May 21, 2016

PLANNING FOR THE LONG RUN: PROGRAMMING WITH PATIENT, PARETO RESPONSIVE PREFERENCES

URMEE KHAN, MAXWELL B. STINCHCOMBE

ABSTRACT. Ethical social welfare functions treat generations equally and respect Pareto improvements to non-null sets of generations. There are 'perfect' optima for such preferences: generations facing the same circumstances choose the same actions; and all circumstances are treated as possible. If each stationary policy determines a unique long-run distribution that is independent of the starting point, then there is a recursive formulation of policies that are simultaneously optimal for all of the preferences studied here. When the ergodic distribution can depend on early events (hysteresis), the curvature of the social welfare function determines the risks that society is willing to undertake and leads to a variant of the precautionary principle.

Date: May 21, 2016. Many thanks to: Aloisio Araujo for pointing out the connections between this work and general equilibrium theory; to Svetlana Boyarchenko for many discussions about irreversible decisions; to Hector Chade, Amanda Friedenberg and Alejandro Manelli for ganging up and forcing one of the authors to abandon a silly insistence on "pragmatism"; to Paolo Piaquadio, Marcus Pivato, Federico Echenique and Chris Chambers for early encouragement and illuminating conversations; to seminar audiences at USC, ASU, UT Austin, Carlos III, Iowa, Kansas State and participants at SSCW 2014 and ERG Winter Workshop 2015 for many helpful questions and discussions.

Key words and phrases: intergenerational equity and Pareto responsiveness; long-run optimality in Markovian decision problems;

Generations to come

> . . . intergenerational solidarity is not optional, but rather a basic question of justice, since the world we have received also belongs to those who will follow us. (Pope Francis [22])

> As we peer into society's future, we — you and I, and our government — must avoid the impulse to live only for today, plundering for our own ease and convenience the precious resources of tomorrow. We cannot mortgage the material assets of our grandchildren without risking the loss also of their political and spiritual heritage. We want democracy to survive for all generations to come . . . . (Eisenhower, 1953)

Timescales and choices

> With 500 million years left of acceptable habitat for humans on Earth, population being stable at 10 billion with an average length of life equal to 73 years, the ratio of people who will potentially live in the future to people living now is approximately 10 million to 1. (Asheim [5])

> If you are planning for a year, sow rice; if you are planning for a decade, plant trees; if you are planning for a lifetime, educate people. (Attributed to Confucius, 500 B.C.)

1. INTRODUCTION

Choices we make today may have enormous impacts on future generations and the world they will live in. Analyses of intergenerational allocations of costs and benefits have a long history. In that tradition, we are interested in societal choices that involve a combination of the following: choice of actions that can have consequences that far outlast the lifespan of decision makers, ethical questions about intergenerational equity, and elements of irreversibility. Our aim is to put forward an approach that, while grappling with the substantive analytical issues, is nonetheless tractable and applicable in optimization settings. We take as our starting point, and generalize upon, a series of earlier attempts to reason about these issues. Our generalized approach is consonant with the values commonly expressed in analyses of equity, and meets the more pragmatic criteria of tractability and applicability. We begin with an example to provide some context to the issues involved.

1.1. **Present Threats and an Old Analysis.** The looming possibility of drastic change to the climate equilibrium and the associated easy access resources from oceans and forests may be a threat to civilization as we know it. We still believe, or hope, that the expected duration of human society is much longer than the timescale of the decisions that affect this possibility. In the presence of decisions with extremely long-lasting effects and the associated mis-match of timescales, notions of patient preferences for long-run optimality become attractive criteria for decision problems that affect society, society being conceived of as an aggregate of the generations that make it up.

The analysis of intergenerational allocations and welfare has a long history, from which we take, as starting point, Sèbastien Le Prestre de Vauban's *Traité de la Culture des Forêts* (in e.g. [42] or [43]), written in the late 1600's. Vauban, Louis XIV's defense minister, noted, during his wide travels, that several aspects of the economics and ecology of forests complicate the analysis of good societal practices for foresty: first, forests, being a free or easy access resource, were systematically over-exploited; second, after replanting, forests start being productive in slightly less than 100 years but don't become fully productive for 200 years; and third, no private enterprise can have so long and multigenerational a time horizon. From these observations, Vauban concluded that the only institutions that could, and should, undertake such projects in society's interest were the government, in the form of the monarchy at the time, and the church.[1] The calculations behind his conclusion assumed that society would be around for at least the next 200 years to enjoy the net undiscounted benefits.

From Vauban's summing of undiscounted costs and benefits as a way of expressing concern for the welfare of future generations, we take the following: if τ represents the random time until the end of society and $\mathbf{u} = (u_0, u_1, u_2, \ldots)$ is a sequence of numerical measures of different generations, then

$$L_\tau(\mathbf{u}) := E\, \tfrac{1}{\tau+1} \textstyle\sum_{t=0}^{\tau} u_t \tag{1}$$

[1]As well as the government and the church, Vauban also argued for the possibility that, in some settings, a market-like solution to the various incentives problem might be found by making large enough stakes in a forest inheritable but not divisible.

is a measure of society's welfare that, conditional on $\tau = t$, treats all generations equally. If $\mathtt{Prob}(\tau \leqslant M)$ is small for large M, then $L_\tau(\cdot)$ is a measure of welfare for a patient society, one confident in its longevity. To capture the $10^7:1$ ratio between people in the present generation and people who may live in future generations cited above, we take limits of the $L_\tau(\cdot)$'s in (1) as $\mathtt{Prob}(\tau \leqslant M)$ goes to 0 for all M.

This class of limits have representations as integrals against a subclass of strongly translation invariant, purely finitely additive measures, and this property will be crucial for our work. Having positive multiples of these limits as tangents is one of two characterizations we offer for the concave social welfare functions studied here. This tangent characterization is the most useful for applications. The second characterization of societal preferences arises from a a set of axioms suitable for studying optimality while maintaining a concern for intergenerational equity. As such, the second characterization provides the axiomatics behind the first one.

In appliction to models involving intergenerational externalities, the class of preferences under consideration gives rise to an issue of underselectiveness in optimization—there are many optima, some of them ethically objectionable. We solve this problem by invoking a conditional equal treatment property that selects a subset of the optima. This can be seen as a form of 'perfectness', that is, as a limit of optima of perturbed versions of the model. The solution takes a recursive form that can be seen as a vanishing discount limit of the usual Bellman equations.

1.2. **Intergenerational Ethics.** There is a long-held view that it is not acceptable to slight future generations (e.g. [38], [35]), and by considering a 'patient' society in the limit, we follow Ramsey [36] in that, "we do not discount later enjoyments in comparison with earlier ones, a practice which is ethically indefensible and arises merely from the weakness of the imagination." Ramsey's theory posited a "Bliss" point, reachable in finite time, and examined the optimality equations for this point. These optimality equations are less generally applicable than one might hope; Chakravarty [14] showed that, in quite simple examples, the divergence of infinite horizon integrals/sums can lead to the Ramsey optimality equations being satisfied by feasible plans with minimal long-run utility.

To overcome such problems, Weiszäcker [46] formulated the notion of overtaking optimality — a path is overtaking optimal if its accumulated benefits are eventually weakly higher than any other feasible path (see Brock [12] for an axiomatization of these preferences). The overtaking criterion captures a notion of patience and has been extensively used in studies of growth theory, especially the 'turnpike' properties of deterministic optimal paths. This literature was surveyed, extended and unified in McKenzie [32], and the first general existence result for overtaking optimality in convex problems was given by Brock and Haurie [11].

We are primarily interested in intergenerational choice problems where uncertainty is central. The concavity of the social welfare functions we study has two roles: applied to distributions over streams of utility, it captures risk aversion; applied to deterministic streams of utility, it captures a preference for smoothing. The property of having tangents in the class described above can be formulated in terms of a variant of the overtaking criterion studied by Denardo and Miller [16]: overtaking of the intergenerational *average* of accumulated benefits.

Our main axiom for these preferences is the following: for any two paths, $\mathbf{u}$ and $\boldsymbol{v}$, of measures of generational utilities, if the long-run average of the difference $\mathbf{u} - \boldsymbol{v}$ is above and strictly bounded away from 0, then $\mathbf{u}$ is strictly prefered to $\boldsymbol{v}$. Unlike the classical overtaking criterion, the average overtaking criterion is immune to bounded (and more) permutations of the "names" of the generations.[2] In the study of intergenerational allocations, indifference between permuted sequences of measures of generational well-being is called, varyingly, "equity," or "weak anonymity," or "intergenerational neutrality."

There is a set of results demonstrating "incompatibility" between this equity requirement and the Pareto principle. We briefly discuss them here and propose a way forward.

1.3. **On the Pareto Criterion.** Diamond [17] showed that there is no continuous function on the space of sequences of utilities that is indifferent to uniformly bounded permutations and also satisfies a version of the Pareto principle. Basu and Mitra [7] and Fleurbaey and Michel [21] showed that this incompatibility extends to all real-valued social welfare functions, continuous or not, and Asheim [5] provides an extensive review of this literature. From the perspective of criteria based on average overtaking, this conflict arises from treating utility improvements accruing to null coalitions as being strict improvements in social welfare.

We normalize sequences of generational measures of well-being to be non-negative, suppose that they are bounded, and denote by $\mathbf{W}$ the result class of sequences.[3] An intergenerational allocation $\mathbf{u}$ in $\mathbf{W}$ strictly average overtakes an allocation $\boldsymbol{v}$

[2] If $\mathbf{u} = (5, 8, 0, 8, 0, \ldots)$ and $\boldsymbol{v} = (0, 8, 0, 8, 0, \ldots)$, then $\sum_{t \leqslant T}(u_t - v_t) \equiv 5$ so that $\mathbf{u}$ overtakes $\boldsymbol{v}$, but by permuting $\boldsymbol{v}$ to $\boldsymbol{v}^{\pi} = (8, 0, 8, 0, 8, 0, \ldots)$, $\sum_{t \leqslant T}(u_t - v^{\pi}_t)$ is equal to -3 for $T = 0, 2, 4, \ldots$ and is equal to $+5$ for $T = 1, 3, 5, \ldots$. By contrast, $\mathbf{u}$, $\boldsymbol{v}$, and $\boldsymbol{v}^{\pi}$ all have the same long-run average, 4.

[3] See Blackorby et al. [10] for a discussion of this normalization.

if $\liminf_T \frac{1}{T+1}\sum_{t=0}^{T}(u_t - v_t) > 0$ and average overtakes it if this lim inf is non-negative. The allocation $\mathbf{0} = (0, 0, 0, \ldots)$ can never strictly average overtake any element of $\mathbf{W}$, but there is an important subclass of allocations that it does weakly overtake, those with benefits accruing to a null coalition.

For a coalition of generations B and a stream of measures of well-being $\mathbf{u}$, $\mathbf{u}+r1_B$ represents giving a utility bump of r to every member of B. A coalition B is a null coalition if $\mathbf{0}$ average overtakes $r1_B$, equivalently, if $\mathbf{u}$ average overtakes $\mathbf{u} + r1_B$, and it is a non-null coalition if the allocation $r1_B$ strictly average overtakes $\mathbf{0}$. There is an intimate connection between this view of coalitions of generations and Hildenbrand's [24] foundational treatment of Pareto optimality with measure spaces of agents — there, as here, the appropriate Pareto criterion involves ignoring null coalitions and recognizing non-null coalitions. Our main axiom — if $\mathbf{u}$ strictly overtakes $\mathbf{v}$ on average, then intergenerational preferences strictly prefer $\mathbf{u}$ to $\mathbf{v}$ — delivers Pareto responsiveness for our social welfare functions: for any non-null coalition B and any utility bump $r > 0$, $\mathbf{u} + r1_B$ strictly average-overtakes $\mathbf{u}$, hence $\mathbf{u}+r1_B$ is strictly prefered; for any null coalition, $\mathbf{u}+r1_B$ average-overtakes $\mathbf{u}$ which in turn average-overtakes $\mathbf{u} + r1_B$, hence $\mathbf{u}$ and $\mathbf{u} + r1_B$ are indifferent.

1.4. **Methodology.** Our methodological point of view is that one cannot fully understand a class of preferences without knowing their implications in the analysis of problems of interest.[4] For this reason, we study the implications of our social welfare functions in three classes of applications: general equilibrium models with infinite time horizons; Markovian decision problems; irreversible decisions with long-run implications. These applications are chosen for their usefulness in highlighting key aspects of the overall problem we set out to solve.

In general equilibrium models with an infinite time horizon, one typically makes assumptions — on preferences in exchange economy models and on preferences and technologies in production economy models — to guarantee that equilibria have prices that can be represented as integrals against countably additive probabilities. Bewley [9] identifies the requisite assumptions as "an asymptotic form of impatience," Brown and Lewis [13] gave a characterization of the form of "myopia." Araujo [2, Thm. 3] sharpened this result, showing that the combination of

[4]The motivation for Asheim [3] and [4] is to examine the implications of equitable intergenerational preferences by examining their behavior in well-understood models. See also [6, p. 206], "By applying ethical criteria to concrete economic models, we learn about their consequences, and this may change our views about their attractiveness."

countably additive and purely finitely additive properties in the tangents to a preference relation can lead to the non-existence of Pareto optima in these models.

We Work with preferences having tangents in the subclass of purely finitely additive measures identified above. This yields an existence result as well as strong variants of the First and Second Welfare Theorems. In terms of dynasty interpretations of infinitely-lived economic agents, the different weights assigned to coalitions of future generations by our tangents give the tradeoffs between the welfare of different non-null subsets of the future generations within the dynasties. Bewley [8] shows that in such dynastic economies, the less patient dynasties end up trading away their long-run utilities and in equilibrium, achieve a stationary utility of zero. Our results with patient dynasties can have the same results if the different dynasties weight different parts of the future differently.

General equilibrium theory is rarely focused on the analysis of externalities, our primary interest. To capture the one-way flow of externalities, both positive and negative, from earlier to later generations, we turn to a class of problems known as Markovian decision problems (MDPs). In these models: the utility at any given time is a function of the state of the system and the choice of present action; externalities are encoded as states of a system; and the present state and present choice of actions determine the distribution of future states. There is a vast literature on the existence and characterization of stationary solutions to MDPs that maximize the long-run average reward (see, for e.g., [1], [27], [19], [45]).

We will see that the patient welfare functions studied here are particularly well-behaved on the class of ergodic sequences, those that arise from the stationary policies that maximize the long run average payoff. Therefore, the set of utility functionals to which the long-run average MDP existence and characterization results apply include the concave, patients ones that we study. Further, the concavity of the social welfare functions allows for a more complete treatment of Markovian decision problems with irreversibilities, and this yields a variant of the precautionary principle.

The insistence on social indifference to boons given to null subsets of a patient society has implications with the general name of underselectiveness. This arises because the translation invariant, purely finitely additive measures that represent the tangents to our social welfare functions are indifferent to, inter alia, ignoring the misbehavior or the mistreatment of any finite number of generations so long as the effects are not irreversible. Present profligacy arises if the early generations squander resources, leaving a barely recoverable mess for the future. Following

Chichilnisky [15], a "dictatorship of the future" arises if the early generations sacrifice nearly all of their own consumption for the purposes of arriving at a richer future more quickly. Chichilnisky's proposal to alleviate this problem is to use the class of nonstationary preferences that Araujo [2, Thm. 3] used to show the non-existence of Pareto optimal allocations in general equilibrium models. We give an alternate, ethically sound method to circumvent both the profligacy and the dictatorship, one that has a conditional equal treatment property.

A solution demonstrates conditional equal treatment if it treats all states as possible and then solves for actions on the assumption that different generations have equal weight. For us, the existence of patient optima with the conditional equal treatment properties outweighs the existence of alternative, non-stationary, ethically suspect optima.

