Preliminary draft, October 2, 2015

PATIENT, PERFECTLY PARETO PREFERENCES: PROGRAMMING AND THE PRECAUTIONARY PRINCIPLE

URMEE KHAN, MAXWELL B. STINCHCOMBE

Abstract. Society is an aggregate of present and future generations. We study stochastic inter-generational programming problems in which similar treatment of generations in similar situations is possible. For such problems, all patient, inequality averse societal welfare functions that are perfectly Pareto responsive have the same optimal policies. When the outcomes of irreversible decisions are partially learnable, the optimal policies for patience preferences yield a variant of the precautionary principle. Under mild conditions, optimal policies exist and there is a single Bellman-like equation characterizing them.

Contents

> ... 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. (Francis, 2015)

> 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, 2010)

> 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)

1. Introduction

In his *Traité de la Culture des Forêts*, written in the late 1600's, Sèbastien Le Prestre de Vauban, Louis XIV's defense minister, noted several aspects of the economics and ecology forests, aspects that complicate the analysis of good societal practices for foresty.[1] First, forests, being a free/easy access resource, were systematically over-exploited in France. Second, after replanting, forests start being productive in slightly less than 100 years but don't become fully productive for 200 years. Third, no private enterprise could conceivably have so long 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.[2] The calculations behind his conclusion assumed that society would be around for at least the next 200 years to enjoy the net benefits — he summed the un-discounted benefits, delayed and large, and compared them to the summed un-discounted costs, early and small.

Present threats to the free/easy access resources from oceans and forests may also be threats 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 the well-being of significant portions of future generations. In the presence of decisions with extremely long-lasting effects and the 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.

We study intergenerational maximization problems in which a conditional equal treatment assumption is plausible, in which it is reasonable to assume that generations in similar situations make similar decisions. For such problems, we study societal welfare functions defined on sequences of numerical measures of generational

[1]See e.g. the edited collections of his writings (Vauban, 1910) or (Vauban, 2007).

[2]As well as the government and the church, Vauban identified 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.

well-being. The concavity of these functions captures inequality aversion, and the immunity to bounded permutations of the measures of well-being captures both equity and patience. These properties alone do not deliver a satisfactory theory.

1.1. **Testing Theory with Applications.** One cannot fully understand a class of preferences without knowing their implications in the analysis of problems of interest.[3] Equal treatment, patience, and inequality aversion do not deliver a satisfactory theory of preferences for Markovian decision problems, one needs more. An example is the 'Rawlsian' preferences on sequences of utilities, preferences represented by the inequality averse (concave), patient (indifferent to finite permutations) function $S_{Rawls}((u_0, u_1, u_2, \ldots)) = \liminf_{t \uparrow \infty} u_t$.

Suppose that the stochastic development of state variables, such as the state of the world's ecosystem, depends on choices in good as well as in bad states, and that sacrifices of present utility in bad states make the good states more likely in the future while sacrifices in the good state make future bad state less likely. Rawslian preferences ignore all but the welfare of those unlucky generations that find themselves in the worst states. This means that optimality for such preferences never requires any sacrifices: in the bad states, sacrifices affect the unluckiest, and are therefore ruled out, no matter how large the benefit to future generations; in the good states, any action, no matter how likely it is to lead to the bad state, is an optimal choice.

Rawlsian preferences have an essential flaw — the value of the social welfare function can be determined by an arbitrarily small proportion of the unluckiest generations in society. By contrast, for patient, inequality averse utility functions such as $\liminf_{T \uparrow \infty} \frac{1}{T+1} \sum_{t=0}^{T} u_t$ or $\liminf_{\beta \uparrow 1} (1 - \beta) \sum_{t=0}^{\infty} u_t \beta^t$, increasing or decreasing the welfare of any non-vanishing proportion of future generations increases or decreases social welfare. We develop the general theory of inequality averse, patient preferences with this kind of Pareto responsiveness.

1.2. **Patient Preferences, Dictatorship and the Pareto Principle.** The concave societal welfare functions under study are defined by the property that they have, as tangent functionals, the limits of expected average utilities of present and future generations, the limits being taken as the expected duration of society becomes arbitrarily large. Preferences of this form are patient in the sense that the ordering is indifferent between $o(t)$ permutations of the generations' measured well-being. The extensive social choice literature on the implications of this indifference to permutations has two main lessons, one concerning the aptly named "dictatorship of the future," the other concerning the Pareto principle.

The "present" is a salient finite set of generations, and a societal welfare function that privileges the future over arbitaily large but still finite versions of the present seems ethically flawed. The dictatorship of the future (Chichilnisky, 1996) is a consequence of indifference to bounded permutations of the generations' well-being, a property that implies that the societal welfare function ignores changes of the welfare

[3]This methodological position is expressed in (Atkinson, 2001, p. 206), "By applying ethical criteria to concrete economic models, we learn about their consequences, and this may change our views about their attractiveness."

of any finite number of generations. It is here that equal treatment enters our analysis — we circumvent such problems by restricting the domain on which the patient preferences are applied to stochastic sequences generated by societal decisions in which generations in similar situations are treated similarly. If all generations in similar situations make similar decisions, then the only generations that can be "mistreated" are those that find themselves in situations that do not occur. For us, the existence of patient optima with conditional equal treatment properties outweighs the existence of alternative, non-stationary, ethically suspect optima.

Difficulties with the Pareto principle arise from a second implication of indifference to bounded permutations (see Fleurbaey and Michel (2003) for extensive references and a thorough examination). The tangent functionals to a concave function S representing such preferences must be, up to a positive scaling factor, integrals against purely finitely additive probabilities. A probability η is purely finitely additive if and only if for all $u = (u_0, u_1, \ldots)$ with $u_t > 0$ and $u_t \to 0$, we have $\int_{\mathbb{N}_0} u_t \, d\eta(t) = 0$. Having tangents with these properties means $S(u) > S(0)$ or $S(v + u) > S(v)$ cannot happen — all moves from 0 toward u or toward from v toward $v + u$ have slope 0, a seeming impossibility result for the existence of Pareto-respecting social welfare functionals. However, if $u_t \to 0$ and η is purely finitely additive, then for all $\epsilon > 0$, a probability 1 part of society is receiving a utility boost less than ϵ. By contrast, we require that $S(v + u) > S(v)$ when and only when a $u \geq 0$ delivers a strictly positive amount to a non-negligible proportion of society.

1.3. **Two Classes of Problems.** We study the workings of social welfare functions in two classes of problems: stochastic dynamic decision problems with Markovian structures; and stochastic dynamic decision problems with partially hidden states relating to the consequences of irreversible decisions.

