alternative catching-up criterion, but not vice versa.

FACT 11. *(a) No valuation satisfies the overtaking criterion, (b) a patient valuation satisfies the catching-up criterion, and (c) a normalized linear function $v : \ell_\infty \rightarrow \mathbb{R}$ that satisfies the alternative catching-up criterion is a patient valuation.*

2.11.7 Wold's condition In his paper on a continuous function representing a preference relation $\succsim$ on the positive orthant of $\mathbb{R}^n$, Wold deduces from a list of plausible axioms that any $y \in \mathbb{R}^n$ is equivalent to a bundle x on the diagonal, i.e., $x_i = x_j \forall 1 \leq i, j, \leq n$. Wold's condition has been widely used to establish numerical representations of preferences over product sets under a variety of different assumptions (see, e.g., [Diamond \(1965\)](#), [Asheim, Mitra, and Tungodden \(2012\)](#), [Mitra and Ozbek \(2013\)](#), [Banerjee \(2014\)](#), [Banerjee and Mitra \(2018\)](#)), and its infinite-dimensional version is our basic assumption that any bounded stream of payoffs is indifferent to some perpetuity. In particular, any valuation satisfies Wold's condition.

2.11.8 The Pareto and intergenerational equity properties A preference relation $\succeq$ on ℓ_∞ satisfies the strong (respectively, weak) Pareto property if and only if for every two distinct elements $g, h \in \ell_\infty$, $g \geq h$ (respectively, $g_t > h_t \forall t$) implies $g \succ h$, and it satisfies the intergenerational equity property (also called the finite anonymity or equal treatment or time neutrality property) if and only if $g \sim \pi g$ for every finite permutation π of $\mathbb{N}$.

The characterization of valuations shows the following fact.

FACT 12. *There does not exist any valuation v that satisfies the weak Pareto and the intergenerational equity properties.*

Fact 12 follows from the more general result that there does not exist any function that aggregates an infinite stream of payoffs into a real number satisfying weak Pareto and intergenerational equity ([Basu and Mitra \(2003, Theorem 1\)](#)).⁸

2.11.9 Measurability of valuations A Borel-measurable valuation is a valuation v such that the map v from $[0, 1]^\mathbb{N}$ ($\subset \ell_\infty$) to $\mathbb{R}$ is measurable when $\mathbb{R}$ is equipped with the σ -algebra of Borel sets and $[0, 1]^\mathbb{N}$ is equipped with the σ -algebra of weak* Borel sets.

Any impatient valuation is Borel measurable. Moreover, if v is an impatient valuation, then the following statements hold:

- (i) The map v is universally measurable on $B = [0, 1]^\mathbb{N}$, i.e., if μ is a Borel probability measure with respect to the product topology on B , then v is μ -measurable.
- (ii) It is measure-linear, i.e., if $f_n : [0, 1] \rightarrow [0, 1]$ is a sequence of Borel measurable functions, then the function $f = v((f_n))$ is measurable with respect to any probability measure μ on B , and we have the identity

$$\int f d\mu \equiv \int v((f_n)) d\mu = v\left(\left(\int f_n d\mu\right)\right).$$

These two properties of an impatient valuation enable one to use an impatient valuation as a von Neumann–Morgenstern utility over streams of payoffs so that the preference over distributions of streams of payoffs coincides with the preference over the deterministic stream of stage payoffs.

⁸This nonexistence result applies to ordinal utilities on ℓ_∞ and not to complete orderings. [Svensson \(1980\)](#) established the general possibility result (for a social welfare relation) that one can find an ordering that satisfies the axioms of strong Pareto and intergenerational equity.

The existence of a Borel-measurable patient valuation that satisfies properties (i) and (ii) above follows from the Zermelo–Frankel axioms together with the axiom of choice and, e.g., the continuum hypothesis or Martin’s axiom (Martin and Solovay (1970)).⁹

However, “most” patient valuations are not Borel measurable. The nonmeasurability makes it conceptually difficult to distinguish between two distinct patient valuations. However, the concept of robust optimization, which is detailed in Section 3, overcomes this difficulty as the topology on the space of valuations that is used in the theory of robust optimization “identifies” all patient valuations.

A related difficulty with measurability of preference relations is raised in Zame (2007).¹⁰

3. ROBUST OPTIMIZATION

This section starts with a subsection on optimization that defines (for a given set of feasible streams of payoffs) the optimal value and an approximate optimizer for a valuation or a subset of valuations. The subset of valuations can be interpreted as an imprecision in the specification of a valuation. To define a small imprecision in the specification of a valuation, we define in Section 3.2 a topology on the set of valuations and state a few properties of the topology.

Section 3.3 defines a robust ε -optimizer at a valuation as an approximate optimizer for a sufficiently small neighborhood of the valuation, discusses the relation between the existence of robust ε -optimizers and the continuity of the optimal value, and states a minmax type condition that is equivalent to the existence of a robust ε -optimizer for every $\varepsilon > 0$. Section 3.4 shows that any bounded stream of payoffs has, for any $\varepsilon > 0$ and an impatient valuation, a robust ε -optimizer.

Section 3.5 remarks that the notion of robustness at a patient valuation provides a unifying view of earlier studies of robust optimization of a patient decision maker, and Section 3.6 remarks on the need to consider robust optimization at non-impatient valuations even if one wishes to confine the analysis to impatient valuations.

Section 3.7 states the implications of a bounded set of streams of payoffs F having a robust ε -optimizer at every valuation v . In this case, (a) there are finitely many streams in F such that for each valuation, one of them is a robust ε -optimizer, and (b) if, in addition, the set F is convex, then there is a continuous function $v \mapsto f^v \in F$ that maps a valuation v to an ε -optimizer at v .

3.1 Optimization

For any valuation v , the maximum (or more precisely, the supremum) of $v(g)$ over all streams g in F is called the *v -optimal value of F* and is denoted by $v(F)$.

⁹As far as I know, the existence of a Borel-measurable patient valuation is not provable in Zermelo–Frankel axioms together with the axiom of choice.

¹⁰Following the terminology of Zame (2007), a preference relation $\succsim$ on $[0, 1]^{\mathbb{N}}$ is *ethical* if $\pi g \succsim \pi h$ whenever $g \succsim h$ and π is a finite permutation and $g > h$ whenever $g_t > h_t \forall t$. Zame (2007, Theorem 2) shows that any ethical preference over $[0, 1]^{\mathbb{N}}$ is not measurable.

An imprecise specification of a valuation is modeled as a set U of valuations. The maximum (or, more precisely, the supremum) of $u(g)$ over all streams g in F and valuations u in U is called the *U -optimal value of F* and is denoted by $U(F)$.

Fix a nonnegative number $\varepsilon \geq 0$, a valuation v , a set of valuations U , a set of streams of payoffs F , and a stream f in F .

The stream $f \in F$ is an *ε -optimizer for v with respect to F* if $v(f)$ (which is at most the v -optimal value of F) is within ε of the v -optimal value of F (i.e., $v(f) \geq v(F) - \varepsilon$ for any $g \in F$).

The stream $f \in F$ is an *ε -optimizer for U with respect to F* if for any valuation u in U , we have that $u(f)$ (which is at most the U -optimal value of F) is within ε of the U -optimal value of F (i.e., $u(f) \geq U(F) - \varepsilon$ for any valuation w in U and any stream g in F). Note that an ε -optimizer for U with respect to F is, for any $u \in U$, an ε -optimizer for u with respect to F .

It follows that if the set F of streams of payoffs has an ε -optimizer for a set of valuations U , then the oscillation of the u -optimal value of F , where u ranges over all valuations in U , is at most ε .

An imprecision in the specification of a valuation is often expressed by stating that a fixed valuation v is a good proxy for the “true” valuation. Such an imprecise specification of the valuation u is modeled as the set of all valuations that are sufficiently similar to the fixed valuation v . This leads to the following important concept of robust optimization. This concept depends on the topology on the space of valuations.

3.2 The topology on the set of valuations

To define nearby valuations, as well as the proximity of one valuation to another one, we need to define a topology on the set V of valuations.

The coarser the topology is, the larger the neighborhoods of a point are. Therefore, the coarser the topology is, the stronger the positive results on the existence of robust optimization are. Hence, we define the topology $\mathcal{T}$ on the space of valuations as the coarsest topology in which the most basic real-valued functions $v \mapsto v(\mathbf{e}_t)$, $t \geq 1$, are continuous.

This topology is the minimal topology in which the denumerably many functions $v \mapsto v(\mathbf{e}_t)$, $t \geq 1$, are continuous.

As the topology $\mathcal{T}$ on V is defined by countably many continuous functions, V is a pseudo-metric space. Namely, there is a function $d : V \times V \rightarrow \mathbb{R}_+$, e.g., $d(u, v) = \max_{t \geq 1} |v(\mathbf{e}_t) - u(\mathbf{e}_t)|$, such that (i) $d(u, v) + d(v, w) \geq d(u, w) \forall u, v, w \in V$, (ii) for every neighborhood U of a valuation u , there is $\varepsilon > 0$ such that any valuation v with $d(v, u) < \varepsilon$ is in U , and (iii) for every valuation v and a positive $\varepsilon > 0$, $\{u : d(u, v) < \varepsilon\} \in \mathcal{T}$.

By defining the equivalence relation $\equiv$ on V by $u \equiv v$ if and only if u is 0-close to v , i.e., $d(u, v) = 0$ (equivalently, $v(\mathbf{e}_t) = u(\mathbf{e}_t) \forall t$), the space of equivalence classes $V/\equiv$ is a metrizable space.

Recall that c is the subspace of ℓ_∞ that consists of all converging sequences. The following remark states a few properties of the topological space $(V, \mathcal{T})$. In particular, it shows, implicitly, that $\mathcal{T}$ is the minimal topology on V in which the functions $v \mapsto v(g)$ are continuous for each $g \in c$.

REMARK 5. The topological space $(V, \mathcal{T})$ is compact.

The impatient valuations are dense in V .

A sequence v^k of valuations converges if and only if the sequence $v^k(\mathbf{e}_t)$ converges $\forall t$.

A sequence v^k of valuations converges if and only if the sequence $v^k(g)$ converges $\forall g \in c$.

For any two distinct impatient valuations $v, u \in V$, there is a converging sequence $g \in c$ such that $v(g) \neq u(g)$.

For any two patient valuations $v, u \in V$ and for any converging sequence $g \in c$, we have $v(g) = u(g)$. Therefore, any neighborhood of a patient valuation includes all patient valuations.

Note that for any neighborhood W of a patient valuation, there is a positive integer k_0 and a positive $0 < r_0 < 1$ such that for all $k \geq k_0$ and $0 < r \leq r_0$, the impatient valuations u_r and u_k are in W .

3.3 Robust optimization at a valuation

Let F be a set of bounded streams of payoffs and let v be a valuation. Recall that the v -optimal value of F , $v(F)$, is defined by $v(F) = \sup_{f \in F} v(f)$ and that the following definition holds.