1.5. **Outline.** §2 provides two complementary treatments of our class of social preferences, as those satisfying a set of Postulates, as those with a special class of tangents. §3 treats the recursive solvability of a wide class of MDPs using our patient preferences. The subsequent section gives four applications, one fits into the recursively solvable class of MDPs, the others demonstrate further the range of implications of the preferences. The final section discusses and concludes.

Throughout, Theorems are about our class of preferences. Corollaries and Propositions concern applications of our preferences.

2. Patient, Inequality Averse Social Welfare Functionals

The inequality averse intergenerational social welfare functions that we work with satisfy a strong form of patience/anonymity. One can develop this class of functions axiomatically or one can specify that the tangents have desirable properties. We begin with the axiomatics, but one could as easily start with the class of tangents.

2.1. **Notation and Setting.** Realizations of intergenerational streams of well-being belong to $\mathbf{W}$, the non-negative elements of ℓ_∞. By assumption, ℓ_∞ is equipped with the sup norm, $\|\mathbf{u}\| := \sup_{t\in\mathbb{N}_0} |u_t|$ where $\mathbf{u} = (u_0, u_1, \ldots) = (u_t)_{t\in\mathbb{N}_0}$ where $\mathbb{N}_0 := \{0, 1, 2, \ldots\}$. The sup norm distance is $d(\mathbf{u}, \mathbf{v}) := \|\mathbf{u} - \mathbf{v}\|$. The interior of $\mathbf{W}$ is denoted $\mathrm{int}(\mathbf{W})$, and $\mathbf{u} \in \mathrm{int}(\mathbf{W})$ if and only if $\inf_t u_t \geqslant r$ for some strictly positive r.

To allow for the study of stochastic dynamic intergenerational problems, the domain for preferences is the mixture set, $\mathcal{M}$, of countably additive Borel probabilities on $\mathbf{W}$ having bounded support, $(\exists B)[p(\{\mathbf{u} \in \mathbf{W} : \|\mathbf{u}\| \leqslant B\}) = 1]$. By assumption, the set $\mathcal{M}$ is given the weak* topology.

2.1.1. *Resultants.* We study 'risk averse expected utility' preferences on $\mathcal{M}$. These are the preferences that can be represented by $p \succ q$ if and only if $\int_{\mathbf{W}} S(\mathbf{u})\,dp(\mathbf{u}) > \int_{\mathbf{W}} S(\mathbf{u})\,dq(\mathbf{u})$ for a continuous, concave intergenerational welfare function $S(\cdot)$. To give the axiomatic form of the inequality aversion contained in concavity, we will use the **resultant** or **expectation** of a $p \in \mathcal{M}$ is denoted $\mathbf{r}(p)$. This is the infinite dimensional version of the expectation of a vector in $\mathbb{R}^k$, and it is defined as the unique point $\mathbf{r}(p) \in \mathbf{W}$ satisfying $\int \langle \mathbf{v}, \mathbf{y}\rangle\,dp(\mathbf{v}) = \langle \mathbf{r}(p), \mathbf{y}\rangle$ for all $\mathbf{y} \in \ell_1$.

2.1.2. *Tangents.* A concave function $S : \mathbf{W} \to [0, \infty)$ has a non-empty set of tangents, denoted $\mathbf{D}S(\mathbf{u})$, at any interior $\mathbf{u}$. The set of tangents determine properties of $S(\cdot)$ — if $L \in \mathbf{D}S(\mathbf{u})$, then concavity implies that for any $\mathbf{v} \in \mathbf{W}$, $S(\mathbf{u}) + L(\mathbf{v} - \mathbf{u}) \geqslant S(\mathbf{v})$. Knowing the class of differences, $\mathbf{v} - \mathbf{u}$, for which each $L(\cdot)$ is non-negative gives general property of $S(\cdot)$.

Tangents are continuous linear functionals on ℓ_∞, and as such, each has an integral representation, $L(\mathbf{u}) = \int_{\mathbb{N}_0} u_t\,d\gamma(t)$, denoted $\langle \mathbf{u}, \gamma\rangle$, where γ is a signed, finitely additive measure on $\mathbb{N}_0$ having finite (variation) norm, $\|\gamma\| := \sup_{\|\mathbf{u}\| \leqslant 1} |L(\mathbf{u})| < \infty$. A net (generalized sequence) γ_α of measures $\mathbb{N}_0$ converges in the weak*-topology to γ if $\langle \mathbf{u}, \gamma_\alpha\rangle \to \langle \mathbf{u}, \gamma\rangle$ for all $\mathbf{u} \in \ell_\infty$. By Alaoglu's theorem, the weak* closure of a norm-bounded set of measures on $\mathbb{N}_0$ is weak* compact.

2.1.3. *Permutations.* We will define the patience of a social welfare functional using indifference to a class of permutations. The set of integers, negative and non-negative is denoted $\mathbb{Z}$ and defined as $\{\ldots, -2, -1, 0, 1, 2, \ldots\}$. A **permutation** is a 1-to-1 function $\pi : \mathbb{N}_0 \to \mathbb{Z}$ that is onto $\mathbb{N}_0$. Given $\mathbf{u} = (u_0, u_1, u_2, \ldots) \in \ell_\infty$ and a permutation π, define $\mathbf{u}^\pi$ as $(u_{\pi^{-1}(0)}, u_{\pi^{-1}(1)}, u_{\pi^{-1}(2)}, \ldots)$. A permutation is **finite** if $\pi(T) = T$ for all but finitely many generations, it is **bounded** if $|\pi(T) - T|$ is uniformly bounded, and it is an $\mathcal{o}(T)$**-permutation** if $\limsup_T \frac{|\pi(T)-T|}{T+1} = 0$. This last class of permutations are closely related to a variant of the overtaking criterion.

2.1.4. *Overtaking.* For $\mathbf{u}, \mathbf{v} \in \ell_\infty$, $\mathbf{u}$ $\mathcal{o}(T)$**-overtakes** $\mathbf{v}$, written $\mathbf{u} \succsim_{\mathcal{o}(T)} \mathbf{v}$, if

$$\liminf_{T\to\infty} \tfrac{1}{T+1} \textstyle\sum_{t=0}^{T} (u_t - v_t) \geqslant 0, \tag{2}$$

and $\mathbf{u}$ **strictly** $\mathcal{o}(T)$**-overtakes** $\mathbf{v}$, written $\mathbf{u} \succ_{\mathcal{o}(T)} \mathbf{v}$, if

$$\liminf_{T\to\infty} \tfrac{1}{T+1} \textstyle\sum_{t=0}^{T} (u_t - v_t) > 0. \tag{3}$$

The classical definition of $\mathbf{u}$ overtaking $\mathbf{v}$ requires the stronger, unaveraged condition $\liminf_{T\to\infty} \sum_{t=0}^{T}(u_t - v_t) \geqslant 0$, but the classical strict overtaking criterion, $\sum_{t=0}^{T}(u_t - v_t) \geqslant \epsilon$ for some strictly positive ϵ and for all sufficiently large T is not sufficient for $\mathcal{o}(T)$-strict overtaking.[5] The $\mathcal{o}(T)$-overtaking criterion is, for obvious reasons, called "average overtaking" in the literature. Lemma 1 (below) gives our reason for using the name "$\mathcal{o}(T)$-overtaking."

2.1.5. *Ergodicity.* The **ergodic** subclass of ℓ_∞ is denoted **Erg** and defined as the set of $\mathbf{u} \in \ell_\infty$ for which the long run average, $\mathbf{lra}(\mathbf{u}) := \lim_{T\to\infty} \frac{1}{T+1}\sum_{t=0}^{T} u_t$, exists. The $\mathcal{o}(T)$-permutations, overtaking and the ergodic class are tightly related.

Lemma 1. *If π is an $\mathcal{o}(T)$-permutation, then for all $\mathbf{u} \in \ell_\infty$, $\mathbf{u} \succsim_{\mathcal{o}(T)} \mathbf{u}^\pi \succsim_{\mathcal{o}(T)} \mathbf{u}$, and for all $\mathbf{u} \in$ **Erg**, $\mathbf{lra}(\mathbf{u}) = \mathbf{lra}(\mathbf{u}^\pi)$.*

Thus, $\mathcal{o}(T)$-overtaking is preserved under $\mathcal{o}(T)$-permutations, a class than strictly includes the finite and the bounded permutations.

2.2. **Formulation by Postulates.** Our assumptions on social preferences are given in terms of a binary relation $\succ$ on the set of probabilities $\mathcal{M}$. We always assume that $\succ$ is **asymmetric**, that is, if $p \succ q$, then it is not the case that $q \succ p$. Define $\sim$ and $\succsim$ on $\mathcal{M}$ by $p \sim q$ if neither $p \succ q$ nor $q \succ p$, and $p \succsim q$ by $p \succ q$ or $p \sim q$. As usual, $\succ$ is **negatively transitive** if for all $p, q, r \in \mathcal{M}$, $[p \succ r] \Rightarrow [\,[p \succ q] \vee [q \succ r]\,]$, and we call a negatively transitive $\succ$ an **asymmetric weak order**.

2.2.1. *Desiderata.* We have two desiderata for our intergenerational preferences: they should be patient, defined as indifference to $\mathcal{o}(T)$-permutations; and they should be exactly Pareto, increases in the utility of null coalitions has no effect on intergenerational preferences while uniform increases accruing to non-null coalitions have a strictly positive effect.

For $\mathbf{u} \in \mathbf{W}$ and $q \in \mathcal{M}$, we write "$\mathbf{u} \succ q$" for $p \succ q$ where $p(\{\mathbf{u}\}) = 1$, with the same convention for "$\succsim$" and "$\sim$." In particular, we write $\mathbf{u} \succ$ (resp. $\succsim$, resp. $\sim$)$\mathbf{v}$ for $p \succ$ (resp. $\succsim$, resp. $\sim$)q where $p(\{\mathbf{u}\}) = 1$ and $q(\{\mathbf{v}\}) = 1$. In this fashion, we restrict $\succ$ to $\mathbf{W}$ by identifying points $\mathbf{u}$ in $\mathbf{W}$ with the associated point masses/Dirac measures, $\delta_{\mathbf{u}}$, in $\mathcal{M}$.

Definition 2.1. *A preference relation $\succ$ on $\mathbf{W}$ is* ***patient*** *if for all $\mathcal{o}(T)$ permutations π and all $\mathbf{u} \in \mathbf{W}$, $\mathbf{u} \sim \mathbf{u}^\pi$.*

[5]See McKenzie [32] for a survey of models of long-run optimal economic growth using the overtaking criterion, Marinacci [31] and Fleurbaey [21] for axiomatic treatments of patience, Asheim [5] for a survey.

The vector $\mathbf{0} = (0, 0, 0, \ldots)$ can never strictly overtake any element of $\mathbf{W}$, but there is an important subclass that $\mathbf{0}$ does overtake. We say that $B \subset \mathbb{N}_0$ is a **null coalition** if $\mathbf{0} \succsim_{\mathcal{O}(T)} 1_B$ and it is a **non-null coalition** if $1_B \succ_{\mathcal{O}(T)} \mathbf{0}$. The appropriate Pareto criterion for patient preferences involves ignoring null coalitions and recognizing non-null coalitions.

Definition 2.2. *A preference relation $\succ$ on $\mathbf{W}$ is **exactly Pareto** if*

(a) *for all $\mathbf{u} \in \mathbf{W}$, all null coalitions B and all $r > 0$, $\mathbf{u} + r1_B \sim \mathbf{u}$, and*
(b) *for all $\mathbf{u} \in \mathbf{W}$, all non-null coalitions B and all $r > 0$, $\mathbf{u} + r1_B \succ \mathbf{u}$.*

2.2.2. *Postulates.* We will invoke the following postulates on $\succ$.

Postulate I. **Weak Order**. $\succ$ is an asymmetric weak order.
Postulate II. **Independence**. For all $p, q, r \in \mathcal{M}$ and all $\alpha \in (0, 1)$, if $p \succ q$, then $\alpha p + (1-\alpha)r \succ \alpha q + (1-\alpha)r$.
Postulate III. **Continuity**. For all $q \in \mathcal{M}$, the sets $\{p \in \mathcal{M} : p \succ q\}$ and $\{p \in \mathcal{M} : p \prec q\}$ are open.
Postulate IV. **Risk and Inequality aversion**. For any $p \in \mathcal{M}$, $\mathbf{r}(p) \succsim p$.
Postulate V. **Respect for overtaking.** $[\mathbf{u} \succ_{\mathcal{O}(T)} \boldsymbol{v}] \Rightarrow [\mathbf{u} \succ \boldsymbol{v}]$.

Discussion. The first three axioms are standard in the expected utility theory of choice under uncertainty: Postulate I is the usual ordering assumption for preference relations; Postulate II is the "linearity in probabilities" assumption for the existence of an expected utility representation for $\succ$; and Postulate III guarantees that the representation is continuous. For expected utility theory, Postulate IV guarantees risk aversion in the form of concavity of the expected utility function, here the concavity of the social expected utility function implies that social preferences are inequality averse on $\mathbf{W}$ and risk averse on $\mathcal{M}$. As noted, Postulate V loosens the classical overtaking criterion. It also directly implies that indifference sets can have no interior: for any $\mathbf{u} \in \mathbf{W}$ and $\epsilon > 0$, $\boldsymbol{v} := \mathbf{u} + \epsilon \cdot 1_{\mathbb{N}_0}$ strictly overtakes $\mathbf{u}$ and $d(\mathbf{u}, \boldsymbol{v}) = \epsilon$.

2.2.3. *Representation.* The next result gives the representation theorem for preferences satisfying the postulates.

Theorem A. *The asymmetric weak order $\succ$ satisfies Postulates I-V if and only if there exists a continuous, concave $S : \mathbf{W} \to [0, \infty)$ such that $[p \succ q] \Leftrightarrow [\int S(\mathbf{u})\, dp(\mathbf{u}) > \int S(\mathbf{u})\, dq(\mathbf{u})]$ with $S(\cdot)$ satisfying the following properties,*

(1) *it is patient, for all $\mathbf{u} \in \mathbf{W}$ and all $\mathcal{O}(T)$-permutations, $S(\mathbf{u}) = S(\mathbf{u}^\pi)$, and*

(2) *it is exactly Pareto, for all* $\mathbf{u} \in \mathbf{W}$ *and all* $r > 0$*, if* B *is a null coalition, then* $S(\mathbf{u} + r1_B) = S(\mathbf{u})$*, and if* B *is a non-null coalition, then* $S(\mathbf{u} + r1_B) > S(\mathbf{u})$.

This result stands in stark contrast to the long history of "impossibility" results for combining patience and the Pareto criterion in inter-generational social welfare analyses. To see what is involved, fix a concave $S(\cdot)$ representing a social ordering satisfying Postulates I-V. For any $\mathbf{u} \in \mathbf{W}$, any $\mathcal{O}(T)$-permutation π, and any $\alpha \in [0,1]$, Lemma 1 tells us that $\mathbf{u} \succsim_{\mathcal{O}(T)} (\alpha\mathbf{u} + (1-\alpha)\mathbf{u}^\pi) \succsim_{\mathcal{O}(T)} \mathbf{u}^\pi \succsim_{\mathcal{O}(T)} \mathbf{u}$. Continuity and respect for overtaking yield the following indifference condition,

$$S(\mathbf{u}) = S(\alpha\mathbf{u} + (1-\alpha)\mathbf{u}^\pi) = S(\mathbf{u}^\pi). \tag{4}$$

Constancy of $S(\cdot)$ on the line joining an interior $\mathbf{u}$ and $\mathbf{u}^\pi$ implies that any $L \in \mathbf{D}S(\mathbf{u})$ must satisfy $L(\mathbf{u} - \mathbf{u}^\pi) = 0$, i.e. $L(\mathbf{u}) = L(\mathbf{u}^\pi)$. Positive continuous linear functionals on ℓ_∞ that are invariant for *finite* permutations have representations as integrals against purely finitely additive positive measures on $\mathbb{N}_0$. A measure η on $\mathbb{N}_0$ is purely finitely additive if and only if for each $\boldsymbol{\nu}$ with $\nu_t \to 0$, $\int \nu_t \, d\eta(t) = 0$. The connection with the classical Pareto criterion comes from taking a strictly positive $\boldsymbol{\nu}$ with $\nu_t \to 0$ and comparing $\mathbf{u}$ and $\mathbf{u} + \boldsymbol{\nu}$. Purely finitely additive measures put mass one "far to the right" of $\mathbb{N}_0$, and for this measure space of agents, $\boldsymbol{\nu}$ is an almost everywhere null endowment — it is non-negative but integrates to 0. The impossibility results arise because they insist on assigning a strictly positive utility to the addition of each such null endowment.

