Collective Dynamic Choice:
The Necessity of Time Inconsistency[†]

By MATTHEW O. JACKSON AND LEEAT YARIV*

We study collective decisions by time-discounting individuals choosing a common consumption stream. We show that with any heterogeneity in time preferences, every Pareto efficient and non-dictatorial method of aggregating utility functions must be time-inconsistent. We also show that decisions made via non-dictatorial voting methods are intransitive. (*JEL* D71, D72, D91)

We make collective decisions of many types in our day-to-day lives: from household savings and consumption decisions, to choosing projects and allocating costs in committees, to deciding on taxes and spending in legislatures. A central challenge in making such decisions is that individuals in the collective may have heterogeneous preferences, most especially over the timing of consumption or expenditure of resources.

This heterogeneity is quite evident in many applications. For instance, in the US Congress there is great variance in the tenure of representatives. At the start of the 113th Congress, in the House or Senate, respectively, 17 percent or 14 percent had no previous experience, 37 percent or 40 percent had between 0–8 years of experience, 26 percent (in both chambers) had between 8–16 years of experience, and 20 percent (in both chambers) had over 16 years of service under their belt (see Glassman and Wilhelm 2013). To the extent that tenure in a chamber is a proxy for a representative's remaining time in service and, in particular, the likelihood of running for reelection, it is reasonable to assume that different tenures would be associated with different time preferences when making long-term decisions. As a different example, in (heterosexual) households, it is well documented that women have significantly higher life expectancies than men in most parts of the world.[1]

Having different survival probabilities per period translates into husbands and wives discounting the value of savings at substantially different rates. Indeed, Schaner (2015) finds evidence of significant differences in time preferences inside households in Western Kenya. More generally, the heterogeneity in discount factors may be driven by attributes other than age or tenure, such as background, personality, health conditions, etc.

In addition, there is another important application of the analysis of collective decisions made by agents discounting the future differently. There is growing evidence from neuroscientific investigations that suggests that individual brains engage in parallel processing and aggregation of motives. In particular, there is evidence that different parts of a brain respond differently to timed rewards.[2] From that perspective, we may gain insight into *individual* decision making by examining aggregation of heterogeneous preferences corresponding to different "motives" in one's own mind.

Despite the heterogeneity of decision-making bodies, economic models are usually built assuming that organizations are representable by "rational" agents who have time-consistent and transitive preferences. For example, the neoclassical firm is assumed to maximize a time discounted stream of expected profits and representative consumers are assumed to maximize time discounted streams of consumption. This embodies an important time consistency: if an agent prefers a consumption of $9 eleven periods from now to a consumption of $8 ten periods from now, then the agent should also prefer $9 in one period to the consumption of $8 immediately. This means that the agent will not reverse a choice as time passes. This is an important form of "rationality" that implies, in particular, that an agent will not require commitment devices to settle disagreements between different "selves."

One of our main results shows, however, that very natural procedures for making collective dynamic decisions are inherently *time-inconsistent*, even if the underlying individuals are perfectly time-consistent. Thus, presuming representative time-consistent preference is with tremendous loss of generality.

Specifically, consider any group of agents who are each individually time-consistent and who are collectively evaluating a stream of common consumption. Individuals have possibly heterogeneous discount factors and instantaneous utility functions that are well-behaved. Preferences are aggregated to a collective utility function that satisfies a minimal restriction: that it respect unanimity, or Pareto efficiency. That is, whenever everyone agrees that one consumption stream is superior to another, the collective utility function exhibits the same preference.[3] We prove that any time-consistent aggregation rule that satisfies this minimal restriction must be dictatorial: it must track the preferences of only one of the

---

Nations Statistics Division Social Indicators, updated in December 2010. In fact, compounded with the age difference at marriage, Browning (2000)'s estimates suggest that in Canada the wife of a 65-year-old man would have, on average, an approximately 50 percent longer expected survival horizon than her husband (see also Drefahl 2010).

[2] See, for example, McClure et al. (2004, 2007); Glimcher and Rustichini (2004); and Hare, Camerer, and Rangel (2009).