DEFINITION 9. An element $f \in F$ is a *robust ε -optimizer at v with respect to F* , $\varepsilon \geq 0$, if there is $\delta > 0$ such that

$$u(f) \geq w(F) - \varepsilon \quad \text{for all valuations } u, w \text{ that are } \delta\text{-close to } v, \quad (5)$$

and, equivalently, if there is a neighborhood U of v such that f is an ε -optimizer for U with respect to F , i.e.,

$$u(f) \geq w(F) - \varepsilon \quad \forall u, w \in U. \quad (6)$$

The next proposition is a simple corollary of the definition of a robust ε -optimizer at a valuation v .

PROPOSITION 2. *If the set F of feasible streams of bounded payoffs has, for every $\varepsilon > 0$, a robust ε -optimizer at a valuation v , then the function $u \mapsto u(F)$ is continuous at v .*

The following example shows that the converse, however, does not hold: there is a bounded set of streams of payoffs for which the u -optimal value is a constant that does not have a robust ε -optimizer at any non-impatient valuation v .

EXAMPLE 1. Let F_1 be the set of all streams $f = (f_1, f_2, \dots)$ with $f_t \in \{-1, 1\}$, $\liminf_{t \rightarrow \infty} \bar{f}_t = -1$, and $\limsup_{t \rightarrow \infty} \bar{f}_t = 1$. $\diamond$

For any valuation v , the v -optimal value of F_1 , $v(F_1)$, equals 1; see Section 8.1. Therefore, the function $v \mapsto v(F_1)$ is a constant function and, hence, continuous. However, if v is a non-impatient valuation, then no $f \in F_1$ is a robust ε -optimizer at v with respect to F_1 for some $\varepsilon > 0$.

The next example shows that the existence of a u -optimizer at any valuation u is insufficient for continuity of the optimal value at a non-impatient valuation.

EXAMPLE 2. Let F_2 be a set that consists of a single stream of payoffs g such that $\liminf_{n \rightarrow \infty} \bar{g}_n + 2\varepsilon < \limsup_{n \rightarrow \infty} \bar{g}_n$, where $\varepsilon > 0$. $\diamond$

The set F_2 consists of a single element. Therefore, it has, for every valuation u , a (unique) u -optimizer. However, it does not have a robust ε -optimizer at any patient valuation v . Moreover, if $v = (1 - \beta)w + \beta u$, where u is a patient valuation, $\beta > 0$, and w is a valuation, then F_2 does not have a robust $\beta\varepsilon$ -optimizer at v .

An important robust optimization property of a set of streams F is that it has a robust ε -optimizer at v for every $\varepsilon > 0$. The following proposition provides a “minmax = maxmin”-type condition on a set F that is equivalent to F having a robust ε -optimizer at v for every $\varepsilon > 0$.

PROPOSITION 3. *The set F has a robust ε -optimizer at v for every $\varepsilon > 0$ if and only if*

$$\sup_{f \in F, W \in \mathcal{N}(v)} \inf_{u \in W} u(f) = \inf_{W \in \mathcal{N}(v)} \sup_{h \in F, u \in W} u(h),$$

where $\mathcal{N}(v)$ denotes the set of all neighborhoods of a valuation v .

3.4 Robust optimization at an impatient valuation

The following proposition shows that a bounded set of streams of payoffs F admits robust optimization at every impatient valuation v . Hence, robust optimization at an impatient valuation is of secondary importance.

PROPOSITION 4. *Let F be a bounded set of streams of payoffs and let v be an impatient valuation. If f is an ε -optimizer for v with respect to F , then, for every $\varepsilon' > \varepsilon \geq 0$, f is a robust ε' -optimizer at v with respect to F . Therefore, F has, for every $\varepsilon > 0$ and every impatient valuation v , a robust ε -optimizer at v .*

3.5 Robust optimization at a patient valuation

A neighborhood of a patient valuation contains, for all sufficiently large n and all sufficiently small r , the n th Cesàro average valuation u_n and the r -discounted valuation u_r . Therefore, if $f \in F$ is a robust ε -optimizer at a patient valuation v with respect to F , then, for all sufficiently large n and all sufficiently small r , $f \in F$ is an ε -optimizer for u_n and for u_r with respect to F , and the oscillation of the u_r -optimal and u_n -optimal values of F is at most ε .

Therefore, the notion of robustness at a patient valuation provides a unifying view of earlier studies of robust optimization of a patient decision maker.

We illustrate the importance of approximate (as opposed to exact) optimizers and the advantage of studying robust optimizers at a patient valuation by considering the following example.

Consider the imprecise specification of an impatient valuation that is obtained by specifying that its averaging weights are sufficiently small, e.g., less than 0.01. This can be modeled as the set U of all impatient valuations that are within a distance of 0.01 from a patient valuation.

The set F of feasible streams of payoffs consists of the perpetuity $\mathbf{1}$, with a constant payoff 1, and the streams f^k , $k \geq 0$, where the payoff is 2 in the first period and in each of the first k even periods, and the payoff is 0 in all other periods. Note that $\mathbf{1}$ is, for every $\varepsilon > 0$, a robust ε -optimizer at any patient valuation with respect to F .

If our objective is to select the “best” stream in F , given that the impatient valuation’s weight on each individual period is less than 0.01, then it seems intuitive that we should select the perpetuity $\mathbf{1}$. This intuition is justified by the observation that the perpetuity $\mathbf{1}$, which is not an optimizer for any impatient valuation in U , is a 0.02-optimizer for¹¹ U (i.e., for any $u \in U$) with respect to F , and (as $u_n \in U$ for $n > 100$ and $u_n(f^k) \rightarrow_{n \rightarrow \infty} 0 < 1 = u_n(\mathbf{1})$) no other stream in F is even a 0.99-optimizer for U with respect to F .

3.6 Robust optimization at a non-impatient valuation

One may argue that impatience is a natural assumption on a preference over streams of payoffs and that it is, therefore, sufficient to confine the analysis to impatient valuations.

However, to model the imprecision in the specification of the impatient valuation, it may be advantageous to fix a non-impatient valuation and then consider all the impatient valuations in its neighborhood.

For example, consider a preference of an impatient decision maker who has a pretty good idea of the “interest rate” between successive points in time, as long as these are not too distant; however, as regards the very distant future, he cannot tell much beyond the fact that the interest rates remain nonnegative; furthermore, he wants to give the very distant future a nonzero weight, say 30%. Such a preference is modeled as an impatient valuation in a sufficiently small neighborhood of a non-impatient valuation that is a mixture of an impatient valuation with weight 70% and a patient valuation with weight 30%.

Preferences that are defined by valuations that are in a sufficiently small neighborhood of a non-impatient valuation arise naturally in decision problems that involve pollution, global warming, etc.

The advantage of using valuations that are not impatient in the description of a small imprecision in the specification of an impatient valuation is analogous to the advantage of using boundary points of a square in the description of a small imprecision in the

¹¹By the time value of money principle, a stream of alternating 0s and 2s is worth at most as much as a constant stream of 1s. For a valuation $u \in U$, the extra 2 in the first period contributes at most $2/100$; hence, we have $u(f^k) \leq u(\mathbf{1}) + 0.02$.

specification of an interior point, e.g., an interior point that is sufficiently close to a fixed boundary point.

3.7 Global robust optimization

The ability to select a robust ε -optimizer at v (with respect to F) that can be changed gradually as the valuation v changes corresponds to the existence of a robust ε -optimizer at v (with respect to F) that varies continuously as a function of the valuation v . Theorem 4 shows that if F has, for any valuation v , a robust ε -optimizer at v with respect to F , then (a) there are finitely many streams of payoffs such that for each valuation, one of them is a robust ε -optimizer and (b) there exists a robust ε -optimizer at v with respect to F that depends continuously on v whenever F is convex.

In this section, we state the implications of a bounded set of streams of payoffs F having a robust ε -optimizer at every valuation v .

THEOREM 4. *Assume that the set F of feasible streams of bounded payoffs has a robust ε -optimizer at every valuation v . Then there is a finite list $f^1, f^2, \dots, f^k$ in F such that the following statements hold:*

- (a) *For every valuation v , there is an index $1 \leq i \leq k$ such that f^i is a robust ε -optimizer at v with respect to F .*
- (b) *There is a continuous function $v \mapsto f^v$ with values in the convex hull of $\{f^1, \dots, f^k\}$ such that every valuation v has a neighborhood U such that $u(f^v) \geq w(F) - \varepsilon$ $\forall u, w \in U$; hence, if f^v is in F , then f^v is a robust ε -optimizer at v with respect to F .*

The next proposition demonstrates that the condition that F has a robust ε -optimizer at every valuation v is essential for the conclusions of Theorem 4 and Proposition 2.

PROPOSITION 5. *For every non-impatient valuation u and a neighborhood U of u , there is a bounded set of streams of payoffs $F \subset c$ such that the following statements hold:*

- (a) *The optimal value of F is not continuous at u . Moreover, there is a sequence of impatient valuations v_n that converges to u such that the sequence $v_n(F)$ does not converge.*
- (b) *The optimal value of F is continuous at any valuation $v \notin U$.*
- (c) *There exists $\eta > 0$ such that for every finite subset $G \subset F$, there is an impatient valuation w such that $w(F) - \eta > 1 + \eta > \max_{g \in G} w(g)$.*

4. ROBUST OPTIMIZATION IN A MARKOV DECISION PROCESS

4.1 Markov decision process

In a discrete-time finite Markov decision process (MDP), play proceeds in stages. At each stage, the process is in one of finitely many states, and the decision maker chooses

an action from a finite set of feasible actions. The action and the current state jointly determine the payoff of the decision maker and the probability of the succeeding state.

Before making the choice, the decision maker observes the current state.

A finite MDP is defined by the list $\Gamma = (S, A, r, p)$, where S is the finite set of states, A is the finite set of actions, $r : S \times A \rightarrow \mathbb{R}$ is the payoff function, and $p : S \times A \rightarrow \Delta(S)$ is the transition function. If action $a \in A$ is taken at stage t and the state in stage t is $s \in S$, then the payoff at stage t is $r(s, a)$ and the (conditional) probability distribution of the state at stage $t + 1$ is $p(s, a)$.

A pure (respectively, behavioral) policy π of the decision maker specifies the action (respectively, the probability distribution over actions) at stage t as a function of the current state and past states and actions. Namely, $\pi : \bigcup_{t \geq 1} (S^t \times A^{t-1}) \rightarrow A$ (respectively, $\rightarrow \Delta(A)$).

Given an initial state $s_1 = s$, a policy π defines a probability distribution P_π^s over the sequences $s_1, a_1, \dots$ of states and actions. The expectation with respect to P_π^s is denoted by E_π^s . For simplicity, we use the same symbol P_π^s to denote also the distribution over the streams of payoffs $g_t = r(s_t, a_t)$.

The set F^s of feasible distributions over streams of payoffs, as a function of the initial state s , is defined by $F^s = \{P_\pi^s : \pi \text{ a behavioral policy}\}$. It equals the convex hull of the sets $\{P_\pi^s : \pi \text{ a pure policy}\}$.

The set $\hat{F}^s$ of feasible streams of payoffs, as a function of the initial state s , is the set of streams of payoffs $\hat{P}_\pi^s = g^{s, \pi}$, where $g_t^{s, \pi} = E_\pi^s r(s_t, a_t)$ and π ranges over all policies in the finite MDP. It equals the convex hull of the sets $\{\hat{P}_\pi^s : \pi \text{ a pure policy}\}$.