Many irreversible decisions have unknown consequences, and one can often gather a great deal of information about the consequences before the decision is made. In such problems, the optimal policies for any responsive patient societal welfare function have two properties: they call for expending societal resources on research until its marginal value goes to 0; and once no more information is forthcoming, call for making/not making the decision as the expected value of change in the utility of future generations is positive/negative.

Some versions of the precautionary principle allocate 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. Other versions of the precautionary principle concern the requisite degree of certainty of benefits before a potentially irreversible action should be taken. Our analysis suggests that optimal policies should do not avoid risky decisions, rather, they take risky decisions only after sufficient study.

The main result for the second class of problems is a single version of the Bellman equation that works for *all* of our preferences. This is a surprising unification. It works because patient preferences that ignore welfare gains only acquired by vanishingly small parts of society have tangent with a property that interacts strongly with Markovian decision processes — with probability 1 for a weakly ergodic decision process, the sequences of utilities are indifferent.

This indifference leads to a single Bellman equation, and this is very good news for applications and analysis, but it seems to be decidedly mixed news for incorporating inequality and risk aversion into our analyses. In decision problems where it is feasible to have different sequences of utilities with the same long run average and different long run variances, our preferences are indifferent. However, one simple expedient can remedy this at the expense of slightly complicating the single Bellman equation: make every generation's well-being dependent, in a concave fashion, on the well-being of their children.

1.4. **Literature.** Das Gupta (several, including SAET 2015 address).

Figuières and Tidball Figuières and Tidball (2012) (taking convex combinations in policy space, that is, weighted average of LRAC optimal and β optimal).

Fleurbaey and Michel (2003) (their Thm 2 is incompatibility of weak Pareto, patience and strong continuity, their Thm 5 is the free ultrafilter i.e. purely finitely additive approach).

Basu and Mitra Basu and Mitra (2003) and Basu and Mitra (2007) are prominent and thorough statements of the "impossibility" of satisfying the Pareto principle.

Marinacci (1998) (concave with various subsets of translation invariant tangents, the new criteria use translation invariant tangents that can ignore Pareto optimality in our sense, this by focusing on far tail, by contrast, ours are "anchored" in the present).

Chichilnisky (1996), weighted combination of β discounting and an integral against an unspecified purely finitely additive probability, non-stationary optimal policies that smooth out to stationarity.

Heal (1997) argues that governments look at most decades ahead, (compare Vauban).

Zuber and Asheim (2012), Asheim (2010), Asheim and Zuber (2014) (various concave functionals with/without the presence of the risk of the end of society).

Asheim, Bucholz and Tungodden have a book/paper "Justifying Sustainability" with insights into patient programming.

Examples with long-term benefits: the armory system and the second industrial revolution; National Parks; extinctions; DDT and CFC bans; breathability of air in the Los Angeles basin; microchips and the internet; hybrid crops; genetic engineering; UN accounting standards that hide these kinds of benefits.

1.5. **Outline.** Two examples. Theory of patient and inequality averse preferences. Commentary. Applications of patient preferences to Mdps. Irreversible hidden state models. Summary and conclusions.

2. Two Examples

We begin with a brief description of the class of preferences under study but defer a more serious investigation of their properties to the next section. The main structural result of the investigation characterizes preferences that respect the Pareto ordering and ignore vanishingly small portions of society in terms of properties of their tangent functionals. The first example in this section previews the second main result — a policy for a Markovian decision problem is S-optimal for one of our patient social welfare function respecting the Pareto principle if and only if the policy maximizes the long-run average utility. We then examine the implications of our social welfare functions for the structure of optimal policies in the presence of risky, potentially irreversible decisions where research and delay may yield information. In this class of problems, optimal policies yield a version of the precautionary principle.

2.1. **Preferences.** The class of bounded sequences of utilities is denoted ℓ_∞ and defined as the set of infinite length vectors of real numbers, $(u_0, u_1, u_2, \ldots)$, with norm $\|u\| = \sup_{t\in\mathbb{N}_0} |u_t| < \infty$ where $\mathbb{N}_0$ is $\{0, 1, 2, \ldots\}$. A **social welfare function**, is a function, $u \mapsto S(u)$, on the non-negative elements of ℓ_∞. We study social welfare functions with five properties. Three properties are standard, continuity, monotonicity, and concavity/inequality aversion. The properties of concave functions are determined by the properties of their tangents, and tangents provide the best method for described the last three properties, patience, respect for the Pareto ordering, and determination by ergodic means.

2.1.1. *Tangents and Patience.* If τ represents the random time until the end of society, then $L_\tau(u) := E\,\frac{1}{\tau+1}\sum_{t=0}^{\tau} u_t$ is a measure of society's welfare. If $Prob(\tau \leq M)$ is small for large M, then $L_\tau(\cdot)$ is a measure of welfare for a patient society, one confident in its longevity. The mappings $u \mapsto L_\tau(u)$ are continuous linear functionals on ℓ_∞, and we denote by $\mathbb{V}$ the set of accumulation points[4] of the sets of $L_\tau(\cdot)$ with $Prob(\tau \leq M) < \epsilon$ for large M and small ϵ.

A social welfare function, $S(\cdot)$, is $\mathbb{V}$-concave if all of its tangents are, possibly after positive re-scaling, elements of $\mathbb{V}$. The class of $\mathbb{V}$-concave functions demonstrate patience in the following strong fashion: a permutation, π, of $\mathbb{N}_0$ is $o(t)$ if $|\pi(t)-t| = o(t)$, that is, if $\limsup_t \frac{|\pi(t)-t|}{t} = 0$; u^π denotes the stream of utilities u with the generational indexes permuted by π; for all order 1 permutations, all $u \in \ell_\infty$, and all $L \in \mathbb{V}$, $L(u) = L(u^\pi)$; any $\mathbb{V}$-concave social welfare function, $S(\cdot)$, inherits this property from its tangents — $S(u) = S(u^\pi)$.

The class of $\mathbb{V}$-concave social welfare functions has appeared in the previous literature: if $S_{\mathbb{V}}(u) = \min_{L\in\mathbb{V}} L(u)$, that is, if the set of tangents at every u is the entire set $\mathbb{V}$, then for every u, $S_{\mathbb{V}}(u) = \liminf_{T\uparrow\infty} \frac{1}{T}\sum_{t=0}^{T-1} u_t$; using techniques developed in Keller and Moore (1992), one can identify a set $\mathbb{V}' \subset \mathbb{V}$ such that for each u, $S_{\mathbb{V}'}(u) = \min_{L\in\mathbb{V}'} L(u)$ is equal to $\liminf_{\beta\uparrow 1}(1-\beta)\sum_{t=0}^{\infty} u_t\beta^t$.