2.3. **Formulation by Tangents.** Let τ be a random variable with $\sum_{T \in \mathbb{N}_0} P(\tau = T) = 1$. If we interpret τ as the random time at which society ends and $P(\tau < M) < \epsilon$ for large M and small ϵ, then the mapping $\mathbf{u} \mapsto L_\tau(\mathbf{u}) := E \, \frac{1}{\tau+1} \sum_{t=0}^{\tau} u_t$ is a measure of welfare for a patient society, one confident in its longevity. The mappings $L_\tau(\cdot)$ are also continuous, linear, positive, and have norm 1. Therefore, for each distribution of τ, there exists a unique probability η_τ such that $L_\tau(\mathbf{u}) = \langle \mathbf{u}, \eta_\tau \rangle$ for all $\mathbf{u}$.

2.3.1. *The Class* $\mathbb{V}$. We denote by $p\mathbb{V}$ the set of weak* accumulation points of the linear functionals L_τ as $P(\tau < M) \to 0$ for all M,

$$p\mathbb{V} = \bigcap \{ \mathrm{cl}\left(\{\eta_\tau : \mathrm{Prob}(\tau \leqslant M) \leqslant \epsilon\}\right) : M \in \mathbb{N}_0, \epsilon > 0\}. \tag{5}$$

The class $p\mathbb{V}$ is non-empty (by the finite intersection property of the class of sets in equation (5)), as well as compact and convex.

A concave function on $\mathbf{W}$ has a non-empty set of tangents at every interior point of $\mathbf{W}$. Our tangents will belong to the following class.

Definition 2.3. *The class $\mathbb{V}$ is the closed convex cone generated by $p\mathbb{V}$.*

2.3.2. *Interior Points and $\mathbb{V}$-Concavity.* Continuous concave functions may fail to have tangent functions at the boundary points of their domain. The Cobb-Douglas functions $f(x_1, x_2) = x_1^{\alpha} x_2^{(1-\alpha)}$ for $0 < \alpha < 1$ and $x_1, x_2 \geqslant 0$ are a case in point. These considerations indicate the need for care in formulating a subset of the concave functions on $\mathbf{W}$ in terms of the existence of tangents.

Definition 2.4. *A function $S : \mathbf{W} \to [0, \infty)$ is $\mathbb{V}$-**concave on** $\text{int}(\mathbf{W})$ if it is continuous and for all $\mathbf{u} \in \text{int}(\mathbf{W})$, the set of tangent functions at $\mathbf{u}$ is a closed, non-empty, norm bounded set of strictly positive elements of $\mathbb{V}$. A function is $\mathbb{V}$-**concave** if the same condition holds for all $\mathbf{u} \in \mathbf{W}$.*

The following $\mathbb{V}$-concave function on $\mathbf{W}$,

$$S_{p\mathbb{V}}(\mathbf{u}) := \min_{\eta \in p\mathbb{V}} \langle \mathbf{u}, \eta \rangle, \tag{6}$$

has a more familiar expression $S_{p\mathbb{V}}(\mathbf{u}) = \liminf_{T\to\infty} \frac{1}{T+1} \sum_{t=0}^{T} u_t$. Because $S_{p\mathbb{V}}(\cdot)$ is the minimum of a collection of linear (not merely affine) functionals, it is homogenous of degree 1.

It is the strict positivity in Definition 2.4 that rules out indifference sets having an interior. To see this take $\mathbf{u} \in \mathbf{W}$ and the interior point $\mathbf{v} := \mathbf{u} + \epsilon \cdot 1_{\mathbb{N}_0}$. Because $\mathbf{v}$ is interior, $\mathbf{D}S(\mathbf{v}) \neq \emptyset$. For any $L \in \mathbf{D}S(\mathbf{v})$, we have $S(\mathbf{v}) + L(\mathbf{u} - \mathbf{v}) \geqslant S(\mathbf{u})$, and $L(\mathbf{u} - \mathbf{v}) = -\epsilon L(1_{\mathbb{N}_0}) < 0$. Combining, we have $d(\mathbf{u}, \mathbf{v}) = \epsilon$ and $S(\mathbf{v}) > S(\mathbf{u})$.

2.3.3. *The Tangents Characterization.* The following relates $\mathbb{V}$-concave functions and preferences satisfying our Postulates.

Theorem B. *If $S : \mathbf{W} \to [0, \infty)$ is $\mathbb{V}$-concave, then the expected utility preferences it represents satisfy Postulates I-V, and if a continuous concave $S : \mathbf{W} \to [0, \infty)$ represents preferences satisfying Postulates I-V, then it is $\mathbb{V}$-concave on $\text{int}(\mathbf{W})$.*

For $\mathbf{u}, \mathbf{v} \in \mathbf{Erg}$, we have seen that $\mathbf{u} \succ_{o(T)} \mathbf{v}$ if and only if $\mathbf{lra}(\mathbf{u}) > \mathbf{lra}(\mathbf{v})$. This implies that the social welfare functions under study have a particularly simple structure when restricted to the class $\mathbf{Erg} \subset \mathbf{W}$ and to probabilities that put mass 1 on $\mathbf{Erg}$.

Corollary B.1. *If a social ordering $\succ$ satisfies Postulates I-V, then*

(1) *for all $\mathbf{u}, \mathbf{v} \in$ **Erg**, $\mathbf{u} \succ \mathbf{v}$ if and only if **lra**$(\mathbf{u}) >$ **lra**$(\mathbf{v})$, and*
(2) *there exists a strictly increasing, concave $\varphi : [0, \infty) \to [0, \infty)$ such that for all $p, q \in \mathcal{M}$ satisfying $p(\mathbf{Erg}) = q(\mathbf{Erg}) = 1$, $p \succ q$ if and only if $\int \varphi(\mathbf{lra}(\mathbf{v}))\, dp(\mathbf{v}) > \int \varphi(\mathbf{lra}(\mathbf{v}))\, dq(\mathbf{v})$.*

This will play a large role in our analysis of optimization problems with ergodic solutions. In strongly ergodic models, it is a probability 1 event under the optimal policy that all paths yield the same long run average. In such models, Corollary B.1 implies that one can simultaneously maximize the expected value of $S(\cdot)$ for *every* social welfare function satisfying Postulates I-V by maximizing the long run average. By contrast, in models with irreversibilities, models where optimal policies give a distribution over the long run average, the concavity of $\varphi(\cdot)$ determines how the social welfare function trades off different risky long run prospects.

3. The Conditional Equal Treatment Property

We now turn to the existence of patient optima satisfying a conditional equal treatment property in models where intergenerational externalities are transmitted by a state variable having a future distribution that depends on the present state and present actions. The conditional equal treatment property is defined by any state x being possible, and that action chosen at that state, $a(x)$, is determined by maximizing a social welfare function that weights the generations equally. This provides an ethically sound solution to the *underselectiveness problem* of the long-run average.[6]

There are two reasons one should expect there to be many optima when maximizing patient social welfare functions in economic models. The first reason is geometric: there can be no strictly concave functions on $\mathbf{W}$ because it is non-separable; maximizing a concave function with flat spots over a convex set determined by inequalities defined by concave functions will therefore often have multiple optima. The second reason concerns the particular form of the flat spots in dynamic models, invariance of the tangents to $o(T)$-permutations: changing the utility of the early generations up or down need have no effect on the long run; patient preferences only pay attention to the long run.

The average reward optimality equations (AROEs) and their conditional equal treatment implications provide a solution to the underselectiveness problem. We first develop the AROEs in a well-known, single sector growth model by perturbing the patient utility functions in a way that becomes more and more egalitarian as we go to the patient limit. We then give more general versions of the AROEs and an alternative, stochastic perturbation interpretation, one that captures the conditional

[6] An alternate approach to underselectiveness is to change the optimality criteria for infinite streams so as to increase the selectiveness/shrink the solution set. The most systematic treatment is in Hernández-Lerma and Vega-Amaya [23], who study seven variants of patient preferences and their relations within a broad class of strongly ergodic MDPs.

equal treatment property from a direction reminiscent of perfection in equilibrium refinement.

3.1. **Patience in Growth Models.** The initial period's endowment is $x_0 > 0$, initial period consumption of $c_0 \leqslant x_0$ yields utility $u(c_0)$, $u(\cdot)$ concave and increasing, and leaves $s_0 = (x_0 - c_0)$ to be invested for next period. The next period's endowment is $x_1 = f(s_0)$, $f(\cdot)$ concave, above the diagonal, i.e. $f(s) > s$, on an interval $(0, \overline{s})$, and eventually below the diagonal, i.e. $f(s) \leqslant s$ on $[\overline{s}, \infty)$. At x_1, the process begins again, consumption is $c_1 \leqslant x_1$, investment is $s_1 = x_1 - c_1$. The next period's endowment is $x_2 = f(s_1)$, and so on.

We analyze the optima for this model with three social welfare functions: the normalized discounted sum $L_\beta(\mathbf{u}) := (1-\beta)\sum_{t=0}^{T} u_t \beta^t$, $\beta < 1$; the $\mathbb{V}$-concave functional $S_{p\mathbb{V}}(\mathbf{u}) := \liminf_{T\to\infty} \frac{1}{T+1}\sum_{t=0}^{T} u_t$; and an arbitrary $\mathbb{V}$-concave function patient $\mathbf{u} \mapsto S(\mathbf{u})$. The last two have identical, and large, sets of optimal strategies. However, as $\beta \uparrow 1$, the Bellman equation for the $L_\beta(\cdot)$ has a limiting form that treats each x as possible and treats the generations equally. The optimal policy for this limit Bellman equation is an optimal policy for the second and third welfare functions, and this policy has the conditional equal treatment property.

There are two intuitions for this result. First, for $\beta < 1$, the value functions for $L_\beta(\cdot)$ treat each initial state x as possible, and as $\beta \uparrow 1$, the relative weights given to generations T and T', i.e. $(\frac{\beta^T}{\beta^T+\beta^{T'}}, \frac{\beta^{T'}}{\beta^T+\beta^{T'}})$, converge to $(\frac{1}{2}, \frac{1}{2})$ for each T, T' pair. Second, $\mathbf{u} \in \mathbf{Erg}$ if and only if $L_\beta(\mathbf{u}) \simeq \liminf \frac{1}{T+1}\sum_{t=0}^{T} u_t$ for all $\beta \simeq 1$, with a similar result for the tangents to any $\mathbb{V}$-concave $S(\cdot)$.

3.1.1. *Optimal Paths with Discounting*. For $0 < \beta < 1$, the normalized discounted value for starting at x_0 is

$$V(\beta, x_0) = \max_{c_0, c_1, \ldots} (1-\beta)\sum_{t=0}^{\infty} u(c_t)\beta^t \text{ s.t. } c_t \in [0, x_t],\ x_{t+1} = f(x_t - c_t). \quad (7)$$

The value function $V(\cdot,\cdot)$ is the unique solution to the functional Bellman equation

$$V(\beta, x) = \max_{c\in[0,x]} \left[(1-\beta)u(c) + \beta V(\beta, f(x-c))\right]. \quad (8)$$

Let c_t^* be an optimal path starting from x_0 with the associated investment levels, $s_t^* = x_t^* - c_t^*$, and production levels, $x_t^* = f(s_{t-1}^*)$. The first order conditions from the Bellman equation yield the following recursive relation that is satisfied along the optimal path,

$$\frac{u'(c_{t+1}^*)}{u'(c_t^*)} = \frac{1}{\beta f'(s_t^*)}. \quad (9)$$

To analyze or of optimal discounted paths, let $s^*(\beta)$ solve $\max_{s\geqslant 0}[\beta f(s)-s]$, define $x^*(\beta) = f(s^*(\beta))$ and $c^*(\beta) = f(s^*(\beta)) - s^*(\beta) = x^*(\beta) - s^*(\beta)$. Note that $c^*(\cdot)$, $s^*(\cdot)$ and $x^*(\cdot)$ are all increasing in β, $0 < \beta \leqslant 1$. The stationary solution to the recursive relation in (9) has $c^*_{t+1} \equiv c^*_t$, which yields $f'(s) = 1/\beta$, i.e. optimal investment is $s^*(\beta)$, optimal consumption is $c^*(\beta)$ and each period begins with stock $x^*(\beta)$. The nonstationary solutions can be divided in two cases: if $x_0 < x^*(\beta)$, then $c^*_t \uparrow c^*(\beta)$, $x^*_t \uparrow x^*(\beta)$, and $s^*_t \uparrow s^*(\beta)$; if $x_0 > x^*(\beta)$, then $c^*_t \downarrow c^*(\beta)$, $x^*_t \downarrow x^*(\beta)$, and $s^*_t \downarrow s^*(\beta)$.

3.1.2. *Underselectiveness for Patient Preferences.* With $\mathbf{u} = (u(c_0), u(c_1), u(c_2), \ldots)$, there are many optimal paths for the utility function $S_{p\mathbb{V}}(\cdot)$ or $S(\cdot)$, a phenomenom called **underselectiveness**. One class of solutions involve extreme forms of "present profligacy," the generations $t = 0, 1, \ldots, T$ feast on the capital stock, operating only under the constraint that $x_{T+1} = \epsilon > 0$, and then society consumes along a path that has $\lim_t x_{T+t} = x^*(1)$. Following Chichilnisky [15], another class of solutions involve a "dictatorship of the (far) future:" set $c_t = 0$ for $t = 0, 1, \ldots, T$ until some large T with $x_{T+1} > x^*(1)$, and then start consuming along some path with $c_{T+t} \downarrow c^*(1)$.

The **average reward optimality equations (AROEs)**, in the form of a limiting version of the Bellman equation (8) as the discount factor approaches 1 provide a conditional equal treatment method to avoid such ethically objectionable optima.[7]

3.1.3. *AROEs.* Let ρ^* denote $u(c^*(1))$, the so-called Golden Rule level of utility. As $\beta \uparrow 1$, $V(\beta, x) \to \rho^*$ for all x. Letting $V_1(\cdot,\cdot)$ be the partial derivative of V with respect to its first argument, a first order Taylor expansion tells us that for β infinitesimally close to 1, $V(\beta, x)$ is infinitesimally close to $\rho^* - (1-\beta)V_1(\beta, x)_{|\beta=1}$. Letting $V_1(1, x)$ denote $V_1(\beta, x)_{|\beta=1}$ in the the Bellman equation and rearranging yields

$$V(\beta, x) \simeq \max_{c\in[0,x]}\left[(1-\beta)u(c) + \beta\left(\rho^* + (1-\beta)\cdot(-V_1(1,x))\right)\right], \text{ or} \quad (10)$$

$$V(\beta, x) - \beta\rho^* \simeq (1-\beta)\max_{c\in[0,x]}\left[u(c) + \beta\cdot(-V_1(1,x))\right]. \quad (11)$$

The left-hand term, $V(\beta, x) - \beta\rho^* = V(\beta, x) - \beta V(1, x)$, can be expanded,

$$V(\beta, x) - V(1, x) + V(1, x) - \beta V(1, x) \simeq (1-\beta)\cdot(-V_1(1,x)) + (1-\beta)V(1,x). \quad (12)$$

[7]For general treatments of this vanishing discount factor approach to the AROEs, see Feinberg et al. [19] and the simplifications of their arguments in Vega-Amaya [45].