THEOREM 5. *Let $\Gamma = (S, A, r, p)$ be a finite MDP. For every probability distribution $q \in \Delta(S)$, the set $\sum_{s \in S} q(s) \hat{F}^s$ has, for every $\varepsilon > 0$ and every valuation v , a robust ε -optimizer at v with respect to $\sum_{s \in S} q(s) \hat{F}^s$.*

Theorem 5 shows that any finite MDP has, for every $\varepsilon > 0$ and every valuation v , a robust ε -optimal policy at v . A mixture of two policies in a MDP is a policy in a MDP; hence, the set of feasible streams of payoffs in a MDP is convex. Therefore, Theorem 4 guarantees the existence, for each $\varepsilon > 0$, of a continuous function that assigns to each valuation v a policy π_v such that π_v is a robust ε -optimal policy at v .

In fact, we prove a stronger result. To state this stronger result, we introduce the following notation. For a valuation u and a stream of payoffs g , we denote by $\underline{u}(g)$, respectively, by $\bar{u}(g)$, the infimum, respectively, the supremum, of $u'(g)$ over all valuations u' that are 0-close to u .

Note that $u(g)$ need not be measurable in g and, therefore, the expectation of $u(g)$ with respect to the probability P_π^s (where π is a policy) need not exist. However, $\underline{u}(g)$ and $\bar{u}(g)$ are measurable in g , and, therefore, the expectation of $\underline{u}(g)$ and $\bar{u}(g)$ with respect to the probability P_π^s exists.

As $u(\hat{P}_\pi^s) \geq E_\pi^s \underline{u}(g)$ and $u(\hat{P}_\pi^s) \leq E_\pi^s \bar{u}(g)$, the next theorem implies Theorem 5.

THEOREM 6. *For any finite MDP, valuation v , and $\varepsilon > 0$, there is a policy π and $\delta > 0$, such that for all valuations u and w that are δ -close to v and any policy σ ,*

$$E_\pi^s \underline{u}(g) \geq E_\sigma^s \bar{w}(g) - \varepsilon.$$

Theorem 6 along with Theorem 4 shows that for any finite MDP, for every $\varepsilon > 0$, there is a continuous map $v \mapsto \pi_v$ from valuations to policies in the MDP such that π_v is a robust von Neumann–Morgenstern ε -optimal policy at v .

The proof of Theorem 6 also proves the following stronger property of a finite MDP. The normed space of all sequences $\omega = (\omega_1, \omega_2, \dots)$ with $\|\omega\|_1 := \sum_{t=1}^{\infty} |\omega_t| < \infty$ is denoted by ℓ_1 .

THEOREM 7. *For any finite MDP, $\omega \in \ell_1$, patient valuation v , and $\varepsilon > 0$, there is a policy π and $\delta > 0$, such that for all valuations u and w that are δ -close to v , any policy σ , and any $\omega' \in \ell_1$ with $\|\omega - \omega'\|_1 < \delta$,*

$$E_{\pi}^s \left(\sum_{t=1}^{\infty} \omega_t g_t + \underline{u}(g) \right) \geq E_{\sigma}^s \left(\sum_{t=1}^{\infty} \omega'_t g_t + \overline{w}(g) \right) - \varepsilon.$$

The proof of Theorem 6 demonstrates, implicitly, how to find a robust ε -optimal policy at a valuation u . We assume without loss of generality that the payoff function of the MDP takes values in $[0, 1]$. The first step is to find a stationary uniformly optimal policy π and the undiscounted value v of the MDP. The second step is to find a positive integer t_{ε} such that $\sum_{t=t_{\varepsilon}}^{\infty} w_t < \varepsilon/2$, where $w_t := v(\mathbf{e}_t)$. The robust ε -optimal Markov policy $\sigma = (\sigma_t)_{t \geq 1}$ plays at stages $t \geq t_{\varepsilon}$ according to the stationary uniformly optimal policy π . The definition of the play of the robust ε -optimal Markov policy σ at stages $t < t_{\varepsilon}$ is defined recursively. Set $v_{t_{\varepsilon}} = v$, $R_t(s_t, a, v_{t+1}) = w_t r(s_t, a) + (1 - \sum_{s \leq t} w_t) \sum_{s' \in S} p(s_t, a)(s') v_{t+1}(s')$, $v_t = \max_a R_t(a, v_{t+1})$, and that $\sigma_t(s_t)$ is an action a that maximizes $R_t(s_t, a, v_{t+1})$.

5. PROOFS OF THE THEOREMS

Note that an additive function $u : \ell_{\infty} \rightarrow \mathbb{R}$ that is monotonic is (by classical arguments) linear. Indeed, by the additivity of u , we have $u(-g) = -u(g)$ and $u(\alpha g) = \alpha u(g)$ for every rational α . By the additivity and monotonicity of u , for every $g, h \in \ell_{\infty}$, $|u(g) - u(h)| \leq \|g - h\|u(\mathbf{1})$ and, therefore, $u(\alpha g)$, $\alpha \in \mathbb{R}$, is continuous in α ; hence, $u(\alpha g) = \alpha u(g) \forall \alpha \in \mathbb{R}$.

5.1 Proof of Theorem 1

Assume that u is an impatient valuation. Define $\omega_t = u(\mathbf{e}_t)$.

By the additivity of u , we have $u(\mathbf{0}) + u(\mathbf{0}) = u(\mathbf{0})$ and, hence, $u(\mathbf{0}) = 0$. The time value of money principle of a valuation along with the definition of ω_t implies that $u(\mathbf{0}) = 0 \leq \omega_t = u(\mathbf{e}_t) \geq u(\mathbf{e}_{t+1}) = \omega_{t+1}$.

Note that $- \|g\| \mathbf{1}_{>n} \leq g - \sum_{t=1}^n g_t \mathbf{e}_t \leq \|g\| \mathbf{1}_{>n}$ and, therefore, using the linearity of u , the definition of ω_t , monotonicity (which follows from the time value of money principle), and the impatience of u , we have

$$\left| u(g) - \sum_{t=1}^n \omega_t g_t \right| = \left| u(g) - u\left(\sum_{t=1}^n g_t \mathbf{e}_t \right) \right| \leq u(\|g\| \mathbf{1}_{>n}) \rightarrow_{n \rightarrow \infty} 0.$$

Therefore, $u(g) = \sum_{t=1}^{\infty} \omega_t g_t$. In particular, using the normalization assumption $u(\mathbf{1}) = 1$, we have $u(\mathbf{1}) = \sum_{t=1}^{\infty} \omega_t = 1$. This completes the proof of the “only if” part of the theorem.

Assume that $u(g) = \sum_{t=1}^{\infty} \omega_t g_t$ with $\omega_t - \omega_{t+1} \geq 0$ and $\sum_{t=1}^{\infty} \omega_t = 1$. Then u is a normalized linear real-valued function on the space ℓ_{∞} with $u(\mathbf{1}_{>n}) = \sum_{t>n} \omega_t \rightarrow_{n \rightarrow \infty} 0$. Since $u(g) = \sum_{t=1}^{\infty} \omega_t g_t = \sum_{t=1}^{\infty} (\omega_t - \omega_{t+1}) t \bar{g}_t$, it follows that if $\bar{g}_t \geq \bar{h}_t \forall t$, then $u(g) \geq u(h)$. This completes the proof of the “if” part of the theorem.

5.2 Proof of Theorem 2

Let u be a patient valuation.

Note that if u is a patient valuation, then $u(\mathbf{e}_t) = 0$. Indeed, by additivity, we have $u(\mathbf{0}) = 0$, and by the monotonicity of a valuation, we have $u(\mathbf{e}_t) \geq 0$. As u is a patient valuation, $1 = u(\mathbf{1}) = u(\sum_{t=1}^n \mathbf{e}_t + \mathbf{1}_{>n}) = \sum_{t=1}^n u(\mathbf{e}_t) + 1$. Therefore, $u(\mathbf{e}_t) = 0 \forall t$.

For $g \in \ell_{\infty}$, we set $\bar{g} := \limsup_{k \rightarrow \infty} \bar{g}_k$ and $\underline{g} := \liminf_{k \rightarrow \infty} \bar{g}_k$.

Let u be a patient valuation and let $g \in \ell_{\infty}$. Fix $\varepsilon > 0$ and let n be sufficiently large so that $\underline{g} - \varepsilon < \bar{g}_k < \bar{g} + \varepsilon \forall k \geq n$.

Let h be defined by $h = \sum_{t=1}^n (\|g\| + \varepsilon) \mathbf{e}_t + (\bar{g} + \varepsilon) \mathbf{1}_{>n}$. Note that for every positive integer s , we have $\bar{h}_s \geq \bar{g}_s$ and, therefore, by the time value of money principle, $u(h) \geq u(g)$.

By the linearity and patience of u , $u(h) = (\bar{g} + \varepsilon)u(\mathbf{1}) = \bar{g} + \varepsilon$. Therefore, $u(g) \leq \bar{g} + \varepsilon$. As this last inequality holds for every $\varepsilon > 0$, we deduce that the right-hand inequality of (4) holds for every patient valuation u and every $g \in \ell_{\infty}$.

Note that the left-hand inequality of (4) holds for $g \in \ell_{\infty}$ if (and only if) the right-hand inequality of (4) holds for $-g$. Indeed, $-u(g) = u(-g) \leq \limsup_{n \rightarrow \infty} -\bar{g}_n = -\liminf_{n \rightarrow \infty} \bar{g}_n$. Therefore, $\underline{g} \leq u(g)$ for every $g \in \ell_{\infty}$.

Assume that u is a linear function that is defined on ℓ_{∞} and satisfies (4). Obviously, $u(\mathbf{1}) = 1$; hence, u is normalized. It remains to show that u satisfies the time value of money principle. Assume that $g, h \in \ell_{\infty}$ with $\sum_{t=1}^n g_t \geq \sum_{t=1}^n h_t \forall n$. Then $\underline{g} - \underline{h} \geq 0$ and, therefore, $u(g - h) \geq 0$ by the left-hand side inequality of (4) and, therefore, as u is linear, $u(g) = u(g - h) + u(h) \geq u(h)$.

5.3 Proof of Theorem 3

Obviously, a convex combination of valuations is a valuation. This proves the straightforward “if” part of the theorem. We proceed to the proof of the “only if” part.

Let u be a valuation and let $\omega_t := u(\mathbf{e}_t)$. As u is a valuation, $\omega_t \geq \omega_{t+1} \geq 0 \forall t$.

As u is additive, $u(\mathbf{0}) = 0$. As u is additive, normalized, and monotonic, $u(\mathbf{1}_{>n})$ is non-increasing in n and $0 \leq u(\mathbf{1}_{>n}) \leq 1$.

Let β be the limit of the non-increasing sequence $u(\mathbf{1}_{>n}) = u(\mathbf{1}) - \sum_{t=1}^n \omega_t$. As $0 \leq u(\mathbf{1}_{>n}) = 1 - \sum_{t=1}^n \omega_t \leq u(\mathbf{1}) = 1$, we have $0 \leq \beta = 1 - \sum_{t=1}^{\infty} \omega_t \leq 1$.