2.1.2. *Tangents and the Pareto Principle.* A set of generations, $B \subset \mathbb{N}_0$, represents a non-vanishing proportion of society if $\underline{\ell}(B) := \liminf_T \frac{1}{T+1}\sum_{t=0}^{T} 1_B(t) > 0$. Every $L \in \mathbb{V}$ has the property that $L(1_B) \geq \underline{\ell}(B)$ (where 1_B is the indicator function for the set B). This implies that every concave functional with tangents belonging to $\mathbb{V}$ is **Pareto** — for all non-vanishing B, all $r > 0$, and all $u \geq 0$, $S(u + r1_B) > S(u)$. A societal welfare function that respects the Pareto ordering is **perfectly Pareto** if it ignores vanishing small parts of society, that is, if $\overline{\ell}(B) := \limsup_T \frac{1}{T+1}\sum_{t=0}^{T} 1_B(t) = 0$ implies that $S(u + r1_B) = S(u)$. The proof of Theorem A (below) shows that patient

[4] A filter of linear functionals, L_α, weak* converges to L if $L_\alpha(u) \to L(u)$ for all u.

societal welfare functions are perfectly Pareto if and only if their tangents that are positive elements of the closed span of $\mathbb{V}$.

2.1.3. *Tangents and Determination by Ergodic Means.* A sequence of utilities is **weakly ergodic** or **Cesaro summable** if $\lim_{T\to\infty} \frac{1}{T+1}\sum_{t=0}^{T} u_t$ exists, and Erg denotes the set of all weakly ergodic elements of ℓ_∞. The following is the essential result linking patient and perfectly Pareto welfare functions and dynamic programming: $u \in$ Erg if and only if for all $L, L' \in \mathbb{V}$, $L(u) = L'(u) = \lim_{T\to\infty} \frac{1}{T+1}\sum_{t=0}^{T} u_t$.

This matters for stochastic dynamic programming because, under mild conditions on the problem, a long-run average optimal policy exists. Under that policy, there is a probability 1 set of outcome paths along which the sequence of utilities belongs to Erg and they all have the same long-run average. Any social welfare function with tangents in $\mathbb{V}$ ranks alternative policies with these properties in exactly the same fashion.

2.2. **Climate Change Choices.** We now turn to the problem of finding S-optimal policies in a particular Markovian decision problem.

2.2.1. *A Stark Model.* 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 foodstuff in the amounts humans are evolved to need: in the undamaged state, the seas and forests are able to produce oxygen concentrations supporting life as we currently know it. In the undamaged state, $x = G$, society chooses the transition probability, r to the damaged state, $x = B$, with $0 < \underline{r} \leq r \leq \overline{r} < 1$. The expected utility of choosing r is $u_G(r)$, and higher choices of r lead to a higher expected utility for the present generation, $u'_G(r) > 0$. In a parallel fashion, in the damaged state, society chooses the transition probability, s to the undamaged state with $0 < \underline{s} \leq s \leq \overline{s} < 1$, and higher choices of s lead to lower expected utility of the present generation, $u'_B(s) < 0$. 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.

Starting from the present, $t = 0$, a **policy**, w, chooses an r and a s as a function of the present state. This choice gives rise to a Markov process, $\Phi^w = (\Phi_t^w)_{t\in\mathbb{N}_0}$, taking either the value G or B. A policy w is S-**optimal** if it maximizes $E\left(S(\Phi)|\Phi_0 = x\right)$ for each x. The easiest way to develop the optimality equations uses the expression of the Markov process Φ^w as a sequence of i.i.d. (independent and identically distributed) **sojourns**, here adapted to a two-state process.

If α is an atom of a Markov process, then from any recurrent state, $x \neq \alpha$, the process will find its way back to α in a random time with finite expectation. The time path of the process during periods between a departure from α and a first return to α is called a **sojourn from** α because the process is temporarily staying someplace other than α. In a the present model, either state can be regarded as an atom.

2.2.2. *The Long-Run Average Optimality Equations.* Because the probabilities r and s are interior, any policy $w = (r, s)$ leads to the process Φ^w having a well-defined long-run average, $\rho = \rho^w := E \lim_{T\to\infty} \frac{1}{T+1} u(\Phi_t^w)$. For a state x, let $h^w(x)$ be the expected deviations from ρ until the state changes to y while using the policy w,

$$h^w(x) = E\left(\sum_{t=0}^{\tau_y - 1}(u(\Phi_t^w) - \rho)|\Phi_0 = x\right) \tag{1}$$

where τ_y is the random time that Φ^w first changes from state x to $y \neq x$ given that $\Phi_0^w = x$. The analogue to the Bellman equation for patient optimization is a functional equation in $h(\cdot)$.

If a policy w is optimal, then after the process leaves a state y, it will spend a random amount of time in state x. During that sojourn, actions will determine current utility and the distribution of the random time until the return to the state y. When the current utilities in the sojourn are above the long-run average, ρ^w, the tradeoff is between increasing current utilities and making the sojourn longer. By contrast, when the current utilities in the sojourn and below the long-run average, the tradeoff is between sacrificing present utility and making the sojourn shorter.

Let E_x^w be the expectation operator when using a policy w and starting from $\Phi_0 = x$. To characterize an optimal w, we need to find numbers ρ, $h(G)$, and $h(B)$ such that

$$\rho + h(G) = \max_{r\in[\underline{r},\overline{r}]} \left[u_G(r) + E_G^w h(\Phi_1)\right], \text{ and} \tag{2}$$

$$\rho + h(B) = \max_{s\in[\underline{s},\overline{s}]} \left[u_B(s) + E_B^W h(\Phi_1)\right]. \tag{3}$$

In parallel with the Bellman equations for a discounted stochastic dynamic programming problem, the problems in (2) and (3) maximize a present utility plus an expected value of where the process will arrive one step into the future. The difference is the replacement of the discounted value function with the expected deviation from the long-run average during the sojourn.