[3] For example, if the group evaluates consumption streams as the sum of participants' utility functions (i.e., a utilitarian collective utility function), then this is satisfied, as it would if the group evaluates alternatives according to the minimum utility or maximum utility.

group's members. In other words, an aggregation rule that is non-dictatorial and respects unanimous choices must be time-*in*consistent.

Although we first state this result for common consumption streams, we show that it easily extends to quite general formats of consumption—including parts that are private, common, or common to some subgroups, and so on; thus allowing outcomes to be the consequence of, say, group bargaining. Furthermore, since the result regards observed ultimate choices of consumption streams, it admits situations in which agents can either commit or not to their consumption over time. Regardless of the underlying procedure by which choices are taken, whenever ultimate consumption can be evaluated by a collective utility function (or by a planner armed with some utility function), our first main result highlights the tension between Pareto efficiency, time consistency, and engaging more than one individual in decisions. From a policy perspective, non-dictatorial collective choices that are rationalizable by some collective utility function either necessitate commitment devices or involve reversals over time. Furthermore, when conducting a welfare analysis of intertemporal policies, the underlying distribution of preferences should be accounted for: estimating "representative preferences" (say, for a household or a firm) may not properly reflect the preferences of the individuals.

A natural way to weaken the requirement of time consistency is to consider aggregation rules that can be represented as discounted utility functions with *time-varying* discount factors. We show that whenever such rules respect unanimity, they correspond to a weighted sum of the individual agents' utility functions. Therefore, combined with results in Jackson and Yariv (2014), whenever non-dictatorial, these rules exhibit a specific form of time inconsistency: a present-bias, one of the leading time inconsistencies identified from lab and field data.[4]

Our last results examine other methods of aggregating preferences: via some sort of vote by agents. These aggregation methods hinge on ordinal comparisons between any two consumption streams rather than on their cardinal evaluation, and may be particularly relevant for applications having to do with political governance. For example, majority voting is widely used in economic and political contexts, and is qualitatively different from aggregation procedures relying on agents' cardinal utility functions. One may conjecture that there would be a "median voter," who might effectively appear as a dictator and determine all decisions because he or she would be pivotal. A median voter would be reassuring, since although being a dictator, the voter would be "representative" of the population. However, as we show, this is not the case. If a voting rule satisfies a "local non-dictatorship" condition (chooses one consumption stream over another whenever at least $n - 1$ agents prefer the first to the second), then it must exhibit voting cycles and hence be intransitive.

To summarize, our two main results are:

- If individual preferences are aggregated via a collective utility function that is non-dictatorial and respects unanimity, then that collective utility function must be time-inconsistent.

---

[4] Present-biased preferences correspond to more impatience closer to the present than with respect to delayed consumption, see the literature review below.

- If preferences are aggregated via any voting rule that is locally non-dictatorial, such as majority voting, then the resulting social welfare ordering must exhibit cycles (intransitivities), unless the set of consumption streams is severely restricted.

Thus, via either a cardinal or ordinal aggregation method, some collective "irrationality" is necessary: either we must sacrifice time consistency to obtain a cardinal ranking, or we must sacrifice transitivity to obtain an ordinal ranking.

Related Literature: Zuber (2011) is the closest paper to ours in terms of our first result on the necessity of time inconsistency in aggregating time preferences. He shows that a planner can only aggregate preferences in a stationary manner if all agents have the same discount factor. Zuber's result is in a significantly different setting: each agent can have an independent and arbitrary consumption stream, whereas our focus is on joint decisions in which at least some consumption is common, and in fact our results hold when *all* consumption is common.

This is more than a superficial difference. Allowing for separate consumption streams allows for an easy conflict of interest among agents as one can directly construct combinations of streams over which there are voting cycles: one can induce any preference for any given individual by separately adjusting that agent's consumption. This is not possible with common consumption streams. With fully common consumption streams all agents prefer earlier to later consumption, and preferences over two timed levels of consumption can be ordered in terms of agents' discount factors. Thus, to find conflicts among the agents is more subtle. Hence, our proof technique is different and novel: adapting results of Koopmans (1960) on individual decision making, and showing specific features regarding the richness of common consumption streams.[5] Commonality of consumption is important for practically all of the applications mentioned above: households often decide how much (or what) to consume out of some combined budget(s), committees and legislatures often make decisions over whether to spend today and tax tomorrow, or vice versa, etc.