If $\beta = 0$, then u is an impatient valuation.

If $\beta = 1$, then u is a patient valuation.

Assume that $0 < \beta < 1$. Define $w : \ell_\infty \rightarrow \mathbb{R}$ by $w(g) := \sum_{t=1}^{\infty} \frac{\omega_t}{1-\beta} g_t$ and define the function $v : \ell_\infty \rightarrow \mathbb{R}$ by $v(g) := (u(g) - \sum_{t=1}^{\infty} \omega_t g_t) / \beta$.

As $\omega_t \geq \omega_{t+1} \geq 0$ and $\sum_{t=1}^{\infty} \omega_t = 1 - \beta$, w is an impatient valuation.

Obviously, $u = (1 - \beta)w + \beta v$. Therefore, it remains to prove that v is an impatient valuation.

As $u(\mathbf{1}) - \sum_{t=1}^{\infty} \omega_t = \beta$, we have $v(\mathbf{1}) = 1$. Therefore, the function v is normalized.

By the linearity of the function $g \mapsto u(g) - \sum_{t=1}^{\infty} \omega_t g_t$, the function v is linear.

To prove that v is a valuation, it remains to prove that v satisfies the time value of money principle.

By the linearity of v , it suffices to prove that if $g \in \ell_\infty$ with $\sum_1^s g_t \geq 0 \forall s$, then $v(g) \geq 0$.

For $g \in \ell_\infty$ and an integer n , we denote by $g_{>n}$ the element of ℓ_∞ whose t th coordinate equals g_t if $t > n$ and equals 0 if $t \leq n$.

Assume that $\sum_1^s g_t \geq 0 \forall s$. Fix $\varepsilon > 0$. As $\omega_t \geq 0$ and $\sum_{t=1}^{\infty} \omega_t < \infty$, there is a positive integer k such that $(k\omega_k + \sum_{t=k+1}^{\infty} \omega_t) \|g\| < \varepsilon$.

As v is linear and $v(\mathbf{e}_k) = u(\mathbf{e}_k) - \omega_k = 0$, $v(g) = v(g_{>k}) = v(\sum_{t=1}^k g_t \mathbf{e}_k + g_{>k})$. Using the definition of v along with the time value of money principle of u , we have $\beta v(\sum_{t=1}^k g_t \mathbf{e}_k + g_{>k}) = u(\sum_{t=1}^k g_t \mathbf{e}_k + g_{>k}) - (\sum_{t=1}^k g_t \omega_k + \sum_{t=k+1}^{\infty} g_t \omega_t) \geq 0 - \varepsilon \geq -\varepsilon$.

As the inequality $\beta v(g) \geq -\varepsilon$ holds for every $\varepsilon > 0$ and $\beta > 0$, we conclude that $v(g) \geq 0$.

As $v(\mathbf{1}_{>n}) = (u(\mathbf{1}_{>n}) - \sum_{t>n} \omega_t) / \beta \rightarrow_{n \rightarrow \infty} 1$, the valuation v is a patient valuation.

5.4 Proof of Theorem 4

Let $F \subset \ell_\infty$ be a set of feasible streams of bounded payoffs and let $\varepsilon > 0$.

Assume that for every valuation v there is a stream g^v in F that is a robust ε -optimizer at v with respect to F . Let $W_v \in \mathcal{N}(v)$ be a neighborhood of v such that

$$u(g^v) \geq w(F) - \varepsilon \quad \forall u, w \in W_v. \quad (7)$$

As the topological space V of all valuations is compact and the set of neighborhoods W_v covers V (i.e., $\bigcup_{v \in V} W_v = V$), there is a finite subcover. Namely, there are finitely many distinct valuations $v_1, \dots, v_k$ such that $\bigcup_{i=1}^k W_{v_i} = V$. Set $f^i = g^{v_i}$ and let v be a valuation. As $\bigcup_{i=1}^k W_{v_i} = V$, there is an index $1 \leq i \leq k$ such that $v \in W_{v_i}$.

By setting $v = v_i$ and $g^v = f^i$ in inequality (7), we deduce that f^i is a robust ε -optimizer at v with respect to F .

This completes the proof of the first part of the theorem.

Let $\alpha_i : V \rightarrow \mathbb{R}_+$ be a continuous function such that $\alpha_i(v) = 0$ if and only if $v \notin W_{v_i}$. The existence of such a function α_i follows from the fact that V is a pseudo-metrizable space. (For example, $\alpha_i(v)$ can be the distance of v from the complement of W_{v_i} .) Note that for every $v \in V$, there is $1 \leq i \leq k$ such that $v \in W_{v_i}$ and, hence, $\alpha_i(v) > 0$. Therefore, $\sum_{i=1}^k \alpha_i(v) > 0 \forall v \in V$.

Next, we define the stream f^v by

$$f^v = \frac{\sum_{i=1}^k \alpha_i(v) f^i}{\sum_{i=1}^k \alpha_i(v)}.$$

As the functions α_i are continuous and $\sum_{i=1}^k \alpha_i(v) > 0$, the function $v \mapsto f^v$ is continuous. Note that f^v is in the convex hull of $\{f^1, \dots, f^k\}$.

Let U be the neighborhood of v consisting of all valuations u such that for all $1 \leq i \leq k$, $\alpha_i(u) > 0$ if and only if $\alpha_i(v) > 0$. That is, $U = \bigcap_{i: \alpha_i(v) > 0} W_{v_i} = \bigcap_{i: u \in W_{v_i}} W_{v_i}$.

Let u and w be two valuations in U . For any $1 \leq i \leq k$ such that $\alpha_i(v) > 0$, we have $u(f^i) \geq w(F) - \varepsilon \forall u, w \in W_{v_i}$; hence, $u(f^i) \geq w(F) - \varepsilon \forall u, w \in U \subset W_{v_i}$. As u is a linear function of the stream of payoffs, we deduce that $u(f^v) \geq w(F) - \varepsilon \forall u, w \in U$.

This completes the proof of Theorem 4.

5.5 Proof of Theorem 6

Let $\Gamma = (S, A, r, p)$ be a discrete-time finite MDP and let $v(s)$, $s \in S$, be the undiscounted value of the MDP with initial state s .

Set $g_t = r(s_t, a_t)$ and $\bar{g}_n = \frac{1}{n} \sum_{t=1}^n g_t$.

Let π be a stationary uniformly optimal policy¹² of the decision maker in Γ . Thus,¹³ for every state $s \in S$ and every policy η ,

$$E_\pi^s \liminf_{n \rightarrow \infty} \bar{g}_n \geq v(s) \geq E_\eta^s \limsup_{n \rightarrow \infty} \bar{g}_n, \quad (8)$$

and for every $\varepsilon > 0$, there is n_ε such that for every state $s \in S$, every $n \geq n_\varepsilon$, and every policy η ,

$$\varepsilon + E_\pi^s \bar{g}_n \geq v(s) \geq E_\eta^s \bar{g}_n - \varepsilon. \quad (9)$$

Fix a valuation u and let $\omega_t = u(e_t)$, $t \geq 1$, be the weights of the valuation u .

To prove the theorem, it suffices to define, for every $\varepsilon > 0$, a neighborhood U of u and a policy τ , such that for every policy η and every $u^* \in U$,

$$7\varepsilon + \underline{u}^*(P_\tau^s) \geq v(s) \geq \bar{u}^*(P_\eta^s) - 7\varepsilon. \quad (10)$$

Recall that $\sum_{t=1}^\infty \omega_t \leq 1$. Set $\omega_\infty = 1 - \sum_{t=1}^\infty \omega_t$ and let t_ε be a sufficiently large positive integer such that $(1 + \|r\|) \sum_{t=t_\varepsilon}^\infty \omega_t < \varepsilon$, where $\|r\| = \max_{s,a} |r(s, a)|$.

Fix $\varepsilon > 0$.

¹²A uniformly optimal policy is a policy π that is optimal in every discounted MDP with a sufficiently small discount rate. The existence of a stationary uniformly optimal policy in a finite MDP is due to Blackwell (1962).

¹³Properties (8) and (9) are easily derived from the fact that π is a stationary uniformly optimal policy. Alternatively, by the construction of an ε -optimal policy in Mertens and Neyman (1981), it follows that the policy π is, for every $\varepsilon > 0$, an ε -optimal policy in the undiscounted MDP. Alternatively, see Neyman (2003, part 4) of Proposition 3).

Let Γ_* be the multistage decision problem (N, Σ, r_*) , where the set of policies Σ coincides with the set of policies of the MDP and the payoff function r_* , as a function of the initial state s and the policy σ , is defined by

$$r_*(s, \sigma) = E_\sigma^s \sum_{1 \leq t < t_\varepsilon} \omega_t g_t + \left(1 - \sum_{1 \leq t < t_\varepsilon} \omega_t\right) E_\sigma^s v(s_{t_\varepsilon}).$$

The payoff r_* depends only on finitely many periods of the play of Γ . Therefore, Γ_* is equivalent to a decision problem with finitely many pure policies; hence, Γ_* has an optimal pure policy.

Let σ be an optimal policy of Γ_* with payoff vector v_* . Namely,

$$r_*(s, \sigma) = E_\sigma^s \sum_{1 \leq t < t_\varepsilon} \omega_t g_t + \left(1 - \sum_{1 \leq t < t_\varepsilon} \omega_t\right) E_\sigma^s v(s_{t_\varepsilon}) = v_*(s), \quad (11)$$

and for every policy η ,

$$r_*(s, \eta) = E_\eta^s \sum_{1 \leq t < t_\varepsilon} \omega_t g_t + \left(1 - \sum_{1 \leq t < t_\varepsilon} \omega_t\right) E_\eta^s v(s_{t_\varepsilon}) \leq v_*(s). \quad (12)$$

Define the policy τ as follows. At stage $t < t_\varepsilon$, $\tau_t(s_1, a_1, \dots, s_t) = \sigma(s_1, a_1, \dots, s_t)$ and at stage $t \geq t_\varepsilon$, $\tau_t(s_1, a_1, \dots, s_{t_\varepsilon}, \dots, s_t) = \pi(s_t)$.

The definition of the policy τ along inequality (8) implies that

$$E_\tau^s \liminf_{n \rightarrow \infty} \bar{g}_n \geq E_\tau^s v(s_{t_\varepsilon}) = E_\sigma^s v(s_{t_\varepsilon}). \quad (13)$$

Let U be the set of all valuations u^* whose valuation weights $\omega_t^* := u^*(e_t)$ are such that

$$\|r\| \sum_{t=1}^{t_\varepsilon + n_\varepsilon} |\omega_t^* - \omega_t| < \varepsilon. \quad (14)$$

Note that U is a neighborhood of u .