In this model, $E_G^w h(\Phi_1) = (1-r)h(G) + rh(B)$ and $E_B^w h(\Phi_1) = sh(G) + (1-s)h(B)$. Therefore, the first order equations (FOCs) for an interior solution to (2) and (3) are

$$u_G'(r) = [h(G) - h(B)] \text{ and } u_B'(s) = [h(B) - h(G)]. \tag{4}$$

In this particular model, $h(G) = (u_G(r) - \rho)\cdot E\,\tau_B$, $E\,\tau_B = (1-r)/r$, $h^w(B) = (u_B(s) - \rho)\cdot E\,\tau_G$, $E\,\tau_G = (1-s)/s$ because 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 $[h^w(G) - h^w(B)]$ to be a large positive number. From this, one expects the right-hand sides of these 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, $w^* = (\underline{r}, \overline{s})$.[5]

By contrast, consider the policies that myopically maximize utility in the damaged state, the policies of the form $(r, \underline{s})$. Any such policy maximizes the expected value of the Rawlsian (lim inf) social welfare function, S_{Rawls} because $S_{Rawls}(u) = u_B(s)$ with probability 1 for *any* policy. It is the failure of the Rawlsian ordering, $S_{Rawls}(\cdot)$, to

[5]In this simple model, it is easy to verify that each $w = (r, s)$ gives rise to the long-run average $\rho = \rho^w(r, s) = \frac{s}{r+s}u_G(r) + \frac{r}{r+s}u_B(s)$. Verifying that the FOCs $\partial\rho/\partial r = 0$ and $\partial\rho/\partial s = 0$ reduce to (4) is routine.

respect the Pareto ordering 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 Rawlsian ordering, and this precludes making tradeoffs between the welfare of different proportions of the generations that make up society.

2.2.3. *The Long-Run $S(\cdot)$-Optimal Equations.* As we argued above, one does not fully understand a class of preferences or the set of assumptions behind them until one knows their implications in the analysis of problems of interest. One of the two main results in this paper shows that a policy for a Markovian decision problem is S-optimal for a $\mathbb{V}$-concave societal welfare function if and only if the policy maximizes the long-run average utility. For the analysis of Markovian models, this means that we can appeal to all of the techniques that have been developed to insure the existence of solutions to equations such as (2) and (3) (Meyn, 1997, See especially) to find S-optimal policies. Here, we give some intuitions for this result.

Following a policy $w = (r, s)$, gives rise, with probability 1, to a sequence of utilities $u = (u_0, u_1, u_2, \ldots)$ that is ergodic/Cesaro summable, $\lim_{T\to\infty} \frac{1}{T+1} \sum_{t=0}^{T} u_t$ exists. Further, with probability 1, the limit is the same along all paths, and it was denoted $\rho = \rho^w$ above. Every L in the class $\mathbb{V}$ has the property that for all u, u' with the same Cesaro sum/long run average, $L(u) = L(u')$. Putting the pieces together, with probability 1, all of tangent functionals treat all of the realizations the same. This in turn means that, restricted to a probability 1 set of realizations, $S(u) = \varphi(\rho)$ where $\varphi(\cdot)$ is concave and strictly increasing. As a result, maximizing $\rho = \rho^w$ is necessary and sufficient for finding an S-optimal policy.

2.3. **Patience, Irreversibility, and the Precautionary Principle.** In many stochastic dynamic decision problems, the decision maker has only partial information about the state variables or the consequences of actions. Let us again consider a specific example to heighten the intuitions about maximizing patient preferences in such contexts. A hidden state X takes the two values $x_l < 0 < x_h$, with strictly positive probabilities $(1 - g)$ and g, and we assume that $E\,X < 0$. When the action $a = 1$ is taken, utility will go up/down by X forever thereafter and no further actions are available. When the action $a = 0$ is taken, utility will be unchanged forever thereafter and 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 perhaps related to X will be observed. There is a random number of signals informative about the value of X. Every attempt to observe an informative signal costs c.

The random number of informative signals is another hidden state, M, distributed with probability p_m, $m = 0, 1, \ldots, \infty$ where $p_\infty > 0$ means that the distribution is incomplete and an infinite amount of information can be gathered because M gives the number of informative signals that can be observed. If $m \leq M$, then the signals are Bernoulli, that is, distributed iid $Bern(q)$ if $X = x_h$, and distributed iid $Bern(1 - p)$ if $X = x_l$, where $\frac{1}{2} < p, q < 1$. If $m > N$, then signals $s_{N+1}, \ldots, s_m$ are distributed iid $Bern(\pi)$ for some $\pi \in (0, 1)$.

For any patient preferences we have the following version of the precautionary principle: if there is a positive probability that beliefs will converge to $E^\beta X > 0$, then in any optimal path, signals will be observed until beliefs β have converged to either $E^\beta X > 0$ or $E^\beta X \leq 0$; and it cannot be optimal to never make a decision.

For $s \in \{0,1\}^n$, define $Bin(r,n)(s)$ as the probability that a binomial with parameters r and n takes the value s. The posterior beliefs after making m observations and observing $s \in \{0,1\}^m$ are

$$\beta(x_h|s,m) = P(M > m)\frac{g \cdot Bin(q,m)(s)}{g \cdot Bin(q,m)(s) + (1-g)\cdot Bin(1-p,m)(s)} \tag{5}$$
$$+ \sum_{n=0}^{M} p_n \frac{g \cdot R(q,n,m)(s)}{g \cdot R(q,n,m)(s) + (1-g)\cdot R(1-p,n,m)(s)}$$

where $R(\rho,n,m) = Bin(\rho,n)((s_i)_{i=1}^n) \cdot Bin(\pi, m-n)((s_i)_{i=n+1}^m)$.

Observation: picking M observations of s followed either by $a = 0$ or $a = 1$ so as to maximize $E\sum_{t\geq 0} u_t\delta^t$ has the property that for $\delta \uparrow 1$, the optimal $M^*(\delta)\uparrow$. If $p_m = 0$ for all $m \geq \overline{M}$, then $M^*(\delta) \uparrow \overline{M}$. If $p_m > 0$ for infinitely many m, then $M^*(\delta)\uparrow\infty$. The proof comes from the observation that for any $b > 0$,

$$\frac{c\cdot\sum_{t=0}^{M}\delta^t}{b\cdot\sum_{t>M}\delta^t} \downarrow 0 \text{ as } \delta\uparrow 1.$$

By the tangency arguments, the same result holds for all (continuous, monotonic, inequality averse, patient, and respecting the Pareto ordering) social welfare functions.

3. Societal Welfare Functions

This section is divided into definitions, statements of the results, and commentary. Throughout: the set of sequences of numerical measures of well-being is denoted $\mathbf{W}$ and defined as the non-negative elements of ℓ_∞; $\mathbf{W}$ is endowed with the sup norm, $\|u\| = \sup_{t\in\mathbb{N}_0}|u_t|$ and the associated norm topology; inequality for $u, v \in \mathbf{W}$ is defined coordinatewise, $u \geq v$ iff $u_t \geq v_t$ for all $t \in \mathbb{N}_0$.