It is important to note that it has long been recognized that aggregating preferences with heterogeneous discount factors introduces challenges. Marglin (1963) and Feldstein (1964) discovered the difficulty of deriving an appropriate aggregate time-independent discount rate for a planner facing a society of heterogeneous agents.[6,7] For example, Gollier and Zeckhauser (2005) show that a representative agent has a time-varying discount factor if there is sufficient uncertainty and

---

[5] An analog of Zuber's theorem can be deduced from our Theorem 2, since a common consumption stream can be nested in his domain, but the reverse is not true.

[6] More recent work in the context of household decisions, Bernheim (1999); Browning (2000); Mazzocco (2007); Hertzberg (2010); Abdellaoui, l'Haridon, and Paraschiv (2010); and Schaner (forthcoming), among others, has considered the implications of preference heterogeneity on intertemporal consumption decisions under particular aggregation protocols. For instance, it illustrates that households may have hyperbolic preferences (in Hertzberg, 2010) and that commitment devices play a role in determining consumption patterns (in Mazzocco 2007, and Schaner forthcoming).

[7] Work by Weitzman (2001); Caplin and Leahy (2004); Blackorby, Bossert, and Donaldson (2005); Gollier and Zeckhauser (2005); and Green and Hojman (2009) examines variations on such planner aggregation issues in more detail. See also Jamison and Jamison (2007), who distinguish between the speed and amount of discounting, and discuss some virtues of hyperbolic discounting. Farmer and Geanakoplos (2009) consider uncertainty as the

heterogeneity in the environment. Jackson and Yariv (2014) show that utilitarian aggregation of common consumption streams must lead to a present-bias; a result that shows that a specific form of time inconsistency is exhibited by a particular, but common, form of collective utility function. They also present lab experimental evidence that subjects exhibit time inconsistencies when acting as social planners.

The novelty in our results on time inconsistency comes from showing that whenever agents differ in their discount factors, even in deterministic settings in which agents have identical instantaneous utility functions, time inconsistency must result from *every* non-dictatorial aggregation method respecting unanimity. Thus, our results illustrate that the phenomena identified in the literature are not unique to specific planner or representative agent formulations, but hold generally, and emerge even when there is only heterogeneity in discount factors.

The problems we study are also related to those pertaining to the aggregation of subjective preferences over lotteries: in a sense, a time-separable utility function (in particular, a time-consistent one) is analogous to a subjective expected utility function: time periods can be interpreted as states and the discount factors as a probability measure over those states (with an appropriate normalization). In that regard, our results are conceptually connected to work by Mongin (1995, 1998), who showed that it is impossible to aggregate heterogeneous subjective probabilities into a common representative probability. Nonetheless, the domains in which the problems are embedded are very different. In particular, time discounting places a very specific restriction on the probabilities across states—they have to be proportional to: 1, δ, δ^2, etc., whereas Mongin's approach makes use of the fact that agents can have arbitrarily different orderings across states to derive conflicts in aggregation.[8] Consequently, the results cannot be mapped into each other, and the techniques we use differ substantially. Even more importantly, the applications and implications are quite different.

With regard to our voting results, formal difficulties in aggregating preferences have been evident since Condorcet's (1785) description of the voting paradox. These difficulties were crystallized via Arrow's Theorem (1950, 1963). Later, obstacles to aggregating convex preference relations over multidimensional ("spatial") alternatives were pointed out by Plott (1967) and McKelvey (1976, 1979). Closest to our theorem on voting rules is that of Boylan and McKelvey (1995), who noted the intransitivities that may arise when majority voting is used in the context of consumption and saving problems and voters have varying time preferences. The current paper contributes to this strand of literature in that we show that, in the context of temporal decisions, aggregation is problematic even with a great deal of structure on individual preferences and the requirement that all agents consume the same stream of consumption. Our results concerning the intransitivity of voting rules apply to a wide class of procedures, containing majority rule as a special case. Moreover, beyond showing that there are issues with voting cycles, we also examine

---

foundations for hyperbolic discounting. Galperti and Stulovici (2014) provide an axiomatic foundation for hyperbolic discounting. Lengwiler (2005) considers the effects of time preference heterogeneity on asset prices.