Fix a valuation $u^* \in U$. By the choice of t_ε , we have $\|r\| \sum_{t=t_\varepsilon}^{t_\varepsilon + n_\varepsilon} \omega_t < \varepsilon$, and, therefore, inequality (14) implies that

$$\sum_{t=t_\varepsilon}^{t_\varepsilon + n_\varepsilon} \omega_t^* \|r\| < 2\varepsilon. \quad (15)$$

By equality (11), the definition of τ , the inequality $\omega_t^* g_t \geq \omega_t g_t - \|r\| |\omega_t^* - \omega_t|$, and inequality (14), we have

$$\begin{aligned} E_\tau^s \sum_{1 \leq t < t_\varepsilon} \omega_t^* g_t &= E_\sigma^s \sum_{1 \leq t < t_\varepsilon} \omega_t^* g_t \\ &\geq E_\sigma^s \sum_{1 \leq t < t_\varepsilon} \omega_t g_t - \|r\| \sum_{1 \leq t < t_\varepsilon} |\omega_t - \omega_t^*| \\ &\geq v_*(s) - \left(1 - \sum_{1 \leq t < t_\varepsilon} \omega_t\right) E_\sigma^s v(s_{t_\varepsilon}) - \varepsilon. \end{aligned} \quad (16)$$

Let $t \geq t_\varepsilon + n_\varepsilon$. Then, using inequality (9) and the definition of τ , we have

$$E_\tau(g_{t_\varepsilon} + \dots + g_t \mid \mathcal{H}_{t_\varepsilon}) \geq (t - t_\varepsilon + 1)(v(s_{t_\varepsilon}) - \varepsilon), \quad (17)$$

where $\mathcal{H}_t$ is the algebra (of subsets of plays) that is generated by $s_1, a_1, \dots, s_t$.

By summation by parts and using the inequality $\omega_t^* \geq \omega_{t+1}^* \forall t \geq t_\varepsilon$, we have

$$\sum_{t=t_\varepsilon}^{\infty} \omega_t^* g_t = \sum_{t=t_\varepsilon}^{\infty} (\omega_t^* - \omega_{t+1}^*) \sum_{s=t_\varepsilon}^t g_s \quad (18)$$

and

$$\sum_{t=t_\varepsilon}^{\infty} \omega_t^* = \sum_{t=t_\varepsilon}^{\infty} (\omega_t^* - \omega_{t+1}^*) (t - t_\varepsilon + 1). \quad (19)$$

Therefore, using (18), the triangle inequality, (15), (9), and (19), we have

$$\begin{aligned} E_\tau \left(\sum_{t=t_\varepsilon}^{\infty} \omega_t^* g_t \mid \mathcal{H}_{t_\varepsilon} \right) &= E_\tau \left(\sum_{t=t_\varepsilon}^{\infty} (\omega_t^* - \omega_{t+1}^*) \sum_{s=t_\varepsilon}^t g_s \mid \mathcal{H}_{t_\varepsilon} \right) \\ &\geq E_\tau \left(\sum_{t=t_\varepsilon+n_\varepsilon}^{\infty} (\omega_t^* - \omega_{t+1}^*) \sum_{s=t_\varepsilon}^t g_s \mid \mathcal{H}_{t_\varepsilon} \right) - \sum_{t=t_\varepsilon}^{t_\varepsilon+n_\varepsilon-1} \omega_t^* \|r\| \\ &\geq \sum_{t=t_\varepsilon+n_\varepsilon}^{\infty} (\omega_t^* - \omega_{t+1}^*) (t - t_\varepsilon + 1) (v(s_{t_\varepsilon}) - \varepsilon) - 2\varepsilon \\ &\geq \sum_{t=t_\varepsilon}^{\infty} (\omega_t^* - \omega_{t+1}^*) (t - t_\varepsilon + 1) (v(s_{t_\varepsilon}) - \varepsilon) - 4\varepsilon \\ &= \sum_{t=t_\varepsilon}^{\infty} \omega_t^* (v(s_{t_\varepsilon}) - \varepsilon) - 4\varepsilon \geq \sum_{t=t_\varepsilon}^{\infty} \omega_t^* v(s_{t_\varepsilon}) - 5\varepsilon. \end{aligned}$$

By taking the expectation, we deduce that

$$E_\tau^s \sum_{t=t_\varepsilon}^{\infty} \omega_t^* g_t \geq \sum_{t=t_\varepsilon}^{\infty} \omega_t^* E_\tau^s v(s_{t_\varepsilon}) - 5\varepsilon. \quad (20)$$

Multiplying inequality (8) by $\omega_\infty^* := 1 - \sum_{t=1}^{\infty} \omega_t^*$ and adding inequality (20), we have

$$\begin{aligned} E_\tau^s \omega_\infty^* \liminf_{n \rightarrow \infty} \bar{g}_n + E_\tau^s \sum_{t=t_\varepsilon}^{\infty} \omega_t^* g_t &\geq \left(\omega_\infty^* + \sum_{t=t_\varepsilon}^{\infty} \omega_t^* \right) E_\tau^s v(s_{t_\varepsilon}) - 5\varepsilon \\ &= \left(1 - \sum_{1 \leq t < t_\varepsilon} \omega_t^* \right) E_\tau^s v(s_{t_\varepsilon}) - 5\varepsilon \\ &\geq \left(1 - \sum_{1 \leq t < t_\varepsilon} \omega_t \right) E_\tau^s v(s_{t_\varepsilon}) - 6\varepsilon. \end{aligned} \quad (21)$$

By summing inequalities (16) and (21), we have

$$E_\tau^s \omega_\infty^* \liminf_{n \rightarrow \infty} \bar{g}_n + E_\tau^s \sum_{t=1}^{\infty} \omega_t^* g_t \geq v_*(s) - 7\varepsilon.$$

For any stream of bounded payoffs g , we have $u^*(g) \geq \omega_\infty^* \liminf_{n \rightarrow \infty} \bar{g}_n + \sum_{t=1}^{\infty} \omega_t^* g_t$ by the characterization of valuations (Theorems 1, 2, and 3), and the map $g \mapsto \omega_\infty^* \liminf_{n \rightarrow \infty} \bar{g}_n + \sum_{t=1}^{\infty} \omega_t^* g_t$ is measurable. Therefore,

$$\underline{u}^*(P_\tau^s) \geq E_\tau^s \omega_\infty^* \liminf_{n \rightarrow \infty} \bar{g}_n + E_\tau^s \sum_{t=1}^{\infty} \omega_t^* g_t \geq v_*(s) - 7\varepsilon,$$

which proves the left-hand inequality of (10).

Fix a policy η of the decision maker. By replacing, in the above equations and inequalities, E_τ^s by E_η^s , $=$ and $\geq$ by $\leq$, ε by $-\varepsilon$, and $\liminf$ by $\limsup$, we have

$$\bar{u}^*(P_\eta^s) \leq E_\eta^s \omega_\infty^* \limsup_{n \rightarrow \infty} \bar{g}_n + \sum_{t=1}^{\infty} \omega_t^* g_t^{s, \eta} \leq v_*(s) + 7\varepsilon,$$

which proves the right-hand inequality of (10).

Explicitly, using inequalities (14) and (12), and $\omega_t^* g_t \leq \omega_t g_t + \|g\| |\omega_t^* - \omega_t|$, we have

$$E_\eta^s \sum_{1 \leq t < t_\varepsilon} \omega_t^* g_t = E_\eta^s \sum_{1 \leq t < t_\varepsilon} \omega_t^* g_t \leq v_*(s) - \left(1 - \sum_{t=1}^{\infty} \omega_t\right) E_\eta^s v(s_{t_\varepsilon}) + \varepsilon. \quad (22)$$

By using (18), the triangle inequality, the right-hand inequality of (8), and (19), we have

$$\begin{aligned} & E_\eta \left(\sum_{t=t_\varepsilon}^{\infty} \omega_t^* g_t \mid \mathcal{H}_{t_\varepsilon} \right) \\ &= E_\eta \left(\sum_{t=t_\varepsilon}^{\infty} (\omega_t^* - \omega_{t+1}^*) \sum_{s=t_\varepsilon}^t g_s \mid \mathcal{H}_{t_\varepsilon} \right) \\ &\leq E_\eta \left(\sum_{t=t_\varepsilon+n_\varepsilon}^{\infty} (\omega_t^* - \omega_{t+1}^*) \sum_{s=t_\varepsilon}^t g_s \mid \mathcal{H}_{t_\varepsilon} \right) + \sum_{t=t_\varepsilon}^{t_\varepsilon+n_\varepsilon-1} \|r\| \omega_t^* \\ &\leq \sum_{t=t_\varepsilon+n_\varepsilon}^{\infty} (\omega_t^* - \omega_{t+1}^*) (t - t_\varepsilon + 1) (v(s_{t_\varepsilon}) + \varepsilon) + 2\varepsilon \\ &\leq \sum_{t=t_\varepsilon}^{\infty} (\omega_t^* - \omega_{t+1}^*) (t - t_\varepsilon + 1) (v(s_{t_\varepsilon}) + \varepsilon) + 4\varepsilon \\ &= \sum_{t=t_\varepsilon}^{\infty} \omega_t^* (v(s_{t_\varepsilon}) + \varepsilon) + \varepsilon \leq \sum_{t=t_\varepsilon}^{\infty} \omega_t^* v(s_{t_\varepsilon}) + 5\varepsilon. \end{aligned}$$

By taking the expectation, we deduce that

$$E_\eta^s \sum_{t=t_\varepsilon}^{\infty} \omega_t^* g_t \leq \sum_{t=t_\varepsilon}^{\infty} \omega_t^* E_\eta^s v(s_{t_\varepsilon}) + 5\varepsilon. \quad (23)$$

The uniform optimality of π implies that for every policy η ,

$$E_\eta^s \limsup_{n \rightarrow \infty} \bar{g}_n \leq E_\eta^s v(s_{t_\varepsilon}). \quad (24)$$

Multiplying inequality (24) by $\omega_\infty^* = 1 - \sum_{t=1}^{\infty} \omega_t^*$ and adding inequality (23), we have

$$\begin{aligned} E_\eta^s \omega_\infty^* \liminf_{n \rightarrow \infty} \bar{g}_n + E_\eta^s \sum_{t=t_\varepsilon}^{\infty} \omega_t^* g_t &\leq \left(\omega_\infty^* + \sum_{t=t_\varepsilon}^{\infty} \omega_t^* \right) E_\eta^s v(s_{t_\varepsilon}) + 5\varepsilon \\ &= \left(1 - \sum_{1 \leq t < t_\varepsilon} \omega_t^* \right) E_\eta^s v(s_{t_\varepsilon}) + 5\varepsilon \\ &\leq \left(1 - \sum_{1 \leq t < t_\varepsilon} \omega_t \right) E_\eta^s v(s_{t_\varepsilon}) + 6\varepsilon. \end{aligned} \quad (25)$$

Inequalities (22) and (25) imply that

$$\bar{u}^*(P_\eta^s) \leq \left(1 - \sum_{t=1}^{\infty} \omega_t^* \right) \limsup_{n \rightarrow \infty} \bar{g}_n^{s, \eta} + \sum_{t=1}^{\infty} \omega_t^* g_t^{s, \eta} \leq v_*(s) + 7\varepsilon,$$

which proves the right-hand inequality of (10).

Any valuation u is a mixture of a patient valuation v and an impatient valuation w . If w is impatient, then for any policy π we have $w(g^{s, \pi}) = w(P_\pi^s)$. If v is a patient valuation, then for any policy π we have $\underline{v}(P_\pi^s) \leq v(g^{s, \pi}) \leq \bar{v}(P_\pi^s)$. Therefore, for any valuation u we have $\underline{u}(P_\pi^s) \leq u(g^{s, \pi}) \leq \bar{u}(P_\pi^s)$.