For the purposes of interpretation, it is useful to keep in mind the dynamic programming applications. These have $u_t = v(a_t, \Phi_t)$ where a_t represents the actions taken by generation t, Φ_t represents the state of the system faced by generation t, and $v(\cdot,\cdot)$ measures the associated welfare. For our patient preferences, the optima for stationary problems are stationary under mild conditions on the programming model, but this is not true for the more frequently used definition of patience.

3.1. Definitions.

A social welfare function is a *continuous* function $S : W \to \mathbb{R}_+$. We study social welfare function defined by the property that their tangents belong to a specific class of continuous linear functionals.

3.1.1. $\mathbb{V}$-*Concave Social Welfare Functions.*

We use $\langle u, \eta\rangle$ to denote the bilinear function $(u,\eta) \mapsto \int u\,d\eta$, $u \in \ell_\infty$ and η a finitely additive probability on $\mathbb{N}_0$. A net (generalized sequence) η_α of probabilities on $\mathbb{N}_0$ converges in the weak*-topology to η if $\langle u,\eta_\alpha\rangle \to \langle u,\eta\rangle$ for all $u \in \ell_\infty$. By Alaoglu's theorem, the set of probabilities on $\mathbb{N}_0$ is weak* compact.

Let τ be a random variable taking values in $\mathbb{N}_0$. The mapping $u \mapsto L_\tau(u) := E\,\frac{1}{\tau+1}\sum_{t=0}^{\tau} u_t$ is continuous, linear, positive, and has norm 1. By the usual representation theorems for continuous linear functionals, there exists a unique probability η_τ such that $L_\tau(u) = \langle u, \eta_\tau \rangle$ for all u.

Definition 3.1. *The class* $\mathbb{V}$ *of continuous linear functionals is*

$$\mathbb{V} = \bigcap \{ \mathrm{cl}\left(\{ \eta_\tau : \mathit{Prob}(\tau \le M) \le \epsilon \} \right) : M \in \mathbb{N}_0, \epsilon > 0 \} \tag{6}$$

where closure, $\mathrm{cl}(\cdot)$, *is with respect to the weak*-topology.*

The class $\mathbb{V}$ is non-empty (by the finite intersection property), as well as compact and convex. The tangents will be assumed to belong to $[\gamma^+) \cdot \mathbb{V}$, that is, to the set $\{\gamma' \cdot \eta : \eta \in \mathbb{V}, \gamma' \ge \gamma\}$.

For a social welfare functional S and $u \in \mathbf{W}$, the **set of tangents to S at u** is denoted $\mathbf{D}S(u)$ and defined as the set of continuous linear functionals, L, on ℓ_∞, such that for all $v \in \ell_\infty$,

$$S(v) \le S(u) + L(v - u). \tag{7}$$

Definition 3.2. *A social welfare functional is* ***$\mathbb{V}$-concave*** *if for all $u \in W$, there exists $\gamma > 0$ such that every $L \in \mathbf{D}S(u)$ belongs to $[\gamma^+) \cdot \mathbb{V}$.*

3.1.2. *Patient Social Welfare Functions.* We define the patience of a social welfare functional using $o(t)$ permutations. The set of integers, negative and non-negative is denoted $\mathbb{Z}$ and defined as $\{\ldots, -2, -1, 0, 1, 2, \ldots\}$. A $o(t)$ **permutation** is a 1-to-1 function $\pi : \mathbb{N}_0 \to \mathbb{Z}$ that is onto $\mathbb{N}_0$ and satisfies $\limsup_t \frac{|\pi(t)-t|}{t} = 0$. For $u = (u_0, u_1, u_2, \ldots)$ and a $o(t)$ permutation π, define u^π as $(u_{\pi^{-1}(0)}, u_{\pi^{-1}(1)}, u_{\pi^{-1}(2)}, \ldots)$.

Definition 3.3. *A social welfare function $S : \mathbf{W} \to \mathbb{R}_+$ is* ***patient*** *if $S(u) = S(u^\pi)$ for all $o(t)$ permutations π.*

Bounded permutations satisfy $|\pi(t) - t| \le M$ for all t and some fixed M e.g. $\pi(t) = t - M$. The bounded permutations are a frequently used special case of the $o(t)$ permutations. We will see that defining patience using only immunity to bounded permutations does not deliver a satisfactory theory.

3.1.3. *Pareto Social Welfare Functions.* For $B \subset \mathbb{N}_0$, $\underline{\ell}(B) := \liminf_T \frac{1}{T+1}\sum_{t=0}^{T} 1_B(t)$ and $\overline{\ell}(B) := \limsup_T \frac{1}{T+1}\sum_{t=0}^{T} 1_B(t)$. Subsets of society with $\overline{\ell}(B) = 0$ are vashingly small, subsets with $\underline{\ell}(B) > 0$ are not.

Definition 3.4. *A social welfare function, S, is* ***Pareto*** *if for all B with $\underline{\ell}(B) > 0$ and all $r > 0$, $S(u + r1_B) > S(u)$, a Pareto S is* ***perfectly Pareto*** *if for all B with $\overline{\ell}(B) = 0$ and all $r > 0$, $S(u + r1_B) = S(u)$.*

A perfectly Pareto social welfare function is sensitive to utility boosts for non-negligible portions of society while ignoring utility boosts to negligible portions of society.

3.2. **Results.** There are two results. The first tells us that patience and the Pareto principle are compatible. The second tells us about that patient preferences are easy to use for Markovian decision problems.

Theorem A. *Any $\mathbb{V}$-concave societal welfare function is patient and perfectly Pareto.*

Completing the arguments suggested by the following observations leads to a proof: let η_T denote the uniform distribution on $\{0, 1, \ldots, T\}$; for τ a random time, $L_\tau(u) = E\,\frac{1}{\tau+1}\sum_{t=0}^{T} u_t$ can be expressed as a convex combination of η_T's, $L_\tau(u) = \sum_T \langle u, \eta_T\rangle P(\tau = T)$; an extreme point of this class of distributions has $P(\tau = T) = 1$ for some T; the set $\mathbb{V}$ is a compact and convex limit set; from (Keller and Moore, 1992, Theorem 3.1), the extreme points of $\mathbb{V}$ have expressions as uniform distributions on $\{0, 1, \ldots, T\}$ for infinitely large T; these linear functionals are patient and perfectly Pareto; finally, $S(\cdot)$ inherits these properties from its tangents. Details are in the appendix.

A sequence $u \in \ell_\infty$ is **weakly ergodic** (aka Cesaro summable) if $\text{Ave}(u) := \lim_T \frac{1}{T+1}\sum_{t=0}^{T} u_t$ exists. The set of weakly ergodic elements is denoted Erg. The closed linear subspace Erg is tightly related to the class $\mathbb{V}$.