Substitute (12) into (11), divide both sides by $(1-\beta)$, then take $\beta \simeq 1$. This yields what is called the average reward optimality equation,

$$\rho^* + (-V(1,x)) = \max_{c\in[0,x]} \left[u(c) + (-V_1(1, f(x-c)))\right], \text{ or} \tag{13}$$

$$\rho^* + h^*(x) = \max_{c\in[0,x]} \left[u(c) + h^*(f(x-c))\right] \tag{14}$$

with $h^*(x) = (-V_1(1,x))$. In this simple model, it can be directly verified that if ρ^* and $h^*(\cdot)$ satisfy (14) and $c^*(x)$ is an argmax, then ρ^* is the maximal long-run average payoff and the argmax policy $c^*(\cdot)$ delivers a long-run average of ρ^*.

3.1.4. *Conditional Equal Treatment.* The function $h(\cdot) = -V_1(1,\cdot)$ gives society's limit sensitivity to discounting as a function of the initial capital stock, the limit being taken as discounting disappears. It is strictly increasing in x and satisfies $h(x^*(1)) = 0$, and we will see why this should be true from an alternate perspective in the more general treatment of the AROEs below. In a direct parallel with the Bellman equations, if a generation starts at $x_0 \neq x^*(1)$ and follows the argmax $c^*(\cdot)$ rule, then the path will satisfy

$$\frac{u'(c^*_{t+1})}{u'(c^*_t)} = \frac{1}{f'(s^*_t)} \text{ rather than } \frac{u'(c^*_{t+1})}{u'(c^*_t)} = \frac{1}{\beta f'(s^*_t)}, \tag{15}$$

and capital will move slowly to the Golden Rule level $x^*(1)$, increasing if $x_0 < x^*(1)$ and decreasing if $x_0 > x^*(1)$.

In moving to the limit along a set of utility functions with increasingly equal weight given to each generation, we have arrived at optimal consumption and investment paths that arise from giving exactly equal weight to each generation. One sees this by comparison of the laws of motion for the discounted case, (9), and the limit case (15). In the discounted case, the $\beta < 1$ arises because the relative weights of generations t and $t+1$ in the social welfare function are $(\frac{1}{1+\beta}, \frac{\beta}{1+\beta})$, and $\frac{1}{1+\beta} > \frac{1}{2}$. In the AROEs, the relative weights are $(\frac{1}{2}, \frac{1}{2})$.

3.1.5. *AROEs and Optimality for all Patient Preferences.* We now show that solving the AROEs in this class of growth models simultaneously yields S-optimality for any $\mathbb{V}$-concave social welfare functionals $S(\cdot)$. There is a strong intuition for this result: by the concavity of $u(\cdot)$ in consumption, randomization in the savings rate can never be optimal; by concavity of $S(\cdot)$ on **W**, randomized non-stationary policies are dominated; for the same reason, among the stationary deterministic policies, any strategy that demonstrates long-run cyclical behavior must be dominated; among the remaining strategies, all outcomes belong to **Erg**, that is, all outcomes have a long-run average utility; finally, by Corollary B.1(1), applied to elements of **Erg**, all of the social welfare functions rank paths by their long-run average, and solving the AROEs maximizes the long-run average. The dependence of this argument on the concavity of $u(\cdot)$ is needlessly limiting.

The tangents to our social welfare functions are positive rescaling of elements of $p\mathbb{V}$, and this set is defined as the weak* accumulation points of the linear functionals $L_\tau(\mathbf{u}) = E\,\frac{1}{1+\tau}\sum_{t=0}^{\tau} u_t$ as $P(\tau < M) \to 0$ for all M. The extreme points of $p\mathbb{V}$ are the accumulation points of the L_τ when τ is a point mass on larger and larger T. Equivalently, the extreme points are accumulation points of the linear functionals $\mathbf{u} \mapsto \langle \mathbf{u}, \eta_T\rangle$ where η_T is the uniform distribution on $\{0, 1, \ldots, T\}$.

A continuous linear functional, $L(\cdot)$, on the set of bounded sequences of utilities is called a **Banach-Mazur** limit if $L(\mathbf{u}) \geqslant 0$ for every $\mathbf{u} \geqslant 0$, $L((1, 1, 1, \ldots)) = 1$ and $L(\mathbf{u}) = L(\mathbf{u}^\pi)$ for every bounded permutation π. Our class of tangents, the cone generated by $p\mathbb{V}$, is the strict subset of Banach-Mazur limits characterized by indifference to the wider class of $o(T)$-permutations. All Banach-Mazur limits can be represented as integrals, $L(\mathbf{u}) = \int u_t \, d\eta(t)$, where η is a non-atomic, purely finitely additive probability measure on $\mathbb{N}_0$ (e.g. Jerison [28] or Robinson [37].)

It is easy to show that a policy maximizes the long-run average if and only if it maximizes $\langle \mathbf{u}, \eta\rangle$ for all η's that are accumulation points of the uniform distributions η_T. If the same choice is optimal for every extreme point in the set of possible tangents, then it is optimal for every convex combination of the extreme points. As a result, one maximizes the long-run average if and only if one maximizes $S(\cdot)$ for every $\mathbb{V}$-concave utility function.

3.2. **General Markovian Decision Problems.** In discrete time Markovian decision problems, at each point in time, a state is observed, a feasible action is chosen, the state and chosen action deliver an in-period reward and determine the distribution of the next period's state. The growth model just covered is a Markovian decision problem with deterministic rather than probabilistic transitions. We assume throughout that in-period rewards are non-negative and uniformly bounded above.

Notationally, the state is a point x in a space X, the chosen action is a point a belonging to the feasible set of actions denoted $A(x)$, $A(x)$ is a subset of a larger set A, and $K = \{(x, a) : a \in A(x)\}$ denotes the graph of the action correspondence. The in-period utility is given by a *bounded* $u : K \to \mathbb{R}_+$, and the next period's distribution is given by a transition probability $(x, a) \mapsto Q(\cdot|x, a)$ from K to $\Delta(X)$, the set of distributions on X. At a minimum, one assumes that the sets X and A are Polish (or measurably isomorphic to a Borel measurable subset of a Polish space), that the correspondence $x \mapsto A(x)$ is non-empty valued, that K is measurable and allows measurable selections, and that both $(x, a) \mapsto u(x, a)$ and $(x, a) \mapsto Q(\cdot|x, a)$ are measurable.

Histories at $t \geqslant 0$ are of the form $h_t = (x_0, a_0, x_1, \cdots, x_{t-1}, a_{t-1}, x_t)$ with $a_k \in A(x_k)$ for $k = 0, 1, \ldots, t-1$. Randomized policies, $\pi = (\pi_t)_{t=0}^{\infty}$, are sequences of measurable functions $h_t \mapsto \pi_t(h_t) \in \Delta(A(x_t))$ where $\Delta(A(x_t))$ denotes the set of distributions on A that put mass 1 on $A(x_t)$. Deterministic policies are randomized policies assigning mass 1 to a single point $f_t(h_t) \in A(x_t)$. Stationary deterministic policies are the main focus, these are deterministic policies with the property that $f_t(h_t) = f(x_t)$ for some fixed measurable function $f : X \to A$ with $f(x) \in A(x)$.

For any stationary deterministic policy, the transition probability for the system is denoted $x \mapsto Q_f(\cdot|x)$ and defined by $Q_f(B|x) = Q(B|x, f(x))$, and for a stationary policy π, it is $Q_\pi(B|x) = \int Q(B|x, a)\, d\pi(a|x)$. Under any stationary policy, the system becomes a Markov chain and the n-step transition probabilities are denoted Q_f^n and Q_π^n. Starting from any initial state x_0, under the minimal assumptions mentioned above, a policy π gives rise to a unique distribution on $X \times (A \times X)^{\mathbb{N}}$. Expectations with respect to this distribution are denoted $E^\pi(\cdot|x_0)$ for policies π, stationary or not, and denoted $E^f(\cdot|x_0)$ for stationary deterministic policies f.

3.3. **Long-Run Average Optimality for MDPs.** Assumptions on (X, A, K, u, Q) sufficient to guarantee the existence of policies delivering maximal long-run average payoffs have been extensively studied (Arapostathis et al. [1] is a slightly dated survey, Jaśkiewicz and Nowak [27] and Feinberg et al. [19] are recent extensions of these results that contain short overviews of more recent work). For the payoffs, one typically invokes continuity or upper semi-continuity assumptions on each $u(x, \cdot)$ and complementary assumptions about each $A(x)$ that guarantee that each $u(x, \cdot)$ achieves a maximum on $A(x)$. One also invokes two classes of joint assumptions on the payoffs and the stochastic structure: one class of assumptions guarantee that, if there are actions that could drive the distribution of the states "off to ∞," then these actions are not utility maximizing; the second class of assumptions guarantee that when taking the actions with higher utilities, the process "stays finite" in a Markovian fashion having a unique ergodic distribution.[8] When each stationary policy gives to a unique ergodic distribution, the so-called Poisson equation holds, and maximization along sojourn versions of this equation yields the AROEs.

[8]See Meyn and Tweedie's [34] comprehensive study of the stochastic stability of Markov chains.

3.3.1. *The Poisson Equation.* Fix a stationary deterministic policy f and suppose that $Q_f(\cdot|x)$ gives rise to a positive Harris chain (a ψ-recurrent chain with a unique stationary distribution [34, §10.1.1, p. 231]) with ρ_f being associated long-run average payoff, $\rho_f = \lim_{T\to\infty} E^f(\frac{1}{T+1}\sum_{t=0}^{T} u(x_t, f(x_t))|x_0)$, and that ρ_f is independent of x_0.[9] The Poisson equation is a functional equation for an $h_f : X \to \mathbb{R}$,

$$(\forall x \in X)[\rho_f + h_f(x) = u(x, f(x)) + \int h_f(y)\, dQ_f(y|x)]. \tag{16}$$

Note that solutions to the Poisson equation are only identified up to the addition of a constant — if $h_f(x)$ is a solution, then so is $h_f(x) + 7$. The solution to the Poisson equation that is most informative for optimization theory is the deviation of payoffs from ρ_f along sojourns,

$$h_f(x_0) = E^f\left(\sum_{t=0}^{\tau_\alpha - 1} [u(x_t, f(x_t)) - \rho_f]\ |x_0\right) \tag{17}$$

where α is a recurrent atom for the chain $(x_0, x_1, x_2, \ldots)$ that starts in x_0 and has transition probabilities $Q_f(\cdot|\cdot)$ and $\tau_\alpha = \min\{t \geqslant 1 : x_t \in \alpha\}$ is the next hitting time for the set α.

The positive Harris chains that arise from the MDPs under study can always be "split" so as to introduce a recurrent **atom** (see [34, Ch. 5] for this construction). Keeping the same notation, $Q_f(\cdot|\cdot)$, for the split chain's transition probability, the atom is a measurable $\alpha \subset X$ with the property that for all $x, x' \in \alpha$, $Q_f(\cdot|x) = Q_f(\cdot|x')$. Letting τ_α be the random time until the next visit to α, the atom is recurrent — for all $x_0 \in X$, the expected time until the next visit to α, $E^f(\tau_\alpha|x_0)$, is finite.

A sojourn from α is a path that starts at an $x \in \alpha$ and spends $\tau_\alpha - 1$ periods away from α. For example, a length 2 sojourn from α is the last two elements of a path (x, y, x'), $x, x' \in \alpha$, $y \notin \alpha$. The distribution of the length 2 sojourns are given by $Q_f((x_1 \in B_1) \cap (x_2 \in B_2)|x_0)$ where $x_0 \in \alpha$, $B_1 \cap \alpha = \emptyset$ and $B_2 \subset \alpha$. Because $Q_f(\cdot|x) = Q_f(\cdot|x')$ for all $x, x' \in \alpha$, a chain with a recurrent atom can be analyzed as a path from the initial state to the atom α followed by a sequence of iid sojourns from α.

[9]See Meyn [33, §4, Theorems 4.1 and 4.2] for weak conditions guaranteeing that this probability 1 ergodicity result holds for all starting points x_0 with general state and action spaces. Here we assume that the utility function $u : K \to \mathbb{R}$ is non-negative and bounded, which makes some of the issues in the existence theory easier.

3.3.2. *Maximization Along Sojourns and the AROEs.* One of the most fruitful approaches to maximizing long-run average payoffs has been to look for conditions guaranteeing the existence of solutions to the AROEs,

$$\rho^* + h^*(x) = \max_{a \in A(x)} \left[u(x, a) + \int h^*(y)\, dQ(y|x, a) \right]. \tag{18}$$

This is because, if we have a solution to this equation and $a^*(x)$ is an argmax, then, under one additional condition,[10] then $a^*(\cdot)$ is a deterministic policy delivering ρ^* and ρ^* is the highest possible long-run average payoff starting from any $x \in X$.

In terms of the solution to the Poisson equation given in (17), the optimal policy specified by the AROEs treat every initial state as being part of a possible sojourn going to the atom α and then maximizes the (equally weighted) sum of generational utilities along the path back to α. The first aspect of this, every state being a possible state for any generation to face, is much like Simon and Stinchcombe's [39] perfection for infinite games. The second aspect of this, equal weighting of generations, captures the egalitarian aspect of patient preferences.

Returning to the deterministic growth model MDP above, $h(x_0) < 0$ for x_0 below the golden rule $x^*(1)$ measures the generations' shortfall from ρ^* until they reach the capital level $x^*(1)$; conversely, starting at $x_0 > x^*(1)$, $h(x_0) > 0$ measures the generations' surplus until they reach $x^*(1)$.

We have already indicated the reasoning behind the following result, and we record it here more formally.

Corollary B.2. *Suppose that for an MDP* (X, A, K, u, Q), (ρ^*, h^*, a^*) *satisfy, for all* $x \in X$,

$$\rho^* + h^*(x) = \max_{a \in A(x)} \left[u(x, a) + \int h^*(y) Q(dy|x, a) \right], \tag{19}$$

that $a^*(x)$ *is an argmax for this problem, and that for any stationary policy* π *and any* $x \in X$, $\frac{1}{n} \int h^*(y)\, dQ^n_\pi(y|x) \to 0$. *Then* $a^*(\cdot)$ *is a deterministic policy that solves* $\max_\pi E^\pi(S(\mathbf{u})|x_0)$ *for any* $x_0 \in X$ *and any* $\mathbb{V}$*-concave social welfare function* $S(\cdot)$.

4. Four Applications

Unlike the growth model analyzed in the previous section, our first application is a dynamic economy model with no externalities. Our second application has externalities of the kind discussed in the previous section, every policy yields a

[10] The condition is that for any stationary policy π and any $x \in X$, $\frac{1}{n} \int h(y)\, dQ^n_\pi(y|x) \to 0$. As part of his examination of the policy iteration approach to solving the AROEs, Meyn gives this result [33, Theorem 2.1] as well as several textbook references for the proof. For an alternative, contraction mapping approach to solving the AROEs and arriving at optimal policies under compactness-continuity and stochastic Lyapunov assumptions, see Vega-Amaya [44].

unique, in this case non-trivial, ergodic distribution on the state space. The difference between this and the growth model is that here the externalities are stochastic. Our last two examples have externalities of the more extreme and long-lasting sort, those that capture irreversible decisions.

The first irreversible example involves the possibility of species extinction, a possibility that can be avoided at a cost. In this case, all policies that are optimal for a discounted factor $\beta < 1$ *minimize* the payoff for patient preferences. The second irreversible example explores the implications of our preferences in a learning model and arrives at a variant of the precautionary principle.