Therefore, Theorem 6 implies Theorem 5, i.e., that the set $\{g^{s, \pi} : \pi \text{ a policy}\}$ has, for every $\varepsilon > 0$ and valuation v , a robust ε -optimizer at v .

Note that the inequalities $\omega_t \geq \omega_{t+1} \geq 0$, $1 \leq t < t_\varepsilon$, were not used in the proof. Therefore, the proof demonstrates that for every finite MDP and a finite sequence of real numbers $\omega_1, \dots, \omega_N$, there is a policy π and neighborhoods U_ε , $\varepsilon > 0$, of the patient valuations such that for any policy η ,

$$E_\pi^s \sum_{t=1}^N \omega_t g_t + E_\pi^s \liminf_{t \rightarrow \infty} \bar{g}_t \geq E_\eta^s \sum_{t=1}^N \omega_t g_t + E_\eta^s \limsup_{t \rightarrow \infty} \bar{g}_t,$$

and for every $u \in U_\varepsilon$,

$$E_\pi^s \sum_{t=1}^N \omega_t g_t + E_\pi^s \underline{u}(g^{s, \pi}) \geq E_\eta^s \sum_{t=1}^N \omega_t g_t + E_\eta^s \bar{u}(g^{s, \pi}) - \varepsilon.$$

6. PROOFS OF THE PROPOSITIONS

6.1 *Proof of Proposition 3*

First, we derive an inequality that does not depend on F having a robust ε -optimizer at v .

Note that for every neighborhood W of v , $\inf_{u \in W} u(F) \leq v(F) \leq \sup_{u \in W} u(F)$. Therefore,

$$\sup_{W \in \mathcal{N}(v)} \inf_{u \in W} u(F) \leq v(F) \leq \inf_{W \in \mathcal{N}(v)} \sup_{u \in W} u(F).$$

As

$$\sup_{h \in F, W \in \mathcal{N}(v)} \inf_{u \in W} u(h) \leq \sup_{W \in \mathcal{N}(v)} \inf_{u \in W} u(F),$$

we conclude that

$$\sup_{h \in F, W \in \mathcal{N}(v)} \inf_{u \in W} u(h) \leq v(F) \leq \inf_{W \in \mathcal{N}(v)} \sup_{u \in W} u(F) = \inf_{W \in \mathcal{N}(v)} \sup_{h \in F, u \in W} u(h).$$

Second, assume that f is a robust ε -optimizer at v with respect to F . Then there is a neighborhood U of v such that for every $u \in U$, we have $u(f) \geq u(F) - \varepsilon$ and $|u(f) - v(F)| \leq \varepsilon$ (and, hence, $u(F) \leq u(f) + \varepsilon \leq v(F) + 2\varepsilon$). Therefore,

$$v(F) - \varepsilon \leq \inf_{u \in U} u(f) \leq \sup_{h \in F, W \in \mathcal{N}(v)} \inf_{u \in W} u(h).$$

Also,

$$\inf_{W \in \mathcal{N}(v)} \sup_{h \in F, u \in W} u(h) \leq \sup_{h \in F, u \in U} u(h) \leq \sup_{u \in U} u(F) \leq v(F) + 2\varepsilon.$$

Therefore,

$$v(F) - \varepsilon \leq \sup_{h \in F, W \in \mathcal{N}(v)} \inf_{u \in W} u(h) \leq \inf_{W \in \mathcal{N}(v)} \sup_{h \in F, u \in W} u(h) \leq v(F) + 2\varepsilon.$$

If F has a robust ε -optimizer at v for every $\varepsilon > 0$, we conclude that

$$\sup_{h \in F, W \in \mathcal{N}(v)} \inf_{u \in W} u(h) = v(F) = \inf_{W \in \mathcal{N}(v)} \sup_{h \in F, u \in W} u(h).$$

In the other direction, assume that

$$\sup_{h \in F, W \in \mathcal{N}(v)} \inf_{u \in W} u(h) = a = \inf_{W \in \mathcal{N}(v)} \sup_{h \in F, u \in W} u(h).$$

The left-hand equality implies that for every $\varepsilon > 0$, there are $f \in F$ and neighborhoods $U \in \mathcal{N}(v)$ such that $u(f) \geq a - \varepsilon/2$ for every $u \in U$. In particular, $v(F) \geq a - \varepsilon/2$.

The right-hand equality implies that for every $\varepsilon > 0$, there is a neighborhood $W \in \mathcal{N}(v)$ such that $u(F) \leq a + \varepsilon/2$ for every $u \in U$. In particular, $v(F) \leq a + \varepsilon/2$.

Therefore, $v(F) = a$ and for every $u \in U \cap W$,

$$v(F) + \varepsilon/2 \geq u(F) \geq u(f) \geq v(F) - \varepsilon/2 \geq u(F) - \varepsilon,$$

and thus f is a robust ε -optimizer at v with respect to F .

6.2 Proof of Proposition 4

Assume that v is an impatient valuation with $v(g) = \sum_{t=1}^{\infty} \omega_t g_t$, where $\omega_t \geq 0$ and $\sum_{t=1}^{\infty} \omega_t = 1$.

Fix $\varepsilon > 0$ and $g^\varepsilon \in F$ with $v(g^\varepsilon) > v(F) - \varepsilon$. We will prove that g^ε is a robust 10ε -optimizer at v with respect to F .

Fix n sufficiently large such that $\sum_{t>n} \omega_t \|F\| < \varepsilon$, where $\|F\| = \sup_{f \in F} \|f\|$. Let W be the neighborhood of v of all valuations u such that $|u(\mathbf{e}_t) - \omega_t| \|F\| < \varepsilon/n \forall t \leq n$.

Then, for every $h \in F$ and $u \in W$, $u(h) \leq u(\sum_{t=1}^n h_t \mathbf{e}_t + \|F\| \mathbf{1}_{>n}) \leq \sum_{t=1}^n \omega_t h_t + \varepsilon + \|F\|(1 - \sum_{t=1}^n \omega_t) \leq \sum_{t=1}^n \omega_t h_t + 3\varepsilon \leq v(h) + 4\varepsilon$. Therefore, $u(F) \leq v(F) + 5\varepsilon$.

Similarly, $u(g^\varepsilon) \geq u(\sum_{t=1}^n g_t^\varepsilon \mathbf{e}_t - \|F\| \mathbf{1}_{>n}) \geq \sum_{t=1}^n \omega_t h_t - \varepsilon - \|F\|(1 - \sum_{t=1}^n \omega_t) \geq \sum_{t=1}^n \omega_t g_t^\varepsilon - 3\varepsilon \geq v(g^\varepsilon) - 4\varepsilon \geq v(F) - 5\varepsilon$.

Therefore, for any $u \in W$, $|u(g^\varepsilon) - v(F)| \leq 5\varepsilon$ and $u(g^\varepsilon) \geq v(F) - 5\varepsilon \geq u(F) - 10\varepsilon$. Therefore, g^ε is a robust 10ε -optimizer at v with respect to F .

6.3 Proof of Proposition 5

Fix a non-impatient valuation u and a neighborhood U of u .

Let $u_1^\perp$ denote the set of all $g \in c$ with $\|g\| = 1$ and $u(g) = 0$.

For every $v \in V \setminus U$, $\sup_{g \in u_1^\perp} v(g) > 0$. For every $\varepsilon > 0$, set $U_\varepsilon = \{v \in V : \sup_{g \in u_1^\perp} v(g) > \varepsilon\} = \bigcup_{g \in u_1^\perp} \{v \in V : v(g) > \varepsilon\}$. As a union of open sets, U_ε is an open set. Note that $U_{\varepsilon'} \supseteq U_\varepsilon$ if $\varepsilon' < \varepsilon$ and $\bigcup_{\varepsilon>0} U_\varepsilon \supseteq V \setminus U$. Therefore, there is $\varepsilon > 0$ such that $U_{\varepsilon'} \supseteq V \setminus U$ for every $\varepsilon' \leq \varepsilon$.

Let $\varepsilon < 1$ be sufficiently small so that $U_\varepsilon \supseteq V \setminus U$. For every $v \in U_\varepsilon$, there is an element $g^v \in u_1^\perp$ and a neighborhood $U_\varepsilon(v)$ of v such that for every $w \in U_\varepsilon(v)$, we have $w(g^v) > \varepsilon$. As $\bigcup_{v \in U_\varepsilon} U_\varepsilon(v) \supseteq V \setminus U$, there is a finite list $v^1, \dots, v^k$ such that $\bigcup_{1 \leq i \leq k} U_\varepsilon(v^i) \supseteq V \setminus U$.

Let $F_\varepsilon(u)$ be the finite set $\{g^{v^i} : 1 \leq i \leq k\}$.

Let $h \in \ell_\infty$ be a stream of payoffs with $\|h\| = \varepsilon$, $\limsup_{t \rightarrow \infty} \bar{h}_t = \varepsilon$, and $\liminf_{t \rightarrow \infty} \bar{h}_t = -\varepsilon$.

Define $u_t = u(\mathbf{e}_t)$ if $t \geq 1$ and $u_0 = 1 - \sum_{t=1}^{\infty} u_t$. As u is a non-impatient valuation, $0 < u_0 \leq 1$.

Let n_ε be sufficiently large so that $\sum_{t>n_\varepsilon} |u_t| < u_0 \varepsilon/4$.

Let H be the set of all streams of payoffs h^n , $n > n_\varepsilon$, where $h_t^n = h_t$ if $n_\varepsilon < t \leq n$ and $h_t^n = 0$ otherwise.

Let g be the stream of payoffs where $g_t = u_t / \sum_{t=0}^{\infty} u_t^2$ if $t \leq n_\varepsilon$ and $g_t = (u_0 - \sum_{t=n+1}^{\infty} u_t) / \sum_{t=0}^{\infty} u_t^2$ if $t > n_\varepsilon$.

Let $F = (g + H) \cup (g + F_\varepsilon(u))$.

To prove (a), (b), and (c), it suffices to construct a sequence of impatient valuations w_n that converges to u and a positive number $\eta > 0$, such that for any finite subset G of F , $\limsup_{n \rightarrow \infty} \sup_{f \in F} \omega_n(f) > \eta + \limsup_{n \rightarrow \infty} \max_{g \in G} \omega_n(g)$.

As u is a non-impatient valuation, $u = (1 - u_0)w + u_0v$, where v is a patient valuation. Therefore, the sequence of impatient valuations $w_n := (1 - u_0)w + u_0\gamma_n$, where $\gamma_n(g) = \bar{g}_n$, converges to u .

By the definition of g , it follows that

$$1 \geq u(g) = \sum_{t=0}^n u_t^2 / \sum_{t=0}^{\infty} u_t^2 - \left(\sum_{t=n+1}^{\infty} u_t \right)^2 / \sum_{t=0}^{\infty} u_t^2 > 1 - u_0\varepsilon/16.$$

By the properties of h , there are sequences of integers n_m that converge to infinity such that $\lim_{m \rightarrow \infty} w_{n_m}(g + h^{n_m}) = u(g) + u_0\varepsilon$. Therefore, $w_{n_m}(g + h^{n_m}) \geq 1 + u_0\varepsilon/2$ for all sufficiently large m .