[8] He also requires some additional continuity conditions to derive his results.

collective utility functions and show the general impossibility of time consistency, which is quite different from any of the above-mentioned papers.

Finally, our analysis of particular classes of aggregation methods (namely, welfare maximization or binary voting rules) has important implications for understanding observed anomalies in individual decision making. The literature documenting time inconsistencies and intransitivities is too vast to cover here.[9] The idea that individuals might be usefully thought of as having some internally heterogeneous preferences appears in a variety of places, possibly the most related of which is a recent paper by Green and Hojman (2009), which provides a general revealed-preference welfare bound analysis allowing for such possibilities.[10] A contribution of the current paper is the insight that viewing individuals as nondegenerate collectives leads *necessarily* to behaviors exhibiting time inconsistency and/or intranstitivities in ways that are in line with empirical observations on a variety of dimensions.[11]

I. The Setting

A. *Agents and Consumption Streams*

A set of agents, $\mathcal{N} = \{1, \ldots, n\}$, must make a collective decision over streams of common consumption.

The agents are infinitely lived and consume in discrete periods $t \in \{1, 2, \ldots\}$. All of our results also hold for finite horizons.[12]

Each period's common consumption c_t lies in $[0, 1]$ and a stream of consumption is denoted by $C = (c_1, c_2, \ldots) \in [0, 1]^\infty$.[13]

For expositional simplicity, we start by describing our results for one-dimensional consumption. All agents have utility functions that are functions of the (same) stream. In terms of interpretation, it is not critical that consumption be common

---

[9] Some seminal references include Herrnstein (1961) and Thaler (1981) for time inconsistencies and Tversky (1969) for intransitivies. These phenomena seem fundamental in that they are observed in species other than humans as well. Time inconsistencies have been documented in rats and pigeons (see Ainslie 1975 and Rachlin 2004). Intransitivies have been observed in bees, as in Shafir (1994) and jays, as in Waite (2001), where the jays exhibit intransitivies in settings very similar to the ones analyzed here (with distance substituting for time). For an overview, see, e.g., Frederick, Loewenstein, and O'Donoghue (2002).

[10] Some notable models with multiple personalities, preferences, or motives of agents include, among others: Thaler and Shefrin (1981); O'Donoghue and Rabin (1999); Bernheim and Rangel (2004); Amador, Werning, and Angeletos (2006); Benabou and Tirole (2005); Brocas and Carrillo (2008); Fudenberg and Levine (2006); Evreny and Ok (2011); Ambrus and Rozen (2009); and Cherepanov, Feddersen, and Sandroni (2009). For a simple non-unitary discount model in a two-motive, CRRA utility setting, see Hori and Futagami (2012). In those settings, various forms of differences in preferences across time or state lead to a conflict between, e.g., current and future selves. This is in contrast to the current setting in which multiple individuals or selves collectively make a choice.

[11] The experimental and empirical evidence regarding whether individuals are themselves time-consistent is mixed. On the one hand, when faced with very simple decisions in the lab, many individual decision makers appear to be time-consistent (see Andreoni and Sprenger 2012a, b). On the other hand, time-inconsistent models of decision making appear to explain well a variety of real-world phenomena, ranging from saving behavior (Laibson 1997 and Beshears et al. 2008) to physical exercise (della Vigna and Malmendier 2006). In conjunction, there is some recent evidence suggesting that individuals may discount different dimensions (such as, say, money and health) using different discount factors (see Chapman and Elstein 1995; Chapman et al. 2001, and references therein).

[12] The Proof of Theorem 1 requires there to be at least five periods, while all other results require a minimum of three periods (which is the minimal number of periods for which time consistency has bite).

[13] The uniform bound on consumption ensures that present values are well-defined (although it is only essential for the results that this be true of several periods, as seen in the proofs). Given the bound, the normalization to $[0, 1]$ is without further loss of generality.

per se, but rather that agents be able to evaluate their personal utility based on a collective decision. This would apply to a variety of examples as discussed above; for example, it would apply to any setting in which some entity (a government, a household, etc.) decides upon the allocation of a budget across time periods and agents can assess their individual resulting utilities conditional on the given budget in each period. So, it need not be the case that the budget is spent on public goods or common consumption. Of course, in the interpretation of multiple motives within a single person, the consumption truly is common.