4.1. **Overview and Comparison of the Applications.** Our first application treats patient preferences in general equilibrium exchange models with sequence commodity spaces. We show that with our preferences, competitive equilibria exist and the First and Second Welfare Theorems hold. The contrast between this result and the results on the need for "myopic" preferences for the existence of Pareto optima clarifies some aspects of the structure of patient preferences.

Because general equilibrium theory has a difficult time with our major focus, externalities, we turn to three models in which externalities operate through changes in a state variable. In the first of these models, the state cycles stochastically across the set of possible states with a long-run distribution determined by actions. The model is solvable using the average reward optimality equations and Corollary B.2 applies directly. We also contrast these results with the results of applying other classes of patient preferences.

Irreversible decisions are the *ne plus ultra* of externalities in intergenerational dynamic problems. Decision processes with irreversibilities do not fit into the class for which average reward optimality equations can be applied because they are not strongly ergodic — the long-run value depends on the initial state and/or the early realizations of the stochastics. In our first irreversible model, we show that for $\mathbb{V}$-concave preferences, optimal policies are very cautious when the irreversible event is unambiguously bad, e.g. extinction of a valuable species. The disctinction between the long-run behavior of the system under discount optimal policies and under patient optimal policies is as sharp as possible: all discount optimal policies deliver minimal payoffs for all patient preferences.

In our second model, the irreversible decision has long-run benefits that may or may not outweigh the long-run costs. Here the optimal decisions depends on the risk aversion encoded in the curvature of the social welfare function, and they embody a version of the precautionary principle. More specifically, we examine the possibility that, at a cost of both resources and delay, it is possible to learn more about the distribution of a once and for all shock to payoffs that will come from, say, the adoption of a new technology. The utility value of the shock may be positive, it may be negative, and this makes information valuable even when it is not definitive. The precautionary principle that results from maximizing patient preferences calls for learning until further learning will have no value, then adopting the technology if it still looks like a good idea, and abandoning it otherwise. Because the initial state and initial stochastic realizations can matter to the long-run payoffs, the curvature of $S(\cdot)$ enters into the expected utility calculations as given in Corollary B.1(2).

4.2. **Patience and Myopia in General Equilibrium Theory.** Bewley [9] studies general equilibrium models with commodity spaces that are uniformly bounded sequences of non-negative consumption vectors in $\mathbb{R}^k$. Bewley's Theorem 1 gives sufficient conditions for the existence of a competitive equilibrium with prices that positive, finitely additive $\mathbb{R}^k$-valued measures, while his Theorems 2 and 3 study conditions for the existence of equilibria with prices in the set of summable $\mathbb{R}^k$-valued sequences. Brown and Lewis [13] and Araujo [2] study the same class of general equilibrium models and give results, respectively, on the role of/need for myopia in preferences for the existence of Pareto optima.

4.2.1. *Notation and Assumptions.* Let ℓ_∞^k denote the set of $x \in (\mathbb{R}^k)^{\mathbb{N}_0}$ with $\sup_{t\in\mathbb{N}_0} \|x_t\| < \infty$. Feasible consumptions and endowments belong to $\mathbf{W}^k := \{x \in \ell_\infty^k : x \geqslant 0\}$. Each agent/dynasty i in a finite set I, has an endowment $\omega_i \in \mathbf{W}^k$. We give $\mathbf{W}^k$ the sup norm.

Assumption A. *Each ω_i is an interior point of $\mathbf{W}^k$.*

This implies that $\omega := \sum_i \omega_i$ is also an interior point of $\mathbf{W}^k$.

Feasible consumption streams are vectors $(x_i)_{i\in I}$ with each $x_i \in \mathbf{W}^k$, and $\sum_i x_i \leqslant \omega$. The preferences of i are given by a utility function $x \mapsto U_i(x)$. We say that a $\mathbb{V}$-concave $S : \mathbf{W} \to \mathbb{R}$ is **$\mathbb{V}$-concave at the boundary** if $\mathbf{D}S(\mathbf{u}) \neq \emptyset$ for all $\mathbf{u} \in \mathbf{W}$ (not just all $\mathbf{u} \in \text{int}(\mathbf{W})$).

Assumption B. $U_i(x) = S_i(\mathbf{u}_i(x))$ *where:* $\mathbf{u}_i = (u_{i,0}(x_0), u_{i,1}(x_1), \ldots)$; $S_i(\cdot)$ *is* $\mathbb{V}$-*concave at the boundary; and the* $u_{i,t}$, $t \in \mathbb{N}_0$ *are uniformly bounded, continuous, concave, strictly increasing period utility functions on* $[0, \omega_t] \subset \mathbb{R}_+^k$ *with* $u_{i,t}(0) = 0$ *and* $\liminf_t \{\max u_{i,t}(\omega_{i,t})\} > 0$.

It is the aggregation of period utilities by a $\mathbb{V}$-concave utility function that makes the preferences patient. The assumption that the range of the utilities $u_{i,t}$ does not disappear guarantees that the agents are non-trivial parts of the economy.

Definition 4.1. *An* ***equilibrium*** *for an exchange economy model* $\mathcal{E} = (\boldsymbol{\omega}_i, U_i)_{i\in I}$ *is a feasible consumption stream,* $(\mathbf{x}_i)_{i\in I}$ *and a price* π *in the topological dual of* ℓ^k_∞ *such that for all* $i \in I$ *and all* $\mathbf{y} \in W^k$,

$$[U_i(\mathbf{y}) > U_i(\mathbf{x})] \Rightarrow [\langle \mathbf{y}, \pi\rangle > \langle \boldsymbol{\omega}, \pi\rangle]. \tag{20}$$

4.2.2. *Equilibrium Existence and the Welfare Theorems.* The following equilibrium existence result follows directly from Bewley [9, Theorem 1] and does not require $\mathbb{V}$-concavity at the boundary.

Proposition 1. *Under Assumptions A and B, an equilibrium exists.*

For finite dimensional production and exchange economies without externalities, the First Welfare Theorem states that competitive equilibria are Pareto optimal, and the Second Welfare Theorem states that all Pareto optima are competitive equilibria after appropriate re-arrangement of the initial endowments. Using preferences with tangents that are integrals against convex combinations of countably and purely finitely additive measures Araujo [2, Theorem 3] shows that Pareto optimal allocations may not exist. Such preferences are not in the class of patient preferences we are using here, and this difference leads to very different results.

Because the tangents of $\mathbb{V}$-concave functions have representations as integrals against purely finitely measures, the price vector in Proposition 1 must be purely finitely additive. Our next result shows that, for such preferences, we have the First and Second Welfare Theorems. To complete the analysis, we present two examples. The first, due to Araujo [2], shows that for preferences with tangents that are a mix of countably and purely finitely additive measures, ϵ-Pareto optimal equilibria may not exist. The second example draws a parallel with optimal allocations of uncertainty: in the allocation of risk, if (say) agents i and j assign probability 1 and probability 0 to an event E, then the optimal allocations give i all of the consumption in the event E and give agent j all of the consumption in the event E^c; if dynasties i and j put mass 1 and mass 0 on the coalition E, we have the same pattern.

Proposition 2. *Under Assumptions A and B, every equilibrium is Pareto optimal, and every Pareto optimal allocation is an equilibrium for an appropriate re-arrangement of the initial endowments.*

The tangents to $\mathbb{V}$-concave preferences have representations as integrals against purely finitely measures. This is crucial to the existence of Pareto optimal points as the following example demonstrates.

Example 4.1 (Araujo)**.** *For* $I = \{1, 2\}$*, let* $\boldsymbol{\omega}_1 = \boldsymbol{\omega}_2 = \mathbf{1}$ *be the constant endowment of one unit of the single good. For an allocation* $\mathbf{x}$*, let* $U_1(\mathbf{x}) = x_1 + \langle \mathbf{x}, \eta \rangle$ *where* η *is a non-negative, purely finitely additive measure (not a probability) that satisfies* $\langle \boldsymbol{\omega}_1, \eta \rangle > 1$*. Let* $U_2(\mathbf{y}) = \langle \mathbf{y}, \gamma \rangle + \langle \mathbf{y}, \eta \rangle$ *where* γ *is countably additive and strictly positive, say* $\gamma_t = (1-\beta)\beta^t$ *so that* $\langle \mathbf{y}, \gamma \rangle = (1-\beta) \sum_{t=0}^{\infty} y_t \beta^t$*.*

Suppose now that $(\mathbf{x}, \mathbf{y})$ *is an individually rational Pareto optimal allocation. Because* $\gamma_t > 0$ *for all* t*, Pareto optimality implies that* $x_{1,t} = 0$ *and* $y_{2,t} = 2$ *for all* $t \geqslant 2$*. Feasibility implies that* $x_1 \leqslant 2$*. However,* $U_1(\boldsymbol{\omega}_1) = 1 + \langle \boldsymbol{\omega}_1, \eta \rangle > 2$ *which contradicts* $U_1(\mathbf{x}) \leqslant 2 + 0$*.*

Continuous linear preferences on $\mathbf{W}^k$ can be decomposed into a countably additive part and a purely finitely additive part. For any $\epsilon > 0$, the value of the countably additive parts are determined on $\{0, 1, \ldots, T\}$ for sufficiently large T, while the value of the purely finitely additive part is entirely determined on $\{T+1, T+2, \ldots\}$. For linear preferences, this implies that ϵ-individually rational and Pareto optimal allocations exist. We conjecture that the same is true for concave utility functions.

The next example demonstrates how the Pareto optimal allocations for dynasties with different $\mathbb{V}$-concave can be extreme. The essential intuition is the same as that of optimal risk-sharing between two people, one of whom assigns mass 1 to an event E while the other assigns mass 1 to it.

Example 4.2. *Let the agents and their endowments be as in the previous example. Suppose that* $U_i(\mathbf{y}) = \langle \mathbf{y}, \gamma_i \rangle$ *where* (γ_1, γ_2) *is an accumulation point of the set*

$$\{(\mathtt{Unif}_{0,T!}, \mathtt{Unif}_{0,(T!)^2}) : T \in \mathbb{N}_0\}.$$

It is possible to demonstrate a set of generations, E*, such that* $\langle 1_E, \gamma_1 \rangle = 1$ *and* $\langle 1_E, \gamma_2 \rangle = 0$*. All Pareto optimal allocations must assign dynasty* 1 *the entire economy's endowment in* E *and must assign dynasty* 2 *the entire economy's endowment in* E^c*.*

4.3. **A Stark Model of Climate Change.** Suppose that the world's ecosystem can be in one of two states, damaged or undamaged: in the damaged state, the seas, forests and the biota that survive are unable to produce oxygen and resources in the amounts humans have become accustomed to: in the undamaged state, the seas and forests are able to produce oxygen concentrations and resources supporting life as we currently know it. Payoffs and actions capture the following tradeoffs: a generation in a good state can sacrifice some present utility in order to lower the future probability of disastrous climate changes; a generation in a bad state must sacrifice some of their present utility in order to raise the the future probability of a return to a better world.

- In the undamaged state, $x = G$, society chooses the transition probability, $r \in [\underline{r}, \overline{r}]$ to the damaged state, $x = B$, $0 < \underline{r}$ and $\overline{r} < 1$. The expected utility of choosing r is $u(G, r)$, and higher choices of r lead to a higher expected utility for a generation in the good state, $\partial u(G, r)/\partial r > 0$.
- In a parallel fashion, in the damaged state, $x = B$, society chooses the transition probability, $s \in [\underline{s}, \overline{s}]$ to the good state, $x = G$, $0 < \underline{s}$ and $\overline{s} < 1$. The expected utility of choosing s is $u(B, s)$, and higher choices of s lead to lower expected utility for a generation in the bad state, $\partial u(B, s)/\partial s < 0$.

4.3.1. *Optimal Patient Policies.* Starting from the present, $t = 0$, a **policy**, f, chooses an r and a s as a function of the present state. This choice gives rise to a Markov process, $x^f = (x^f_t)_{t \in \mathbb{N}_0}$, taking either the value G or B and the associated stochastic stream of utilities $\mathbf{u} = (u(x_0, f(x_0)), u(x_1, f(x_1)), u(x_2, f(x_2)), \ldots)$. A policy f is **S-optimal** if it maximizes $E^f(S(\mathbf{u})|x_0)$ for $x_0 = G$ and $x_0 = B$.

Because the probabilities r and s are interior, any policy $f = (r, s)$ leads to the process x^f having a well-defined long-run average, $\rho = \rho_f := E \lim_{T \to \infty} \frac{1}{T+1} u(x_t, f(x_t))$ that is independent of the starting point. The AROEs are

$$\rho + h(G) = \max_{r \in [\underline{r}, \overline{r}]} \left[u(G, r) + E^f(h(x_1)|x_0 = G) \right], \text{ and} \tag{21}$$

$$\rho + h(B) = \max_{s \in [\underline{s}, \overline{s}]} \left[u(B, s) + E^f(h(x_1)|x_0 = B) \right]. \tag{22}$$

As discussed above, the difference between these equations and the discounted Bellman equations is the equal weights given to generations along sojourns.

In this model, $E^f(h(x_1)|x_0 = G) = (1 - r)h(G) + rh(B)$ and $E^f(h(x_1)|x_0 = B) = sh(G) + (1 - s)h(B)$. The first order equations (FOCs) for an interior solution to (21) and (22) are

$$\partial u(G, r)/\partial r = [h(G) - h(B)] \text{ and } \partial u(B, s)/\partial s = [h(B) - h(G)]. \tag{23}$$

Using the sojourn-based solution from (17), $h(G) = (u_G(r) - \rho) \cdot E\,\tau_B$, $E\,\tau_B = (1 - r)/r$, $h(B) = (u_B(s) - \rho) \cdot E\,\tau_G$, $E\,\tau_G = (1 - s)/s$ because the waiting times till transitions, τ_B and τ_G, are geometric distributions.

We assume that the payoffs in the good state are higher than those in the bad state, which leads to $u_G(r) - \rho > 0 > u_B(s) - \rho$. When the expected values of the times until transitions between states are large, we expect $[h(G)-h(B)]$ to be a large positive number. From this, one expects the right-hand sides of the FOCs to be too large (in absolute value) for interior solutions. This would imply that the optimal policy is as careful as possible in the good state and works as hard as possible to return to the good states when in the bad state, that is, $f^* = (\underline{r}, \overline{s})$.

4.3.2. *Optimal Discounted Policies.* In this class of models, the strongly ergodic ones, policies that are optimal for the patient preferences are nearly optimal for the discounted preferences when the discount factor, β, is close to 1. Specifically, as $\beta \uparrow 1$, the optimal discounted policies converge to the optimal patient policies, the normalized discounted values, $V(\beta, G)$ and $V(\beta, B)$ converge to ρ^*, and the long-run distribution associated with using the optimal discounted policies converges to the long-run patient distribution.