For every $n_1 \in \mathbb{N}$ and $f \in F_\varepsilon(u)$, we have $\lim_{n \rightarrow \infty} w_n(g + h^{n_1}) = u(g + h^{n_1}) \leq 1 + u_0\varepsilon/4$ and $\lim_{n \rightarrow \infty} w_n(g + f) \leq 1$. Therefore, for every finite subset G of F , $\limsup_{n \rightarrow \infty} \max_{g \in G} w_n(g) \leq 1 + u_0\varepsilon/4$. This completes the proof of properties (a) and (c) of the set F .

Let $v \in U_\varepsilon$. Define $\alpha(F, v) := v(g) + \max_{f \in F_\varepsilon(u)} v(f)$. By the definitions of U_ε and of the finite set $F_\varepsilon(u)$, there is $f^* \in F_\varepsilon(u)$ such that $\alpha(F, v) = v(g + f^*) > v(g) + \varepsilon \geq v(g + h) \forall h \in H$.

Fix $0 < \eta < v(f^*) - \varepsilon$. Let $U(v)$ be the set of all valuations w such that $|w(g) - v(g)| + |w(f) - v(f)| < \eta$ for all $f \in F_\varepsilon(u)$. As $F_\varepsilon(u)$ is a finite subset of c and g is a fixed element of c , $U(v)$ is a neighborhood of v .

The definitions of $U(v)$ and f^* imply that if $w \in U(v)$, then $w(g + f) \leq v(g + f^*) + \eta = \alpha(F, v) + \eta$ for all $f \in F_\varepsilon(u)$.

As $\|h\| \leq \varepsilon$ for every $h \in H$, the properties of f^* and η imply that $w(g + h) \leq w(g) + \varepsilon \leq v(g) + \varepsilon + \eta < \alpha(F, v)$ for all $h \in H$.

The definitions of F , $U(v)$, and f^* imply that $w(F) \geq w(g + f^*) \geq v(g + f^*) - \eta = \alpha(F, v) - \eta$.

As F is the union of $g + H$ and $g + F_\varepsilon(u)$, and $g + f^* \in F$, we conclude that $\alpha(F, v) - \eta \leq w(F) \leq \alpha(F, v) + \eta$. This completes the proof of property (b) of the set F .

7. PROOFS OF THE LEMMAS

7.1 Proof of Lemma 1

Assume that $g, h \in \ell_\infty$ with $\sum_{t=1}^s h_t \geq \sum_{t=1}^s g_t \forall s$ and let $u : \ell_\infty \rightarrow \mathbb{R}$ be a monotonic, impatient, and additive function that satisfies $w_t := u(\mathbf{e}_t) \geq u(\mathbf{e}_{t+1})$. Then, as in the proof of Theorem 1, $u(g) = \sum_{t=1}^\infty w_t g_t = \sum_{t=1}^\infty (w_t - w_{t+1}) t \bar{g}_t \leq \sum_{t=1}^\infty (w_t - w_{t+1}) t \bar{h}_t = \sum_{t=1}^\infty h_t w_t = u(h)$.

7.2 Proof of Lemma 2

Fix $g \in \ell_\infty$. Let U be the one-dimensional subspace of ℓ_∞ that is spanned by g . Let φ be the linear functional on U that satisfies $\varphi(g) = \bar{g}$; hence, $\varphi(\theta g) = \theta \bar{g} \forall \theta \in \mathbb{R}$.

Define the function $p : \ell_\infty \rightarrow \mathbb{R}$ by the equality $p(h) = \bar{h}$. Then p is sublinear (i.e., $p(g + h) \leq p(g) + p(h)$ and $p(\theta g) = \theta p(g)$ for all $g, h \in \ell_\infty, \theta \in \mathbb{R}_+$) and $\varphi(h) \leq p(h) = \bar{h}$ for all $h \in U$.

Therefore, by the Hahn–Banach theorem, there is a linear functional u on ℓ_∞ such that $u(h) \leq p(h) = \bar{h} \forall h \in \ell_\infty$ and $u(g) = \varphi(g) = \bar{g}$.

It remains to show that $\underline{h} \leq u(h)$ for all $h \in \ell_\infty$, which follows from $\underline{h} = \liminf_{n \rightarrow \infty} \bar{h}_n = -\limsup_{n \rightarrow \infty} -\bar{h}_n = -\limsup_{n \rightarrow \infty} \overline{(-h)}_n = -p(-h) \leq -u(-h) = u(h)$.

Applying the above-proved part to the element $-g$ of ℓ_∞ shows that there is a linear functional v on ℓ_∞ such that $v(-g) = \limsup_{n \rightarrow \infty} -\bar{g}_n$; hence, $v(g) = \liminf_{n \rightarrow \infty} \bar{g}_n$ and $v(-h) \geq \liminf_{n \rightarrow \infty} -\bar{h}_n$; hence, $v(h) \leq \bar{h} \forall h \in \ell_\infty$ and $v(h) = -v(-h) \geq -\overline{(-h)} = \underline{h}$.

7.3 Proof of Lemma 3

Define the following two linear operators on ℓ_∞ . The linear operator $O : \ell_\infty \rightarrow \ell_\infty$ is defined by the equality $Oh = (h_1, h_3, h_5, \dots)$, i.e., $(Oh)_t = h_{2t-1}$, and the linear operator $E : \ell_\infty \rightarrow \ell_\infty$ is defined by the equality $Eh = (h_2, h_4, h_6, \dots)$, i.e., $(Eh)_t = h_{2t}$.

Let $0 \leq g \in \ell_\infty$ with $\underline{g} < \bar{g}$. Let u and v be two patient valuations such that $u(g) = \underline{g}$ and $v(g) = \bar{g}$.

Therefore, $u(g) - v(g) < 0$ and $u(\mathbf{e}_t) = v(\mathbf{e}_t) = 0 \forall t$.

Define the function $w : \ell_\infty \rightarrow \mathbb{R}$ by $w(h) = u(Oh)/2 + u(Eh)/2$. We claim that w is normalized, linear, monotonic, and satisfies $w(\mathbf{e}_t) \geq w(\mathbf{e}_{t+1})$, but w does not satisfy the time value of money principle.

First, note that $u \circ O$ and $u \circ E$ are normalized, linear, and monotonic, and, therefore, so is their average w . As $w(\mathbf{e}_t) = 0 \forall t$, $0 = w(\mathbf{e}_t) \geq w(\mathbf{e}_{t+1}) = 0 \forall t$.

Next, define h by $Oh = g$ and $Eh = -g$, i.e., $h = (g_1, -g_1, g_2, -g_2, \dots)$. Note that $\sum_{t=1}^{2n} h_t = 0$ and that $\sum_{t=1}^{2n-1} h_t = g_n \geq 0$. But $2w(h) = u(Oh) + u(Eh) = u(g) - v(g) < 0$. Therefore, w does not satisfy the time value of money principle.

8. PROOFS OF THE PROPERTIES OF THE SETS IN THE EXAMPLES

8.1 Properties of the set F_1 in Example 1

Let v be a patient valuation. We will prove¹⁴ that $v(F_1) = 1$.

Let $n_k > 0$, $k \geq 0$, be an increasing sequence of positive integers such that $\lim_{k \rightarrow \infty} n_k / n_{k+1} = 0$. Let j be a positive integer and let f^i , $0 \leq i < j$, be the stream of payoffs with $f_t^i = 1$ if $n_k < t \leq n_{k+1}$ and $k = i \bmod j$, and $f_t^i = 0$ otherwise.

Note that $\sum_{0 \leq i < j} f^i = \mathbf{1}_{>n_0}$, $\mathbf{1} - 2f^i \in F_1$, and $v(f^i) \geq 0$. Therefore, as $v(\sum_{0 \leq i < j} f^i) = v(\mathbf{1}_{>n_0}) = 1$, there is i such that $v(f^i) \leq 1/j$ and, hence, $v(\mathbf{1} - 2f^i) \geq 1 - 2/j$. Therefore, $v(F_1) = 1$.

Obviously, by the definitions of the n th Cesàro average u_n and the set F_1 , for any $f \in F_1$, we have $\liminf_{n \rightarrow \infty} u_n(f) = \liminf_{n \rightarrow \infty} \bar{f}_n = -1$. Therefore, no $f \in F_1$ is a robust 1-optimizer at v with respect to F_1 .

Similarly, if v is a non-impatient valuation, then, by choosing n_0 sufficiently large, we deduce that $v(F_1) = 1$, and no $f \in F_1$ is a robust ε -optimizer at v with respect to F_1 whenever $\varepsilon < \lim_{n \rightarrow \infty} v(\mathbf{1}_{>n})$.

¹⁴We thank Bruno Ziliotto for the proof.

8.2 Properties of the set F_2 in Example 2

Let u be a non-impatient valuation. Then $u = (1 - \beta)w + \beta v$, where w is an impatient valuation, v is a patient one, and $\beta > 0$.

The impatient valuations $(1 - \beta)w + \beta u_n$, where u_n is the n th Cesàro average valuation, converge, as $n \rightarrow \infty$, to the valuation u .

Recall that $F_2 = \{f\}$ and $\liminf_{n \rightarrow \infty} \bar{f}_n + 2\varepsilon = \liminf_{n \rightarrow \infty} u_n(f) + 2\varepsilon < \limsup_{n \rightarrow \infty} \bar{f}_n = \limsup_{n \rightarrow \infty} u_n(f)$.

Then $\liminf_{n \rightarrow \infty} ((1 - \beta)w + \beta u_n)(f) = (1 - \beta)w(f) + \beta \liminf_{n \rightarrow \infty} \bar{f}_n < (1 - \beta)w(f) + \beta \limsup_{n \rightarrow \infty} \bar{f}_n - 2\beta\varepsilon = \limsup_{n \rightarrow \infty} ((1 - \beta)w + \beta u_n)(f) - 2\beta\varepsilon$. Therefore, f is not a robust $\beta\varepsilon$ -optimizer at v with respect to F_2 .

APPENDIX A: IMPATIENT ROBUST OPTIMIZATION

Confining the theory of robust optimization to impatient valuations leads to the following modification of the definition of a robust optimizer.

For any set U , we denote by U^* the set of all impatient valuations in U . Let v be a valuation and let $\varepsilon \geq 0$. A small imprecision in the specification of an impatient valuation is modeled as the set of impatient valuations in a small neighborhood of a valuation v , and v need not be an impatient valuation.

A stream f in F is an *impatient-robust ε -optimizer at v with respect to F* if there is a neighborhood U of v such that

$$u(f) \geq w(F) - \varepsilon \quad \forall u, w \in U^*.$$

A robust ε -optimizer at v with respect to F is obviously an impatient-robust ε -optimizer at v with respect to F . We now show that the converse holds as well.

Fix a stream f and a neighborhood U of a valuation v . The infimum of $u(f)$ over all $u \in U^*$ equals the infimum of $u(f)$ over all $u \in U$, and the supremum of $w(f)$ over all $w \in U^*$ equals the supremum of $w(f)$ over all $w \in U$. Therefore, if f is an impatient-robust ε -optimizer at v with respect to F , then $u(f) \geq w(F) - \varepsilon$ for all $u, w \in U$; hence, f is a robust ε -optimizer at v with respect to F .