Lemma 1. $u \in \text{Erg}$ *if and only if for all* $L, L' \in \mathbb{V}$, $L(u) = L'(u) = \text{Ave}(u)$.

Definition 3.5. *A social welfare function is* ***determined by erdogic means on*** Erg *if for all* $u, v \in \text{Erg}$, $S(u) > S(v)$ *iff* $\text{Ave}(u) > \text{Ave}(v)$.

Theorem B. *If* $S : \mathbf{W} \to \mathbb{R}_+$ *is a* $\mathbb{V}$*-concave societal welfare function, then there exists a concave strictly increasing* $\varphi : \mathbb{R}_+ \to \mathbb{R}_+$ *such that for all* $u \in \text{Erg}$, $S(u) = \varphi(\text{Ave}(u))$. *In particular, any* $\mathbb{V}$*-concave societal welfare function is determined by ergodic means on* Erg.

The proof works because any concave function is the lower envelope of the affine functions that majorize it, and for $\mathbb{V}$-concave functions, the linear parts of the majorizing affine functions, after positive re-scaling, satisfy $L(u) = \text{Ave}(u)$ for all $u \in \text{Erg}$. Details are in the appendix.

4. Commentary

Weakening the patience/equity criterion by requiring that the social welfare function be indifferent between u and u^π only for bounded permutations leads to a theory violating the conditional equal treatment property. The domain on which such a theory can be useful replaces the set Erg with the much smaller set of almost convergent sequences, and this smaller set excludes the use of "most" Markov processes. Finally, the "linearity" result of Theorem B, that $S(u) > S(v)$ iff $\text{Ave}(u) > \text{Ave}(v)$ for $u, v \in \text{Erg}$, does not destroy the inequality aversion built into the concavity of $S(\cdot)$.

4.1. **Violating the Conditional Equal Treatment Property.** Our definition of patience requires that $S(u) = S(u^\pi)$ for all $o(t)$ permutations. This is a more stringent requirement than the requirement usually found in the literature, that $S(u) = S(u^\pi)$ only for the bounded permutations, those with $|\pi(t) - t|$ uniformly bounded. The corresponding class of tangent functionals grows from $\mathbb{V}$ to the the set of all Banach-Mazur limits, and the use of this larger class of tangents allows social welfare functions for which optimality requires violations of the equal treatment property.

A continuous linear $L : \ell_\infty \to \mathbb{R}$ with the properties $L(u) \geq 0$ for all $u \geq 0$, $L((1, 1, 1, \ldots)) = 1$, and $L(u^\pi) = L(u)$ for all u and bounded permutations π is called a **Banach-Mazur limit** (Banach, 1978, Ch. II.3). Every Banach-Mazur limit can be represented as an integral against a purely finitely additive probability, $L(u) = \langle u, \eta \rangle$, and the set of all such functionals is denoted $\mathbb{BM}$.

Elements of $\mathbb{BM} \backslash \mathbb{V}$ include accumulation points of uniform distributions on $\{T', \ldots, T\}$ where $(T - T')$ becomes unboundedly large and $\frac{T'}{T}$ becomes arbitrarily close to 1. Let η be such an accumulation point. Because η ignores all but the "tail" of the consumption streams, it simultaneously fails to respect the Pareto criterion and fails to be immune to $\mathcal{o}(t)$ permutations.[6] As a result, there are well-behaved stationary Markovian decision problems for which the optima of e.g. the social welfare function $S(u) = \langle u, \eta \rangle$ are non-stationary in a particularly disturbing fashion. The optima involve "feasts" during a sequence of intervals $\{T'_n, \ldots, T_n\}$ at the expense of "starvation" of the generations between T_n and T'_{n+1}. Further, the set B of starving generations has $\underline{\ell}(B) = 1$ because $\frac{T'}{T}$ becomes arbitrarily close to 1.

4.2. **The Almost Convergent Sequences.** Starting with $U^0 = (u_0, u_1, u_2, \ldots)$, for each $j \in \mathbb{N}_0$, define the points $U^j \in \ell_\infty$ as follows.

U^0	u_0	u_1	u_2	u_3	$\cdots$
U^1	$\frac{u_0+u_1}{2}$	$\frac{u_1+u_2}{2}$	$\frac{u_2+u_3}{2}$	$\frac{u_3+u_4}{2}$	$\cdots$
U^2	$\frac{u_0+u_1+u_2}{3}$	$\frac{u_1+u_2+u_3}{3}$	$\frac{u_2+u_3+u_4}{3}$	$\frac{u_3+u_4+u_5}{3}$	$\cdots$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\cdots$
U^j	$\frac{u_0+\cdots+u_j}{j+1}$	$\frac{u_1+\cdots+u_{j+1}}{j+1}$	$\frac{u_2+\cdots+u_{j+2}}{j+1}$	$\frac{u_3+\cdots+u_{j+3}}{j+1}$	$\cdots$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\ddots$

Definition 4.1. $u = (u_0, u_1, u_2, \ldots) \in \ell_\infty$ *is* ***almost convergent to*** r *if for every* $\epsilon > 0$*, there exists* J *such that for all* $j \geq J$*, every element of* U^j *is within* ϵ *of* r*.*

It is clear that any almost convergent u belongs to Erg, however, the set of almost convergent utilities is too small to contain the outcomes of even the simplest non-degenerate Markov processes.

Example 4.1. *Suppose that for all* $t \in \mathbb{N}_0$*,* u_t *is either* 0 *or* 1 *and that* u_t *follows a Markov process: if* $u_t = 0$ *then* $u_{t+1} = 1$ *with probability* $\alpha \in (0, 1)$*; and if* $u_t = 1$ *then* $u_{t+1} = 0$ *with probability* $\beta \in (0, 1)$*. With probability* 1*, the realizations of this Markov chain belong to* Erg *and* $\text{Ave}(u) = \frac{\alpha}{\alpha+\beta}$*. However, with probability* 1*, for each* $j \in \mathbb{N}_0$*, each realization contains infinitely many length* $(j + 1)$ *sequences of* 0*'s and contain infinitely many length* $(j + 1)$ *sequences of* 1*'s. Therefore each* U^j *contains infinitely many elements equal to* 0 *and infinitely many elements equal to* 1*, hence cannot be almost convergent.*

[6] The Polya index of Marinacci (1998) has this kind of tail-only sensitivity.

The following characterization of almost convergent sequences is (Lorentz, 1948, Theorem 1). It should be compared to Lemma 1.