We can see why this is true by examining the Bellman equations for the normalized discounted values,

$$V(\beta, G) = \max_{r \in [\underline{r}, \overline{r}]} [(1-\beta)u(G, r) + \beta(rV(\beta, B) + (1-r)V(\beta, G)] \tag{24}$$

$$V(\beta, B) = \max_{s \in [\underline{s}, \overline{s}]} [(1-\beta)u(B, s) + \beta(sV(\beta, G) + (1-s)V(\beta, B)] . \tag{25}$$

The FOCs for interior r and s reduce to

$$\tfrac{\partial u(G,r)}{\partial r} = \tfrac{\beta}{(1-\beta)} [V(\beta, B) - V(\beta, G)] \text{ and } \tfrac{\partial u(B,s)}{\partial s} = \tfrac{\beta}{(1-\beta)} [V(\beta, G) - V(\beta, B)] . \tag{26}$$

These are similar to the FOCs in (23), and the similarity becomes more pronounced after dividing $[V(\beta, B) - V(\beta, G)]$ by $(1-\beta)$ and using the Taylor approximations. With these, the FOCs are approximated by

$$\tfrac{\partial u(G,r)}{\partial r} = \frac{\beta}{(1-\beta)}(1-\beta) [(-V_1(\beta, G)) - (-V_1(\beta, B))] \text{ and} \tag{27}$$

$$\tfrac{\partial u(B,s)}{\partial s} = \tfrac{\beta}{(1-\beta)}(1-\beta) [(-V_1(\beta, B)) - (-V_1(\beta, G))] , \tag{28}$$

which, as $\beta \uparrow 1$, reduce to

$$\frac{\partial u(G, r)}{\partial r} = \beta [h(G) - h(B)] \text{ and } \frac{\partial u(B, s)}{\partial s} = \beta [h(B) - h(G)] \tag{29}$$

as calculated above. A slightly finer-scaled analysis shows that as $\beta \uparrow 1$, the sacrifice of each generation for the good of future generations increases, reducing the optimal r in the good state and increasing the optimal s is the bad state.

4.3.3. *Other Intergenerationally Equitable Preferences.* In his examination of intergenerational equity issues in growth models with exhaustible resources, Solow [40] noted Rawls's unwillingness to extend maximin preferences from situations of uncertainty to intergenerational settings because of the fundamental asymmetry that time's arrow is one-directional — early generations actions affect later generations but not the reverse. Solow then proceeded to flesh out other objections. The most important of these was a 'poverty trap,' and the argument had three parts: if the initial stock of capital in a growth model is low, then the initial generation has the lowest well-being; if one maximizes the utility of the worst off, then the first generation cannot be called on to make sacrifices; inductively, no generation will save for future well-being and the economy will stay at the initial low level.

We have seen, in the growth model of §3.1, that optimal paths for the patient and equitable preferences $S_{\liminf}(\mathbf{u}) = \liminf_T \frac{1}{T+1}\sum_{t=0}^{T} u_t$ leads, slowly, to the Golden Rule, maximal possible long-run consumption. The $S_{p\mathbb{V}}(\cdot)$ preferences are invariant with respect to all $o(T)$-permutations, but there is a patient preference relation between the Rawlsian $R(\mathbf{u}) := \inf_t u_t$ and $S_{p\mathbb{V}}$, namely the un-averaged lim inf preferences, $L(\mathbf{u}) := \liminf_t u_t$. These preferences are invariant with respect to all permutations, arguably a more egalitarian approach, and the optimal paths in the growth model also converge to the Golden Rule consumption. However, the good behavior of the $L(\cdot)$ preferences can disappear in the presence of uncertainty.

Consider the policies in this climate model that myopically maximize utility in the damaged state. These are policies of the form $(r, \underline{s})$. Any such policy maximizes the expected value of the $L(\mathbf{u})$ because $L(\mathbf{u}) = u(B, s)$ with probability 1 for *any* policy. It is the failure of this ordering to respect improvements in the welfare of non-null coalitions that is at work here — with probability 1, along any path, there is a non-negligible portion of the generations in the good state, their utility does not enter in the $L(\cdot)$ ordering, and this precludes making tradeoffs between the welfare of different proportions of the generations that make up society.

4.4. **Patience and Extinction.** We analyze a simple model of a fishery where species extinction is possible, but avoidable at a cost. In this model, despite the short run behavior being very similar, there is a stark difference between the long-run, ergodic distributions implied by discount optimal policies and the optimal policies for patient preferences.

Suppose that there are two states, f and e, corresponding to the fishery being viable and the fish being extinct. We suppose further that the sets of available actions are $A(f) = [0,1]$, $A(e) = \{0\}$, that utilities are $u(e, 0) = 0$, that using a higher action in the viable state f is more profitable, $\partial u(f, a)/\partial a > 0$, but that higher actions make it more likely that the fish will become extinct, and extinction is absorbing $p_{e,e}(1) = 1$. Specifically, assume

- $u(f, a) = 10(1 + \sqrt{a})$,
- $p_{f,e}(a) = \begin{cases} 0 & \text{if } a \leqslant a^\circ \\ \frac{1}{2}(a - a^\circ)^2 & \text{if } a > a^\circ \end{cases}$

One can think of the bound a° as the minimal size of a marine reserve necessary to guarantee that the fish do not go extinct. While this model has a very simple, two state dynamic structure to the fish population, hence it applies more to shrimp than to tuna, the essential lessons will remain valid with a more complicated population dynamic.[11]

4.4.1. *The Patient Solution.* Any policy that repeatedly runs any uniformly positive risk of extinction has $\mathbf{lra}(\mathbf{u}) = 0$. Therefore, from Corollary B.1, for any $\mathbb{V}$-concave $S(\cdot)$, the S-optimal strategy has $a^P = a^\circ$ and $\mathbf{lra}(\mathbf{u}) = 10(1+\sqrt{a^\circ})$. Here, the optimal long-run distribution has fish and a positive value for all of the patient preferences.

4.4.2. *The Discounted Solution.* Define $q_{f,e}(a) = 1 - p_{f,e}(a)$. The Bellman equation for discount factor β has $V(\beta, e) \equiv 0$, and

$$V(\beta, f) = \max_{a \in [0,1]} 10(1 + \sqrt{a}) + \beta\left[q_{f,e}(a)V_\beta(1) + p_{f,e}(a)0\right].$$

The FOCs are

$$\frac{5}{\sqrt{a}} = \beta v \cdot (a - a^\circ)$$

where $v = V(\beta, f)$.

We now argue that

$$(\ddagger)\ (a^* - a^\circ) \propto \frac{(1-\beta)}{10(1+\sqrt{a^*})},$$

which means that as $\beta \uparrow 1$, the optimal action converges downwards to a°. The policy $a^*(\beta)$ decreases continuously to a° as $\beta \uparrow 1$. However, the long-run ergodic distribution for the policy $a^*(\beta)$, $\beta < 1$, always puts mass 1 on exinction, and it is here that we see the important distinction between discounted preferences and patient preferences.

> The discounted optimal policy for any $\beta < 1$ *minimizes* the payoffs for any patient preferences satisfying Postulates I-V.

[11] See Huang and Smith [25, §1] for the bioeconomic appropriateness of modeling shrimp as an annual industry with simpler dynamics.

One sees this simply by noting that putting mass 1 on long-run extinction means that the long-run average payoff is always 0.

To see why (‡) holds, note that if $a^* = a^*(\beta)$ is the solution, then

$$V(\beta, f) = u(a^*)/(1 - \beta \cdot q_{f,e}(a^*))$$

where $u(a^*) = 10(1 + \sqrt{a^*})$. Substituting into the FOCs yields

$$\beta \cdot 10(1 + \sqrt{a^*})(a^* - a^\circ) = 1 - \beta(1 - \frac{1}{2}(a^* - a^\circ)) = 1 - \beta + \frac{1}{2}\beta(a^* - a^\circ).$$

Dividing both sides by $(a^* - a^\circ)$ yields

$$\beta \cdot 10(1 + \sqrt{a^*}) = \left(\frac{1 - \beta}{a^* - a^\circ}\right) + \frac{1}{2}\beta(a^* - a^\circ).$$

As $\beta \uparrow 1$, the only way to arrange this to stay true is to have $(a^* - a^\circ) \propto \frac{(1-\beta)}{10(1+\sqrt{a^*})}$.

4.5. **Irreversibility and the Precautionary Principle.** Species loss/extinction represents an irreversible negative shock to the well-being of future generations. We now consider problems in which the irreversible decision will deliver a risky shock to well-being, one that may be either positive or negative. The question we analyze is the amount of research that should be done before making the decision. Again, to highlight the role of patience, we work in a simplified model.

We assume that the present stochastic path of utilities, $\mathbf{u}$, has a long-run average, $\mathbf{lra}(\mathbf{u})$. There is a hidden state X with $\mathtt{Prob}(X < 0) > 0$ and $\mathtt{Prob}(X > 0) > 0$. When the action $a = 1$ is taken, generational well-being will go up/down by X forever thereafter and no further actions are available. When the action $a = 0$ is taken, the possibility of taking the decision is closed off, and utility will be $\mathbf{u}$, no further actions are available. Until either $a = 0$ or $a = 1$ is chosen, the action s is available. When $a = s$ is chosen, a signal that is stochastically related to X will be observed. The action s corresponds to researching into the consequences of the irreversible decision. We assume that there is a random, unknown number of informative signals available, and when they are exhausted, this is observed and any further signals are stochastically independent of X. Further, every attempt to observe an informative signal costs c.

If M, the number of informative signals is known to be 0, then the decision for a society with preferences represented by $S(\cdot)$ is given by the comparison of $S(\mathbf{u})$ and $E\, S(\mathbf{u} + X1_{\mathbb{N}_0})$. By Corollary B.1, this reduces to the comparison

$$\varphi(\mathbf{lra}(\mathbf{u})) \lesseqgtr E\, \varphi(\mathbf{lra}(\mathbf{u}) + X). \tag{30}$$

Here the curvature of $S(\cdot)$ encoded in $\varphi(\cdot)$ determines the attitude toward risk, and one might expect more risk tolerance for higher values of $\mathbf{lra}(\mathbf{u})$.

More generally, if is the posterior distribution after beliefs have converged, e.g. after it is known that there are no more informative signals, then the comparison is $\varphi(\mathbf{lra}(\mathbf{u})) \lesseqgtr E\,\varphi(\mathbf{lra}(\mathbf{u})+X)$. Let B^- denote the set of beliefs for which $E\,\varphi(\mathbf{lra}(\mathbf{u})+X) \leqslant \varphi(\mathbf{lra}(\mathbf{u}))$ and B^+ the set for which $E\,\varphi(\mathbf{lra}(\mathbf{u})+X) > \varphi(\mathbf{lra}(\mathbf{u}))$. The optimal policy takes the following form.

> Research until posterior beliefs either converged to B^+ or to B^-. If they have converged to B^+, take the irreversible decision, otherwise close it off.

- Convergence to either the set B^+ or B^- entails knowing that future information will not move beliefs out of the set.
- A definite decision is made.
- The costs do not enter the analysis.

Thus, we have a precautionary principle, but one with some subtleties. It asks that irreversible decisions be delayed until uncertainty is reduced, not to nothing as in Sunstein's [41] 'straw man' version of the precautionary principle, but until society is as sure as possible that the expected benefits, $X > 0$, outweigh the expected costs, $X < 0$. Even if it is not definitive, information that can change the optimal action is worth waiting for in general, and patience magnifies that effect. One form that this convergence may take is that the research may reveal ways to mitigate negative consequences, lowering $\texttt{Prob}(X < 0)$.

To have the costs enter the analysis more sensibly, one should consider a version of this model in which each generation is faced with their own irreversible decision(s). In such an analysis, costs become a permanent component of generational well-being, and society optimally trades off between this component and the future additions/subtractions to the utility path.

5. Summary and Conclusions

This paper has studied the lengthening of the horizon for optimization, the equal treatment of generations inhabiting that longer horizon, and how these interact with dynamic externalities, up to and including externalities in their strongest form, irreversibilities. Viewing society as an aggregate of present and future generations, intergenerational equity captures societal patience. We have offered a resolution to the conflict between intergenerationally equitable preferences and Pareto responsiveness by specifying a class of social welfare functions that have both a tangent formulation and an axiomatic foundation. Further, we have used the ideas behind perfect equilibrium to show that there are conditional equal treatment solutions, and these provide an ethical resolution to the underselectiveness problems that have plagued previous attempts to solve long-run optimization problems.

Application to general equilibrium theory allows us to understand one set of properties of the tangents, the relative weights that can be given to different parts of the future by different dynasties. Applications to two general classes of Markovian decision problems afford further insights. A Markovian decision problem is *strongly ergodic* if every choice of a stationary policy π leads to a unique long-run distribution μ_π for which equal long-run average payoffs is a probability 1 event. By contrast, the problem is *pathwise ergodic* if, with probability 1, each path has a long-run average, but the average can be different on different paths. Pathwise ergodicity captures the possibility that early events and responses to them have influence the long run path of the system.[12]

5.1. **Tangents and Postulates.** Our main postulate on preferences is that if $\mathbf{u}$ strictly average overtakes $\mathbf{v}$, then $\mathbf{u} \succ \mathbf{v}$. For an $r > 0$ and B a set of generations, $\mathbf{u} := \mathbf{v} + r \cdot 1_B \geqslant \mathbf{v}$. When $\mathbf{0}$ average overtakes $r \cdot 1_B$, we not only have $\mathbf{u}$ average overtaking $\mathbf{v}$, we also have $\mathbf{v}$ average overtaking $\mathbf{u}$, in which case our preferences judge $\mathbf{v}+r\cdot 1_B$ and $\mathbf{v}$ as indifferent. It is only when $r\cdot 1_B$ strictly average overtakes $\mathbf{0}$ that our preferences judge $\mathbf{v}+r\cdot 1_B$ as strictly better than $\mathbf{v}$. The tangent formulation provides further understanding of and justification for this pattern of indifference and strict preference.

Our class of tangents is a subset of Banach-Mazur limits, that can be represented by integrals against non-atomic, purely finitely additive measure on $\mathbb{N}_0$. Hence our patient preferences evaluate utility allocations to the non-atomic measure space that models society by integrating utility allocations against the non-atomic, finitely additive measure η. Non-atomic measures come with rich classes of null sets, and

[12]Some decisions clearly set history on different paths. For example, in 1953, after Stalin's death, Dwight D. Eisenhower argued that the world found itself at "… one of those times in the affairs of nations when the gravest choices must be made, if there is to be a turning toward a just and lasting peace." He talked of the long-run consequences of present choices, "Every gun that is made, every warship launched, every rocket fired signifies, in the final sense, a theft from those who hunger and are not fed, those who are cold and are not clothed. This world in arms is not spending money alone. It is spending the sweat of its laborers, the genius of its scientists, the hopes of its children. … This, I repeat, is the best way of life to be found on the road the world has been taking. This is not a way of life at all, in any true sense. Under the cloud of threatening war, it is humanity hanging from a cross of iron."

the essential intuition behind our resolution is that we only pay attention to non-null subsets of society.

The cumulative distribution function for a random variable, τ, having a purely finitely additive distribution η must satisfy $\mathtt{Prob}(\tau \leqslant M) = F_\eta(M) = 0$ for every integer M. This indicates that the probability η on puts all of its mass "to the right of $\mathbb{N}_0$." Future generations are captured in the mass to the right and this is our idealization of the 10 million to 1 "ratio of people who will potentially live to the people living now." The purely finitely additive aspect of this idealization leads to what has been, in the past, identified as an incompatibility result for intergenerational equity and respect for the Pareto ordering.

A probability on $\mathbb{N}_0$ is purely finitely additive if and only if all $\boldsymbol{v}$ with $v_t \to 0$ integrate to 0. These integrals are the tangents to the preferences under study. If $\boldsymbol{v} \geqslant 0$ and $v_t \to 0$ with $v_t > 0$ for infinitely many t, then for any $\mathbf{u}$, we have $u_t + v_t > u_t$ for infinitely many generations. From the point of view suggested by our tangents, for any $\epsilon > 0$, no matter how small, the subset of society that is doing ϵ better in $\mathbf{u} + \boldsymbol{v}$ has mass 0. By contrast, the tangents require that $S(\mathbf{u} + \boldsymbol{v}) > S(\mathbf{u})$ when (and only when) a $\boldsymbol{v} \geqslant 0$ delivers a strictly positive amount of extra utility to a non-negligible subset of society.