In the robustness result for a finite MDP, we alluded to stringent robustness conditions that are called for when the decision maker chooses between different feasible distributions over streams of payoffs. We introduce the formal definition.

Let $\mathcal{P}$ be a set of distributions P over streams of payoffs. For every valuation u and distribution P , we denote by $\underline{u}(P)$ the expectation of $\underline{u}(f)$ with respect to the distribution P , and we denote by $\bar{u}(P)$ the expectation of $\bar{u}(f)$ with respect to the distribution P . The supremum of $\bar{u}(P)$ over all $P \in \mathcal{P}$ is denoted by $\bar{u}(\mathcal{P})$. Let v be a valuation.

A distribution P in $\mathcal{P}$ is a *robust ε -optimizer at v with respect to $\mathcal{P}$* if there is a neighborhood U of v such that

$$\underline{u}(P) \geq \bar{w}(\mathcal{P}) - \varepsilon \quad \forall u, w \in U.$$

A distribution P in $\mathcal{P}$ is an *impatient-robust ε -optimizer at v with respect to $\mathcal{P}$* if there is a neighborhood U of v such that

$$\underline{u}(P) \geq \bar{w}(\mathcal{P}) - \varepsilon \quad \forall u, w \in U^*.$$

A robust ε -optimizer at v with respect to $\mathcal{P}$ is obviously an impatient-robust ε -optimizer at v with respect to $\mathcal{P}$. The converse need not hold.

For example, let g be a stream of payoffs with $\liminf_{n \rightarrow \infty} \bar{g}_n = -1 < \limsup_{n \rightarrow \infty} \bar{g}_n = 1$. Let $\mathcal{P}$ be the set consisting of the single distribution P with $P(g) = 1/2 = P(-g)$. For any impatient valuation u , $u(P) = 0 = \underline{u}(P)$. Therefore P is a V^* -robust ε -optimizer in $\mathcal{P}$. In particular, for any valuation v , P is an impatient-robust v - ε -optimizer in $\mathcal{P}$. However, if w is a patient valuation, then $\bar{w}(P) = 1$. Therefore, if v is a patient valuation, P is not a robust v - ε -optimizer P in $\mathcal{P}$.

APPENDIX B: THE CONTINUOUS-TIME THEORY

In continuous-time theory, a bounded stream of payoffs is a bounded measurable function $[0, \infty) \ni t \mapsto g_t \in \mathbb{R}$. The linear space of bounded streams of payoffs is denoted by L_∞ , and $\mathbf{1}_{\leq T}$ is the stream g with $g_t = 1$ if $t \leq T$ and $g_t = 0$ if $t > T$. Similarly, one defines $\mathbf{1}$ and $\mathbf{1}_{>T}$ in analogy to the definitions in the discrete-time case.

A *valuation* is an additive function $v: L_\infty \rightarrow \mathbb{R}$ that is normalized, i.e., $v(\mathbf{1}) = 1$, and satisfies the time value of money principle: if $\int_0^T g_t dt \geq \int_0^T h_t dt \forall T \geq 0$, then $v(g) \geq v(h)$. A valuation v is *impatient* if $v(\mathbf{1}_{>T}) \rightarrow_{T \rightarrow \infty} 0$; it is *patient* if $v(\mathbf{1}_{>T}) = 1 \forall T$ (equivalently, $v(\mathbf{1}_{>T}) \rightarrow_{T \rightarrow \infty} 1$).

The characterizations of impatient valuations, patient valuations, and valuations are analogous to those in the discrete-time case.

A real-valued function u that is defined on L_∞ is an impatient valuation if and only if there is a function $[0, \infty) \ni t \rightarrow w_t \in \mathbb{R}$, with $\int_0^\infty w_t dt = 1$, that is non-increasing on $(0, \infty)$ and such that

$$u(g) = \int_0^\infty g_t w_t dt.$$

A real-valued function u that is defined on L_∞ is a patient valuation if and only if it is a linear function on L_∞ such that

$$\liminf_{T \rightarrow \infty} \frac{1}{T} \int_0^T g_t dt \leq u(g) \leq \limsup_{T \rightarrow \infty} \frac{1}{T} \int_0^T g_t dt.$$

A real-valued function u that is defined on L_∞ is a valuation if and only if it is a convex combination of an impatient valuation and a patient one.

Similarly, the analogous results of the other theorems and propositions hold also in the continuous-time case.

The topology on the valuation in the continuous-time case is the minimal one where for every $g \in C$, where C consists of all elements $g \in L_\infty$ such that the limit $\lim_{t \rightarrow \infty} g_t$ exists, the function $v \mapsto v(g)$ is continuous.

REFERENCES

Asheim, Geir B., Tapan Mitra, and Bertil Tungodden (2012), “Sustainable recursive social welfare functions.” *Economic Theory*, 49, 267–292. [316]

- Banerjee, Kuntal (2014), “On the representation of preference orders on sequence spaces.” *Social Choice and Welfare*, 43, 497–506. [316]
- Banerjee, Kuntal and Tapan Mitra (2018), “On Wold’s approach to representation of preferences.” *Journal of Mathematical Economics*, 79, 65–74. [316]
- Basu, Kaushik and Tapan Mitra (2003), “Aggregating infinite utility streams with inter-generational equity: The impossibility of being Paretian.” *Econometrica*, 71, 1557–1563. [316]
- Beardon, Alan F. and Ghanshyam B. Mehta (1994), “The utility theorems of Wold, Debreu and Arrow–Hahn.” *Econometrica*, 62, 181–186. [304]
- Blackwell, David (1962), “Discrete dynamic programming.” *Annals of Mathematical Statistics*, 2, 724–738. [327]
- Brock, William A. (1970), “An axiomatic basis for the Ramsey–Weizsäcker overtaking criterion.” *Econometrica*, 38, 927–929. [304]
- Brown, Donald J. and Lucinda M. Lewis (1981), “Myopic economic agents.” *Econometrica*, 49, 359–368. [304]
- Chambers, Christopher P. and Federico Echenique (2018), “On multiple discount rates.” *Econometrica*, 86, 1325–1346. [304, 309]
- Chichilnisky, Garciela (1996), “An axiomatic approach to sustainable development.” *Social Choice and Welfare*, 13, 231–257. [304, 309, 312, 313]
- Chichilnisky, Graciela (1997), “What is sustainable development?” *Land Economics*, 73, 467–491. [304, 309]
- Debreu, Gerard (1959), “Topological methods in cardinal utility theory.” Cowles Foundation Discussion Paper 76, Cowles Foundation for Research in Economics, Yale University. [307, 309]
- Debreu, Gerard (1960), “Topological methods in cardinal utility theory.” In *Mathematical Methods in the Social Sciences* (Kenneth J. Arrow, Samuel Karlin, and Patrick Suppes, eds.), 16–26, Stanford University Press, Stanford. [309, 314]
- Diamond, Peter A. (1965), “The evaluation of infinite utility streams.” *Econometrica*, 33, 170–177. [304, 309, 314, 315, 316]
- Doyle, John R. (2013), “Survey of time preference, delay discounting models.” *Judgment and Decision Making*, 8, 116–135. [308]
- Drugeon, Jean-Pierre and Thai Ha Huy (2021), “A not so myopic axiomatization of discounting.” *Economic Theory*, 73, 1645–1692. [304, 309]
- Falk, Armin, Anke Becker, Thomas Dohmen, Benjamin Enke, David Huffman, and Uwe Sunde (2018), “Global evidence on economic preferences.” *Quarterly Journal of Economics*, 133, 1645–1692. [308]

- Fishburn, Peter. C. and Ward Edwards (1997), “Discount-neutral utility models for denumerable time streams.” *Theory and Decision*, 43, 139–166. [308, 309]
- Fishburn, Peter C. (1966), “Additivity in utility with denumerable product sets.” *Econometrica*, 34, 500–503. [304, 307, 309, 310, 314]
- Forges, Francois (1986), “Correlated equilibria in repeated games with lack of information on one side: A model with verifiable types.” *International Journal of Game Theory*, 15, 65–82. [315]
- Galor, Oded and Ömer Özak (2016), “The agricultural origins of time preference.” *American Economic Review*, 106, 3064–3103. [308]
- Koopmans, Tjalling C. (1960), “Stationary ordinal utility and impatience.” *Econometrica*, 28, 287–309. [304, 309, 314, 315]
- Koopmans, Tjalling C. (1972), “Representation of preference orderings over time.” In *Decision and Organization* (C. Bartlett McGuire and Roy Radner, eds.), 79–100, North-Holland, Amsterdam. [304, 309, 314]
- Koopmans, Tjalling C., Peter A. Diamond, and Richard E. Williamson (1964), “Stationary utility and time perspective.” *Econometrica*, 32, 82–100. [304, 309, 314, 315]
- Lauwers, Luc (1995), “Time-neutrality and linearity.” *Journal of Mathematical Economics*, 24, 347–351. [304, 309, 315]
- Martin, Donald A. and Robert M. Solovay (1970), “Internal Cohen extensions.” *Annals of Mathematical Logic.*, 2, 143–178. [317]
- Mertens, Jean-Francois and Abraham Neyman (1981), “Stochastic games.” *International Journal of Game Theory*, 10, 53–66. [327]
- Mitra, Tapan and M. Kemal Ozbek (2013), “On representation of monotone preference orders in a sequence space.” *Social Choice and Welfare*, 41, 473–487. [316]
- Neyman, Abraham (2003), “From Markov chains to stochastic games.” In *Stochastic Games and Applications* (Abraham Neyman and Sylvain Sorin, eds.), 9–25, Kluwer Academic Publishers, Amsterdam. [327]
- Olson, Mancur and Martin J. Bailey (1981), “Positive time preference.” *Journal of Political Economy*, 89, 1–25. [304, 308, 309]
- Rothbard, Murray N. (1990), “Time preference.” In *The New Palgrave* (Murray Milgate, John Eatwell, and Peter Newman, eds.), 270–275, Macmillan, London. [308]
- Svensson, Lars-Gunnar (1980), “Equity among generations.” *Econometrica*, 48, 1251–1256. [316]
- von Böehm-Bawerk, Eugen (1912), *Positive Theorie des Kapitaless*. Dritte Auflage, especially Buch IV, Abschnitt I, Gegenwart und Zukunft in der Wirtschaft, 426–486. English Translation in *Capital and Interest*, Vol. II, *Positive Theory of Capital*, Book IV, Section I, 257–289, South Holland, Illinois, 1959. [304, 308]

von Weizsäcker, Carl Christian (1967), “Lemmas for a theory of approximate optimal growth.” *Review of Economic Studies*, 34, 143–151. [304]

Wold, Herman O. A. (1943–1944), “A synthesis of pure demand analysis, I, II, and III.” *Scandinavian Actuarial Journal*, 26, 85–118, 220–263, 69–120. [304]

Zame, William (2007), “Can intergenerational equity be operationalized?” *Theoretical Economics*, 2, 187–202. [317]

Co-editor Federico Echenique handled this manuscript.

Manuscript received 17 March, 2021; final version accepted 20 February, 2022; available online 3 March, 2022.