Theorem C (Lorentz)**.** *u is almost convergent if and only if for every $\eta, \eta' \in \mathbb{BM}$, $\langle u, \eta\rangle = \langle u, \eta'\rangle$.*

As argued above, it is crucial to test theories with applications. Defining patience using tangents immune only to bounded permutations results in a theory that can be fruitfully applied, but only to models in which all optima involve almost convergent sequences of utilities. This is a slightly larger class than e.g. growth models with convergent, deterministic sequences of utilities, but seems smaller than one would wish.

4.3. **Concavity and Inequality Aversion.** The result of Theorem B, that $S(u) > S(v)$ iff $\text{Ave}(u) > \text{Ave}(v)$ for $u, v \in \text{Erg}$, could be interpreted as a linearity result, a repudiation of the inequality and risk aversion built into the concavity of $S(\cdot)$. An analysis of the programming applications, those with $u_t = v(a_t, \Phi_t)$ where a_t represents the actions taken by generation t, Φ_t represents the state of the system faced by generation t, shows that this is misleading in at least two senses.

The first, and perhaps most obvious, way in which this is misleading comes from the observation that Φ_t is stochastic. Variability in Φ_t and concavity of $v(a_t, \cdot)$ makes the expected value of u_t lower. It is the long-run expected value of the u_t that determines $S(u)$ for patient preferences, and riskier paths yield lower expected utility.

The second way in which this is misleading involves entrainment and/or hysteresis (For a review of these are other concepts related to irreversibilities, see Perrings and Brock, 2009). In decision problems with "urn-like" components, early decisions and early stochastic events determine the long run path of the system.[7] 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.

5. Applications: Programming and the Precautionary Principle

5.1. MDPs as in Tweedy PIA.

Theorem D. *In an MDP with a stabilizing policy, a policy is LRAC optimal iff it is S-optimal for all concave, uniformly patient S.*

Details and proof to be added: stabilizing policies may not exist, they require that the optima actions and states not wander off; while there are hopes of upper hemi-continuity kinds of results for policies as the discount factor goes to 1, there are

[7]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."

interesting extinction and related examples of the stark differences in the long-run distributions.

5.2. **State Doubling, or Caring for the Next Generation.** The social welfare functions under study are equivalent to a monotonic transformation of the **linear** functional $u \mapsto L(u)$ on Erg. Therefore, the concavity of $S(\cdot)$ can have no bite on this domain, smoothing of inter-generational consumption patterns cannot matter.

A solution is to make generation t's well-being depend on both their state and actions as well as the state and actions taken by generation $t+1$ and then to maximize $S(\cdot)$. In particular, if each generation dislikes their offspring doing worse than they do, this pushes the solutions toward smoothing.

5.2.1. *In the Climate Change Model.* More care for the next generation has monotone comparative statics effects pushing for every lower r and higher s. Can incorporate resilience and analyze the Poisson-Bellman equations: efforts near tipping points between different basins of attraction are, optimally, much much higher.

5.2.2. *Stochastic Fishery.* Comparison of the optimality equation for a stochastic fishery and a stochastic fishery where u_t depends on the consumption of t and $t+1$. Will see emergence of buffer stock, a more cautious approach.

6. Patience and Precaution

6.1. **Research Until You Know.** Include Deutsch's Deutsch (2011) argument as the guess that irreversibility can be undone. This segues nicely to the increased value of options.

6.2. **Option Values.** In long-run problems, decisions one can recover from are costless. Deutsch's argument *redux*.

7. Summary and Conclusions

Pareto and patience are compatible.

Testing the compatibility with applications leads to $o(t)$ permutations as the definition of patience.

Under mild conditions, there is one Bellman equation for all the patient preferences, but this requires recurrence.

When recurrence involves two (or more) large basins of attraction with different utilities, the optimal efforts near tipping points possibly leading to a worse basin are much much higher.

The urn models/hysteresis analyses require a more "hands on" analysis.

Appendix A. Proofs

We work in a κ-saturated, nonstandard enlargement of a superstructure $V(Z)$ where 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 (§II.1 and II.8 Hurd and Loeb, 1985) or (Ch. 3 Lindstrøm, 1988). The essential result that we use is (Theorem 3.1 Keller and Moore, 1992): 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 $\langle u, \eta \rangle = {}^\circ\langle {}^*u, U_{T',T} \rangle$ where $U_{T',T}$ is the *-uniform distribution on $\{T', \ldots, T\}$.

Proof of Theorem A. Let S be a $\mathbb{V}$-concave societal welfare function.

Patience. Pick arbitrary $u \in \ell_\infty$ and arbitrary order 1 perturbation π. For $L' \in \mathbf{D}S(u)$ and $L'' \in \mathbf{D}S(u^\pi)$, we have

$$S(u^\pi) \leq S(u) + L'(u^\pi - u) \text{ and } S(u) \leq S(u^\pi) + L''(u - u^\pi). \tag{8}$$

Since L and L' are positive scalings of elements of $\mathbb{V}$, it is sufficient to show that for any $L \in \mathbb{V}$, $L(u - u^\pi) = 0$. Since $\mathbb{V}$ is compact and convex, it is sufficient to show that $L_{ext}(u - u^\pi) = 0$ for every extreme $L_{ext} \in \mathbb{V}$. From Keller and Moore (1992), if L_{ext} is an extreme point of $\mathbb{V}$, then there exists an integer T, that is, $T \in {}^*\mathbb{N}_0 \setminus \mathbb{N}_0$, such that for all $u \in \ell_\infty$, $L_{ext}(u) = {}^\circ\langle {}^*u, \eta_T \rangle$ where η_T is the uniform distribution on the *integers $\{0, 1, \ldots, T\}$. The integral of ${}^*u$ against η_T is $\frac{1}{T+1}\sum_{t=0}^T {}^*u_t$, and the integral of $*\, u^\pi$ is $\frac{1}{T+1}\sum_{t=0}^T {}^*u_t^\pi$. It is therefore sufficient to show that

$$d := \frac{1}{T+1}\left|\sum_{t=0}^T {}^*u_t - \sum_{t=0}^T {}^*u_t^\pi\right| \simeq 0. \tag{9}$$

The essential idea is that a measure 1 set of the terms in the first sum also appear in the second sum and cancel.