5.2. **Strongly Ergodic Problems.** The classical growth model and the stark climate change model are examples of problems in which every policy leads to a single limit distribution. For this class of models our theoretical contribution is to show how broad a class of preferences are simultaneously covered by maximization of the long-run average. Our analysis depends on the tangent formulation: one maximizes the long-run average if and only if one maximizes each affine function in the class $\mathbb{V}$; if an action maximizes each tangent to a concave function, then it maximizes the concave function itself.

From the observation that a single choice can be optimal for so many utility functions, one can get a sense of how underselectiveness arises. Our solution to underselectiveness can be most clearly seen in the growth model: a limit form of the discounted Bellman equation provides the average reward optimality equations; these equations treat every state as possible; and on stochastic paths to the long-run states, treats the welfare of generations equally.

In the optimal policies for these examples, we see the tradeoffs between two aspects of the preferences we study. The concern for future generations pushes for investment in the growth model, while the concern for equal treatment of present generations pushes for equality of marginal utilities hence equality of consumption across generations, hence pushes for present consumption. In the climate change model, optimal policies cannot call for equalization of marginal utilities. Instead, they call for sacrifices by the well-off in the interests of making it less likely that future generations will be less well-off, and they also call for sacrificies by those who are less well-off to make it more likely that future generations are well-off.

5.3. **Pathwise Ergodicity, (Pre)Caution and Option Values.** When different long run paths have different long run utilities and these are stochastic, the concavity of $S(\cdot)$ induces risk aversion over the choice of paths. Risk aversion and equal treatment of generations does not lead to the "paralysis" as suggested in Sunstein's [41] straw man version of the precautionary principle. Rather, with the weight of future generations' well-being on the scale, optimal policies pursue progress while being very cautious about shutting off future options. The first effect, a preference for knowledge because of its capacity to expand options, comes from the magnification of those enjoying option value to improvements in the set of possibilities. The second effect comes through the similar magnification of the damages of irreversible losses.

As society puts more and more weight on the future, policies that avoid bad absorbing states and make possible better outcomes for future generations become more and more attractive, ceteris paribus. This is a precautionary approach that is conservative in the sense that it conserves a larger set of options for future generations. To put it another way, optimal policies pursue progress, but not at the cost of irreversibly losing what we have.

From the traditional viewpoint of optimality with discount factors that reduce the weight given to long-run welfare of future generations effectively to 0, the cautionary aspects of waiting for information can seem disconcerting. To us, this seems to be at the heart of the arguments about sustainability. Our emphasis on the importance of long horizon thinking leads a society to trade off the speed and scale of present developments for long-run flexibility.

Other versions of the precautionary principle prescribe allocating the expense of the research to those proposing the potentially irreversible action. Our analysis suggests that such a policy that may be optimal if potential benefits are privately appropriable while potential costs are public, but not otherwise.

REFERENCES

[1] Aristotle Arapostathis, Vivek S. Borkar, Emmanuel Fernández-Gaucherand, Mrinal K. Ghosh, and Steven I. Marcus. Discrete-time controlled Markov processes with average cost criterion: a survey. *SIAM J. Control Optim.*, 31(2):282–344, 1993.

[2] A. Araujo. Lack of Pareto optimal allocations in economies with infinitely many commodities: the need for impatience. *Econometrica*, 53(2):455–461, 1985.

[3] Geir B Asheim. Rawlsian intergenerational justice as a markov-perfect equilibrium in a resource technology. *The Review of Economic Studies*, 55(3):469–483, 1988.

[4] Geir B Asheim. Unjust intergenerational allocations. *Journal of economic theory*, 54(2):350–371, 1991.

[5] Geir B. Asheim. Intergenerational equity. *Annual Review of Economics*, 2(1):197–222, 2010.

[6] Anthony B. Atkinson. The strange disappearance of welfare economics. *Kyklos*, 54(2-3):193–206, 2001.

[7] Kaushik Basu and Tapan Mitra. Aggregating infinite utility streams with intergenerational equity: the impossibility of being Paretian. *Econometrica*, 71(5):1557–1563, 2003.

[8] Truman Bewley. An integration of equilibrium theory and turnpike theory. *J. Math. Econ.*, 10(2):233–267, 1 September 1982.

[9] Truman F Bewley. Existence of equilibria in economies with infinitely many commodities. *Journal of Economic Theory*, 4(3):514–540, 1972.

[10] Charles Blackorby, Walter Bossert, and David Donaldson. Intertemporal population ethics: critical-level utilitarian principles. *Econometrica*, 63(6):1303–1320, 1995.

[11] W. A. Brock and A. Haurie. On existence of overtaking optimal trajectories over an infinite time horizon. *Math. Oper. Res.*, 1(4):337–346, 1976.

[12] William A Brock. An axiomatic basis for the Ramsey-Weizsäcker overtaking criterion. *Econometrica*, 38(6):927–929, 1970.

[13] Donald J. Brown and Lucinda M. Lewis. Myopic economic agents. *Econometrica*, 49(2):359–368, 1981.

[14] Sukhamoy Chakravarty. The existence of an optimum savings program. *Econometrica*, 30(1):178–187, 1962.

[15] Graciela Chichilnisky. An axiomatic approach to sustainable development. *Social choice and welfare*, 13(2):231–257, 1996.

[16] E. V. Denardo and B. L. Miller. An optimality condition for discrete dynamic programming with no discounting. *Ann. Math. Statist.*, 39:1220–1227, 1968.

[17] Peter A. Diamond. The evaluation of infinite utility streams. *Econometrica*, 33(1):170–177, 1965.

[18] Nelson Dunford and Jacob T. Schwartz. *Linear operators. Part I*. Wiley Classics Library. John Wiley & Sons, Inc., New York, 1988. General theory, With the assistance of William G. Bade and Robert G. Bartle, Reprint of the 1958 original, A Wiley-Interscience Publication.

[19] Eugene A. Feinberg, Pavlo O. Kasyanov, and Nina V. Zadoianchuk. Average cost Markov decision processes with weakly continuous transition probabilities. *Math. Oper. Res.*, 37(4):591–607, 2012.

[20] Peter C. Fishburn. *The foundations of expected utility*, volume 31 of *Theory and Decision Library*. D. Reidel Publishing Co., Dordrecht, 1982.

[21] Marc Fleurbaey and Philippe Michel. Intertemporal equity and the extension of the Ramsey criterion. *Journal of Mathematical Economics*, 39(7):777–802, 2003.

[22] Pope Francis. *Encyclical letter Laudato Si' of the Holy Father Franci on care for our common home*. The Holy See: Vactican Press, 2015.

[23] O. Hernández-Lerma and O. Vega-Amaya. Infinite-horizon Markov control processes with undiscounted cost criteria: from average to overtaking optimality. *Appl. Math. (Warsaw)*, 25(2):153–178, 1998.

[24] Werner Hildenbrand. Pareto optimality for a measure space of economic agents. *International Economic Review*, 10(3):363–372, 1969.

[25] Ling Huang and Martin D Smith. The dynamic efficiency costs of common-pool resource exploitation. *The American Economic Review*, 104(12):4071–4103, 2014.

[26] Albert E. Hurd and Peter A. Loeb. *An introduction to nonstandard real analysis*, volume 118 of *Pure and Applied Mathematics*. Academic Press Inc., Orlando, FL, 1985.

[27] Anna Jaśkiewicz and Andrzej S. Nowak. On the optimality equation for average cost Markov control processes with Feller transition probabilities. *J. Math. Anal. Appl.*, 316(2):495–509, 2006.

[28] Meyer Jerison. The set of all generalized limits of bounded sequences. *Canad. J. Math*, 9(1):79–89, 1957.

[29] Gordon Keller and L. C. Moore, Jr. Invariant means on the group of integers. In *Analysis and geometry*, pages 1–18. Bibliographisches Inst., Mannheim, 1992.

[30] Tom Lindstrøm. An invitation to nonstandard analysis. In *Nonstandard analysis and its applications (Hull, 1986)*, volume 10 of *London Math. Soc. Stud. Texts*, pages 1–105. Cambridge Univ. Press, Cambridge, 1988.

[31] Massimo Marinacci. An axiomatic approach to complete patience and time invariance. *Journal of Economic Theory*, 83(1):105–144, 1998.

[32] Lionel W. McKenzie. Turnpike theory. *Econometrica*, 44(5):841–865, 1976.

[33] Sean P. Meyn. The policy iteration algorithm for average reward Markov decision processes with general state space. *Automatic Control, IEEE Transactions on*, 42(12):1663–1680, 1997.

[34] Sean P Meyn and Richard L Tweedie. *Markov chains and stochastic stability*. Springer Science & Business Media, 2012.

[35] Arthur C. Pigou. *The Economics of Welfare*. Macmillan, London, 1920.

[36] Frank Plumpton Ramsey. A mathematical theory of saving. *The Economic Journal*, 38:543–559, 1928.

[37] Abraham Robinson. On generalized limits and linear functionals. *Pacific J. Math*, 14:269–283, 1964.

[38] Henry Sidgwick. *The methods of ethics*. Macmillan & Co., London, 1874.

[39] Leo K. Simon and Maxwell B. Stinchcombe. Equilibrium refinement for infinite normal-form games. *Econometrica*, 63(6):1421–1443, 1995.

[40] Robert M Solow. Intergenerational equity and exhaustible resources. *The Review of Economic Studies*, 41:29–45, 1974.

[41] Cass R Sunstein. The paralyzing principle. *Regulation*, 25:32–37, 2002.

[42] Sèbastien Le Prestre de Vauban. Traité de la culture des forêts. In *Vauban, sa famille et ses écrits, ses oisivetés et sa correspondance: analyse et extraits*, volume 2. Berger-Levrault, Paris, 1910.

[43] Sèbastien Le Prestre de Vauban. Traité de la culture des forêts. In Hélène Vérin, editor, *Les Oisivetés de Monsieur de Vauban. Édition intégrale*. CDHTE-Cnam, SeaCDHTE, Seyssel, Champ Vallon, 2007.

[44] Oscar Vega-Amaya. The average cost optimality equation: a fixed point approach. *Bol. Soc. Mat. Mexicana (3)*, 9(1):185–195, 2003.

[45] Óscar Vega-Amaya. On the vanishing discount factor approach for Markov decision processes with weakly continuous transition probabilities. *J. Math. Anal. Appl.*, 426(2):978–985, 2015.

[46] Carl Christian Von Weizsäcker. Existence of optimal programs of accumulation for an infinite time horizon. *The Review of Economic Studies*, 32(2):85–104, 1965.

Appendix A. Proofs

For the proofs using nonstandard analysis, we work in a κ-saturated, nonstandard enlargement of a superstructure $V(Z)$ where the base set, Z, contains $\mathbb{R}$ and ℓ_∞, and κ is a cardinal greater than the cardinality of $V(Z)$. For nearstandard $r \in {}^*\mathbb{R}^k$, ${}^\circ r \in \mathbb{R}^k$ denotes the standard part of r [26, §II.1 and II.8] or [30, Ch. 3]. The essential result that we use is [29, Theorem 3.1]: if η is an extreme point in the set of Banach-Mazur limits, then there exists in interval subset of ${}^*\mathbb{N}_0$, $\{T', T'+1, \ldots, T\}$ with $(T - T') \simeq \infty$ such that for all $\mathbf{u} \in \ell_\infty$, $\langle \mathbf{u}, \eta \rangle = {}^\circ\langle {}^*\mathbf{u}, \eta_{T',T} \rangle$ where $\eta_{T',T}$ is the *-uniform distribution on $\{T', \ldots, T\}$.

For bounded sequences, the Hardy-Littlewood Tauberian theorem tells us that $\lim_{T\to\infty} \frac{1}{T+1} \sum_{t=0}^{T} u_t = c$ if and only if $\lim_{\beta\uparrow 1}(1-\beta)\sum_{t=0}^{\infty} u_t \beta^t = c$. We record the nonstandard formulation of this and include a proof for completeness.

Lemma 2. *The following are equivalent:*

(1) $\mathbf{u} \in \mathbf{Erg}$;
(2) *for all* $\eta, \eta' \in p\mathbb{V}$, $\langle \mathbf{u}, \eta \rangle = \langle \mathbf{u}, \eta' \rangle$; *and*
(3) *for all* $\beta, \gamma \simeq 1$, $\beta, \gamma < 1$, $(1-\beta)\sum_{t=0}^{\infty} u_t \beta^t \simeq (1-\gamma)\sum_{t=0}^{\infty} u_t \gamma^t$.

Proof. (1) $\Leftrightarrow$ (2). For $T \in {}^*\mathbb{N}_0$, let η_T denote the uniform distribution on $\{0, 1, \ldots, T\}$ so that $\langle \mathbf{u}, \eta_T \rangle = \frac{1}{T+1} \sum_{t=0}^{T} u_t =: Ave_T(\mathbf{u})$.

By definition, $\mathbf{u} \in \mathbf{Erg}$ iff $\lim_{T_n} Ave_{T_n}(\mathbf{u})$ exists and is equal along all sequences $T_n \to \infty$ in $\mathbb{N}_0$. The existence of this limit is equivalent to the statement that for all infinite T, T', $Ave_T(\mathbf{u}) \simeq Ave_{T'}(\mathbf{u})$. From [29, Theorem 3.1], the extreme points of $\mathbb{V}$ have an expression as $L(\mathbf{u}) = {}^\circ\langle {}^*\mathbf{u}, \eta_T \rangle$ where η_T is the uniform distribution on $\{0, 1, \ldots, T\}$, T infinite. Integrating a linear functional to the same value on all extreme points of its domain implies integrating to the same value on all points.

$(2) \Rightarrow (3)$ Suppose, without loss, that for all infinite t, $\langle \mathbf{u}, \eta_t \rangle \simeq 1$ where η_t is the uniform distribution on $\{0, 1, \ldots, t\}$. The following calculation shows that for $\beta \simeq 1$, the density $y_t = (1-\beta)\beta^t$ is, to within an infinitesimal, a convex combination of uniform distributions on infinite intervals $\{0, 1, \ldots, t\}$,

$$R := (1-\beta)\textstyle\sum_{t=0}^{\infty} u_t \beta^t = (1-\beta)^2 \sum_{t=0}^{\infty} \beta^t (t+1) \langle \mathbf{u}, \eta_t \rangle. \tag{31}$$

There exists infinite $\tau < \tau'$ such that $\sum_{t=0}^{\tau} \beta^t (t+1) \simeq \sum_{t=\tau'}^{\infty} \beta^t (t+1) \simeq 0$. Therefore, because $\|\mathbf{u}\|$ is finite,

$$R \simeq (1-\beta)^2 \textstyle\sum_{t=\tau}^{\tau'} \beta^t (t+1) \langle \mathbf{u}, \eta_t \rangle. \tag{32}$$

For each infinite t, $\langle \mathbf{u}, \eta_t \rangle \simeq 1$, and the sum of the weights $(1-\beta)^2 \beta^t (t+1)$, $t \in \{\tau, \tau+1, \ldots, \tau'\}$ is infinitesimally close to 1. Therefore, $R \simeq 1$.

$(2) \Leftarrow (3)$ Now suppose that for all $\beta \in {}^*(0,1)$, $\beta \simeq 1$, ${}^\circ(1-\beta)\sum_{t=0}^{\infty} u_t \beta^t = 1$, where the "1" is a harmless normalization. Let $\mathcal{G}$ denote the set of standard functions, $g : [0,1] \to \mathbb{R}$ such that for all $\beta < 1, \beta \simeq 1$,

$${}^\circ(1-\beta)\textstyle\sum_{t=0}^{\infty} u_t \beta^t\, {}^*g(\beta^t) = \int_0^1 g(x)\, dx. \tag{33}$$

We detour to develop properties of the class of "good" functions $\mathcal{G}$.

It is clear that $\mathcal{G}$ is a vector space of functions containing the constants. Therefore, to show that $\mathcal{G}$ contains the polynomials, it is sufficient to show that it contains the monomials, $g(x) = x^k$. For $g(x) = x^k$, we have $\int_0^1 g(x)\, dx = \frac{1}{k+1}$. We now show that for $g(x) = x^k$, $(1-\beta)\sum_{t=0}^{\infty} u_t \beta^t g(\beta^t) \simeq \frac{1}{k+1}$. For $\beta \simeq 1, \beta < 1$, we have $\beta^{k+1} \simeq 1$ and $\beta^{k+1} < 1$. By assumption, ${}^\circ(1-\beta^{k+1})\sum_{t=0}^{\infty} u_t (\beta^{k+1})^t = 1$. Therefore,

$$(1-\beta)\sum_{t=0}^{\infty} u_t \beta^t g(\beta^t) = (1-\beta)\sum_{t=0}^{\infty} u_t (\beta^{k+1})^t \tag{34}$$

$$= \frac{(1-\beta)}{(1-\beta^{k+1})}(1-\beta^{k+1})\textstyle\sum_{t=0}^{\infty} u_t (\beta^{k+1})^t \tag{35}$$

$$\simeq \frac{(1-\beta)}{(1-\beta^{k+1})} \simeq \frac{1}{k+1} \tag{36}$$

where the last "$\simeq$" follows from l'Hôpital's rule.

Therefore, $\mathcal{G}$ contains all of the polynomials. It is also closed under uniform convergence, so by the (Stone-)Weierstrass theorem, it contains all of the continuous functions. Consider the function

$$g(x) = \begin{cases} 0 & \text{if } x < 1/e \\ 1/x & \text{if } 1/e \leqslant x \leqslant 1. \end{cases} \tag{37}$$

Note that $\int_0^1 g(x) = 1$, and that, if $g \in \mathcal{G}$, then taking $\beta = e^{1/N}$ for infinite N, $(1-\beta)\sum_{t=0}^{\infty} u_t \beta^t g(\beta^t) = \frac{1}{N}\sum_{t=0}^{N} u_t \simeq 1$.

The function g is not continuous, but for every $\epsilon > 0$, it is sandwiched between continuous functions that are within ϵ of g outside of the interval $(1/e-\epsilon, 1/e+\epsilon)$, and that integrate to within ϵ of the integral of g on the interval $(1-\epsilon, 1+\epsilon)$. Since ϵ is arbitrary and the continuous functions belong to $\mathcal{G}$, for any $\beta \simeq 1$, $\beta < 1$, $(1-\beta)\sum_{t=0}^{\infty} u_t \beta^t \simeq 1$. □

Proof of Lemma 1. Suppose that π is $\mathcal{o}(T)$ and fix $\mathbf{u} \in \ell_\infty$. For any infinite $T \in {}^*\mathbb{N}_0$, let $t^\dagger$ be the largest $t \in \{0, 1, \ldots, T\}$ such that $\pi(t) < 0$ and let $t^\ddagger$ be the minimum $t \in \{0, 1, \ldots, T\}$ such that $\pi(t) > T$. We show that $t^\dagger/(T+1) \simeq 0$ and $t^\ddagger/(T+1) \simeq 1$. If ${}^\circ t^\dagger/T = \alpha > 0$, then $t^\dagger$ is infinite and $|\pi(t^\dagger) - t^\dagger|/t^\dagger = 1$, contradicting π being $\mathcal{o}(T)$. Note that $t^\ddagger$ is infinite, and if ${}^\circ t^\ddagger/T = (1-\epsilon) < 1$, then ${}^\circ|\pi(t^\dagger) - t^\dagger|/t^\dagger \geqslant \epsilon$, contradicting π being $\mathcal{o}(T)$. This is sufficient to show that $\mathbf{u} \succsim_{\mathcal{o}(T)} \mathbf{u}^\pi$ because, letting $T' = \{t \in \{0, 1, \ldots, T\} : 0 \leqslant \pi(t) \leqslant T\}$,

$$\begin{aligned}
\left|\frac{1}{T+1}\sum_{t=0}^{T}(u_t - u_t^\pi)\right| &= \left|\frac{1}{T+1}\sum_{t=0}^{T}(u_t - u_{\pi^{-1}(t)})\right| \\
&\leqslant \left|\frac{1}{T+1}\sum_{t\in T'}(u_t - u_t)\right| + \frac{1}{T+1}\sum_{t\notin T', t\leqslant T}|u_t - u_t^\pi| \\
&\leqslant 0 + 2\|u\| \cdot \frac{\#\{t \notin T'\}}{T+1} \simeq 0
\end{aligned}$$

The arguments for $\mathbf{u}^\pi \succsim_{\mathcal{o}(T)} \mathbf{u}$ are essentially identical.

For second part of the claim, note that the arguments just given show that for any infinite T and $\mathbf{u} \in \mathbf{Erg}$, $Ave_T(\mathbf{u}) = Ave_T(\mathbf{u}^\pi)$. □

Proof of Theorem A. Suppose that $\succ$ satisfies Postulates I-V. We draw heavily on Theorem 4, Ch. 3 in [20] which characterizes expected utility preference relations $\succ$. Our Postulates I and II are Fishburn's A1 and A2, our Postulate III is a strong form of his A3 and it directly implies his A4* and A5*.

Our mixture set, $\mathcal{M}$, the set of measures with bounded support is closed under finite convex combinations. As we work with the Borel σ-field on $\mathbf{W}$, our Postulate III implies that the domain for our probabilities contains all preference intervals, which then implies that $\mathcal{M}$ is closed under taking conditional measures on preference intervals, completing the verification of Fishburn's A0.2.

Having verified A1-A5 and A0.2, Fishburn's result shows that there exists an integrable $S : \mathbf{W} \to \mathbb{R}$ such that $p \succ q$ iff $\int S(\boldsymbol{\nu})\, dp(\boldsymbol{\nu}) > \int S(\boldsymbol{\nu})\, dq(\boldsymbol{\nu})$. Restricted to the closed set of point masses, $\{\delta_{\mathbf{u}} : \mathbf{u} \in \mathbf{W}\}$, Postulate III implies that $S(\cdot)$ is continuous. By considering measures with two point supports and their resultants, Postulate IV implies that $S(\cdot)$ is concave. For the normalization that $S : \mathbf{W} \to [0, \infty)$, note that there is no loss in setting $S(\mathbf{0}) = 0$ — for all $\mathbf{u} \in \mathbf{W}$, $\mathbf{u} \succsim_{\mathcal{O}(T)} \mathbf{0}$ so that Postulate V implies $S(\mathbf{u}) \geqslant S(\mathbf{0})$.

We now show that $S(\cdot)$ is exactly Pareto. If B is a null coalition, then for any $\mathbf{u} \in \mathbf{W}$ and any $r > 0$, $\mathbf{u} \succsim_{\mathcal{O}(T)} (\mathbf{u} + r1_B) \succsim_{\mathcal{O}(T)} \mathbf{u}$ so that Postulate V implies $S(\mathbf{u}) = S(\mathbf{u} + r1_B)$. If B is a non-null coaltion, then for any $\mathbf{u} \in \mathbf{W}$ and any $r > 0$, $(\mathbf{u} + r1_B) \succ_{\mathcal{O}(T)} \mathbf{u}$ so that Postulate V implies $S(\mathbf{u} + r1_B) > S(\mathbf{u} + r1_B)$.

We now show that $S(\cdot)$ is patient. From Lemma 1, if π is an $\mathcal{O}(T)$-permutation, then for all $\mathbf{u} \in \ell_\infty$, $\mathbf{u} \succsim_{\mathcal{O}(T)} \mathbf{u}^\pi \succsim_{\mathcal{O}(T)} \mathbf{u}$ so that Postulate V implies $S(\mathbf{u}) \geqslant S(\mathbf{u}^\pi) \geqslant S(\mathbf{u})$.

Now suppose that there exists a continuous, concave $S : \mathbf{W} \to [0, \infty)$ such that $[p \succ q] \Leftrightarrow [\int S(\mathbf{u})\, dp(\mathbf{u}) > \int S(\mathbf{u})\, dq(\mathbf{u})]$ with $S(\cdot)$ satisfying the properties (1) and (2). Verification of Postulates I through IV is routine. To verify Postulate V, suppose that $\mathbf{u} \succ_{\mathcal{O}(T)} \mathbf{v}$ and let $r' = \liminf_{T\to\infty} \frac{1}{T+1} \sum_{t=0}^{T} (u_t - v_t)$. By definition, $r' > 0$. Set $r = r'/2$ and note that $\mathbf{u} \succ_{\mathcal{O}(T)} (\mathbf{v} + r1_{\mathbb{N}_0}) \succ_{\mathcal{O}(T)} \mathbf{v}$ and that $\mathbb{N}_0$ is a non-null coaltion. By property (2), $S(\mathbf{u}) > S(\mathbf{v} + r1_{\mathbb{N}_0}) > S(\mathbf{v})$. □

Proof of Theorem B. We first show that any $\mathbb{V}$-concave $S(\cdot)$ is perfectly Pareto. Pick an arbitrary $\mathbf{u} \in \mathbf{W}$.

- If B is a null coalition, we must show that $S(\mathbf{u} + r1_B) \geqslant S(\mathbf{u}) \geqslant S(\mathbf{u} + r1_B)$ for all $r > 0$. For all $\eta \in p\mathbb{V}$, $\langle 1_B, \eta \rangle = 0$. Because $\mathbb{V}$ is the closed cone containing $p\mathbb{V}$, for all $L \in \mathbb{V}$ and all $r \in \mathbb{R}$, $L(r1_B) = 0$. For any $L \in \mathbf{DS}(\mathbf{u})$, $S(\mathbf{u}) + L((\mathbf{u} + r1_B) - \mathbf{u}) \geqslant S(\mathbf{u} + r1_B)$ so that $S(\mathbf{u}) \geqslant S(\mathbf{u} + r1_B)$. Similarly, for any $L' \in \mathbf{DS}(\mathbf{u} + r1_B)$, $S(\mathbf{u} + r1B) + L'(\mathbf{u} - (\mathbf{u} + r1_B)) \geqslant S(\mathbf{u})$ so that $S(\mathbf{u} + r1_B) \geqslant S(\mathbf{u})$.
- If B is a non-null coalition, we must show that $S(\mathbf{u}+r1_B) > S(\mathbf{u})$ for all $r > 0$. For all $\eta \in p\mathbb{V}$, $\langle 1_B, \eta \rangle > 0$. Because $\mathbb{V}$ is the closed cone containing $p\mathbb{V}$ and any $L' \in \mathbf{DS}(\mathbf{u} + r1_B)$ is a strictly positive element of $\mathbb{V}$, $L'(r1_B) > 0$. Since $S(\mathbf{u} + r1_B) + L'(\mathbf{u} - (\mathbf{u} + r1_B)) \geqslant S(\mathbf{u})$, we have $S(\mathbf{u} + r1_B) \geqslant S(\mathbf{u}) + L'(r1_B)$ so that $S(\mathbf{u} + r1_B) > S(\mathbf{u})$.

We now show that any $\mathbb{V}$-concave $S(\cdot)$ is patient. Pick an arbitrary $\mathbf{u} \in \mathbf{W}$ and an arbitrary $\mathcal{O}(T)$-permutation. We must show that $S(\mathbf{u}) = S(\mathbf{u}^\pi)$. From Lemma 1, we know that $\mathbf{u} \succsim_{\mathcal{O}(T)} \mathbf{u}^\pi \succsim_{\mathcal{O}(T)} \mathbf{u}$ and that for every $\eta \in \mathbb{V}$ and every $\mathbf{z} \in N$, $\langle \mathbf{z}, \eta \rangle = 0$.

Therefore, for any $L \in \mathbf{DS}(\mathbf{u})$ and any $L' \in \mathbf{DS}(\mathbf{u}^\pi)$, $L(\mathbf{u} - \mathbf{u}^\pi) = L'(\mathbf{u} - \mathbf{u}^\pi) = 0$. Therefore $S(\mathbf{u}) + 0 \geqslant S(\mathbf{u}^\pi)$ and $S(\mathbf{u}^\pi) + 0 \geqslant S(\mathbf{u})$.

For the second part, assume that $\succ$ satisfies Postulates I-V and that integrals against $S(\cdot)$ represent the preferences on $\mathcal{M}$. We must show that $S(\cdot)$ is $\mathbb{V}$-concave on int($\mathbf{W}$). Pick an arbitrary $\mathbf{u} \in$ int($\mathbf{W}$) and $L \in \mathbf{DS}(u)$. If $L \notin \mathbb{V}$, then there exist $\boldsymbol{v}_1, \boldsymbol{v}_2 \in \mathbf{Erg}$ with $\mathbf{lra}(\boldsymbol{v}_1) = \mathbf{lra}(\boldsymbol{v}_2)$ and $L(\boldsymbol{v}_1 - \boldsymbol{v}_2) \neq 0$. Reversing the role of $\boldsymbol{v}_1$ and $\boldsymbol{v}_2$ if necessary, $L(\boldsymbol{v}_1 - \boldsymbol{v}_2) < 0$. Because $\mathbf{u}$ is interior, for some $r > 0$, $\boldsymbol{v} := \mathbf{u} + r(\boldsymbol{v}_1 - \boldsymbol{v}_2)$ is interior.

We first prove the intermediate claim that $\mathbf{u} \succsim_{\mathcal{O}(T)} \boldsymbol{v} \succsim_{\mathcal{O}(T)} \mathbf{u}$. For any infinite T, $\frac{1}{T+1}\sum_{t=0}^{T} v_t = \frac{1}{T+1}\sum_{t=0}^{T} u_t + r\frac{1}{T+1}\sum_{t=0}^{T}(v_{1,t} - v_{2,t}) = \frac{1}{T+1}\sum_{t=0}^{T} u_t + 0$. Therefore, $\liminf_T \frac{1}{T+1}\sum_{t=0}^{T} v_t = \liminf_T \frac{1}{T+1}\sum_{t=0}^{T} u_t$.

Returning to the argument, by Postulates V and continuity, $S(\mathbf{u}) = S(\boldsymbol{v})$. However, by the properties of tangents for concave functions, $S(\mathbf{u}) + L(\boldsymbol{v} - \mathbf{u}) \geqslant S(\boldsymbol{v})$. However, $L(\boldsymbol{v} - \mathbf{u}) = rL(\boldsymbol{v}_1 - \boldsymbol{v}_2) < 0$, a contradiction. $\square$

Proof of Proposition 1. As this is an exchange economy model, we need only verify (i)-(iv) in Bewley's Theorem 1. (i) is the assumption that the consumption sets are convex and Mackey closed, which is immediate. (ii) is the assumption that the preference relations are transitive, reflexive and complete, which is satisfied because the preferences are given by utility functions. (iii) is the assumption that for all $i \in I$ and all consumption vectors, $\boldsymbol{x}$, the set $\{\boldsymbol{z} \in W^k : U_i(\boldsymbol{z}) \geqslant U_i(\boldsymbol{x})\}$ is convex and Mackey closed. Convexity follows from the concavity of $U_i(\cdot)$, and a convex subset of the dual of a Banach space is closed for all the topologies between the weak*-topology and the norm topology if and only if it is norm closed [18, Cor. V.2.14]. Therefore, norm continuity of the $U_i(\cdot)$ delivers the necessary closure. (iv) is the the assumption that for all $i \in I$ and all consumption vectors, $\boldsymbol{x}$, the set $\{\boldsymbol{z} \in W^k : U_i(\boldsymbol{z}) \leqslant U_i(\boldsymbol{x})\}$ is norm closed, which follows directly from the norm continuity of $U_i(\cdot)$. $\square$

Proof of Proposition 2. For convex sets, weak* and norm closure are equivalent. $\square$

DEPT. OF ECONOMICS, U. C. RIVERSIDE AND U. T AUSTIN