Let T' denote the integer part of $\sqrt{T}$. By the triangle inequality,

$$d \leq \frac{1}{T+1}\sum_{t=T'}^T |{}^*u_t - {}^*u_t^\pi| + \frac{1}{T+1}\sum_{t=0}^{T'-1} |{}^*u_t - {}^*u_t^\pi|. \tag{10}$$

Since $|{}^*u_t - {}^*u_t^\pi| \leq 2\|u\|$, the second term is bounded above by $\frac{\sqrt{T}}{T+1}2\|u\|$, and this is infinitesimal. Because π is of order 1, the proportion of the terms in the first term that fail to cancel each other out is also infinitesimal. In more detail, a sequence $x_t \to 0$ in $\mathbb{R}$ iff for all infinite t, ${}^*x_t \simeq 0$. Thus, a permutation π is order 1 iff for all infinite integers t, $\frac{|{}^*\pi(t)-t|}{t} \simeq 0$. Therefore, letting T' denote the integer part of $\sqrt{T}$, $\delta := \max_{t\in\{T',T\}} \frac{|{}^*\pi(t)-t|}{t} \simeq 0$. Since $\pi(T') \in [(1-\delta)T', (1+\delta)T']$ and $\pi(T) \in [(1-\delta)T, (1+\delta)T]$, at most $2\delta T$ of the terms fail to cancel each other out, and $2\delta\frac{T}{T+1} \simeq 0$.

Perfectly Pareto. Suppose first that $\underline{\ell}(B) = \liminf_T \frac{1}{T+1}\sum_{t=0}^T 1_B(t) > 0$. For every infinite T, $\frac{1}{T+1}\sum_{t=0}^T 1_B(t) = \langle {}^*1_B, \eta_T \rangle \geq \underline{\ell}(B)$. Because the extreme points of $\mathbb{V}$ have representations as ${}^\circ\langle \cdot, \eta_T \rangle$ for T infinite, for any $L \in \mathbb{V}$, $L(1_B) \geq \underline{\ell}(B)$. For any $u \in \mathbf{W}$, $r > 0$ and $L \in \mathbf{D}S(u)$, we have $S(u) \leq S(u + r1_B) + L(-r1_B)$. By linearity, $L(-r1_B) = -rL(1_B)$ so that $S(u) + (> 0) \leq S(u + r1_B)$, that is, $S(u) < S(u + r1_B)$.

Now suppose that $\overline{\ell}(B) \limsup_T \frac{1}{T+1}\sum_{t=0}^T 1_B(t) = 0$. For any infinite T, $\frac{1}{T+1}\sum_{t=0}^T 1_B(t) = \langle {}^*1_B, \eta_T \rangle = 0$, hence for any extreme $L_{ext} \in \mathbb{V}$, $L(1_B) = 0$, which implies that the same is true for all $L \in \mathbb{V}$. For $L' \in \mathbf{D}S(u)$ and $L'' \in \mathbf{D}S(u + r1_B)$, we have

$$S(u + r1_B) \leq S(u) + L'(r1_B) \text{ and } S(u) \leq S(u + r1_B) + L''(-r1_B). \tag{11}$$

Since $L'(r1_B) = L''(r1_B) = 0$, we have $S(u + r1_B) \leq S(u) \leq S(u + r1_B)$. □

Proof of Lemma 1.

Proof. $u \in \text{Erg}$ iff the sequence $\langle u, \eta_T \rangle$, $T \in \mathbb{N}_0$, converges to $\text{Ave}(u)$, and this in turn holds iff for any infinite $T \in {}^*\mathbb{N}_0$, $\langle {}^*u, \eta_T \rangle \simeq \text{Ave}(u)$. If $u \in \text{Erg}$, then every extreme $L_{ext} \in \mathbb{V}$ satisfies $L(u) = \text{Ave}(u)$, hence every $L \in \mathbb{V}$ has this property. If $u \notin \text{Erg}$, then there exists infinite T, T' such that $^\circ\langle {}^*u, \eta_T \rangle \neq {}^\circ\langle {}^*u, \eta_T \rangle$, and both of these functionals belong to $\mathbb{V}$. □

Proof of Theorem B. Let S be a $\mathbb{V}$-concave societal welfare function. Every $L' \in \mathbf{D}S(u)$ is of the form γL for some $L \in \mathbb{V}$. Because $\gamma L \in \mathbf{D}S(u)$ is a tangent to S at u, there exists a constant $\kappa \in \mathbb{R}$ such that $\kappa + \gamma L(v) \geq S(v)$ for all $v \in \mathbf{W}$ with equality for $v = u$. Let KG denote the set of (κ, γ) such that $\kappa + \gamma L(v) \geq S(v)$ for all $v \in \mathbf{W}$ and $L \in \mathbf{D}S(u)$ for some u. From Lemma 1 for all $u \in \text{Erg}$ and all $L \in \mathbb{V}$, $L(u) = \text{Ave}(u)$. For $r \geq 0$, define $\varphi(r) = \min\{\kappa + \gamma r : (\kappa, \gamma) \in KG\}$ and note that for $u \in \text{Erg} \cap \mathbf{W}$, $S(u) = \varphi(\text{Ave}(u))$. The function $\varphi(\cdot)$ is concave because it is the lower envelope of a collection of affine functions, and it is strictly increasing because, by Theorem A, $S(\cdot)$ is Pareto. □

References

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

Geir B. Asheim and Stéphane Zuber. Probability adjusted rank-discounted utilitarianism. 2014.

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

Stefan Banach. *Théorie des opérations linéaires.* Chelsea Publishing Co., New York, second edition, 1978.

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

Kaushik Basu and Tapan Mitra. Utilitarianism for infinite utility streams: A new welfare criterion and its axiomatic characterization. *Journal of Economic Theory*, 133(1):350–373, 2007.

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

David Deutsch. *The beginning of infinity: explantions that transform the world.* Viking, New York, NY, 2011.

Charles Figuières and Mabel Tidball. Sustainable exploitation of a natural resource: a satisfying use of chichilniskys criterion. *Economic Theory*, 49(2):243–265, 2012.

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

Pope Francis. Laudato si: On care for our common home. *Encylical Letter. Retrieved August*, 10:2015, 2015.

Geoffrey Heal. Discounting and climate change; an editorial comment. *Climatic Change*, 37(2):335–343, 1997.

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. ISBN 0-12-362440-1.

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.

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.

George G. Lorentz. A contribution to the theory of divergent sequences. *Acta mathematica*, 80(1):167–190, 1948.

Massimo Marinacci. An axiomatic approach to complete patience and time invariance. *journal of economic theory*, 83(1):105–144, 1998.

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.

Charles Perrings and William Brock. Irreversibility in economics. *Annual Review of Resource Economics*, 1(1):219–238, 2009.

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.

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.

Stéphane Zuber and Geir B. Asheim. Justifying social discounting: the rank-discounted utilitarian approach. *Journal of Economic Theory*, 147(4):1572–1601, 2012.