In Section V, we show that our results extend to settings in which different individuals can consume different streams, some of which might be private or public.

B. *Individual Agents*

Each agent's preferences are represented by a time additive discounted utility function. Agent i has a discount factor $\delta_i \in (0, 1)$ and an increasing[14] and twice continuously differentiable instantaneous utility function $u_i : [0, 1] \rightarrow \mathbb{R}$ such that a consumption stream $\mathbf{C} = (c_1, c_2, \dots) \in [0, 1]^\infty$ is evaluated as

$$(1) \qquad\qquad U_i(\mathbf{C}) = \sum_t \delta_i^{t-1} u_i(c_t).$$

Given that u_i is increasing and twice continuously differentiable, without loss of generality, we normalize all instantaneous utility functions such that $u_i(0) = 0$ and $u_i(1) = 1$. We let the set $\mathcal{U}$ denote the preferences (δ_i, u_i) satisfying the conditions above.

A *society* of n individuals is a profile $(\delta_1, u_1; \dots; \delta_n, u_n) \in \mathcal{U}^n$. We sometimes slightly abuse notation and let U_i denote the corresponding (δ_i, u_i), so that a society can be denoted by $U = (U_1, \dots, U_n) \in \mathcal{U}^n$.

We examine the case in which all agents have "standard" (time discounted, additively separable) utility functions in order to show the difficulties in aggregation even with extremely well-behaved underlying preferences. Of course, allowing time inconsistency and/or intransitivities of these agents would lead to such conclusions in the aggregate a fortiori.[15]

C. *Collective Decisions*

We consider two different ways of aggregating preferences: by a collective utility function and by a collective preference ordering.

Collective Utility Functions.—A *collective utility function* is a function $V : \mathcal{U}^n \times [0, 1]^\infty \rightarrow \mathbb{R}$.

---

[14] Increasing is taken to mean strict.

[15] This specification has become the standard in the literature due to its time consistency properties. Thus, despite the strong restrictions imposed by the additive separability, it is important to understand its properties. Note that any enriched domain of possible consumption streams must face the impossibility results that we obtain here.

A collective utility function provides a "planner's" utility function for a society. Examples include taking a weighted average of the agents' utility functions, $V[U](\mathbf{C}) = \sum_i w_i U_i(\mathbf{C})$, as in a utilitarian approach, or considering the minimum of agents' utilities, $V[U](\mathbf{C}) = \min_i U_i(\mathbf{C})$, as in a Rawlsian approach.

We sometimes abuse notation and, for a given society $U = (\delta_1, u_1; \dots. \delta_n, u_n)$, write $V(\mathbf{C})$ instead of $V[U](\mathbf{C})$ to denote the collective utility for stream $\mathbf{C}$, omitting the explicit dependence on U when it is fixed.

Social Welfare Orderings and Voting.—The collective decision making of a society might not be representable by a collective utility function. For example, when collective decisions are taken by a vote, they may result in choices between any pair of alternatives, which are not rationalizable by any collective utility function (particularly if choices turn out to be intransitive).

As such, it is also useful to consider a social welfare ordering as representing collective behavior. This is a binary relation that represents the decision society would make between any given pair of consumption streams.

We denote the (weak) binary preference relation of a society by $R[U]$ for any $U = (\delta_1, u_1, \dots, \delta_n, u_n) \in \mathcal{U}^n$. In some cases the social welfare orderings are complete and reflexive, but that need not be the case.

The induced strict preference relation $P[U]$ is defined as usual by

$$\mathbf{C} \ P[U] \ \mathbf{C}' \ \text{if}$$

$$\mathbf{C} \ R[U] \ \mathbf{C}' \ \text{and not} \ \mathbf{C}' \ R[U] \ \mathbf{C}.$$

One prominent example of such a preference relation is the case in which $\mathbf{C}P[U]\mathbf{C}'$ if a strict majority of individuals prefer $\mathbf{C}$ to $\mathbf{C}'$, which corresponds to standard majority rule.

Note that any collective utility function induces a social welfare ordering, but clearly not the reverse.

II. Separable Aggregation of Preferences and Present-Bias

Before presenting our general results, we consider the important class of *time-separable* collective utility functions. We show that this class of utility functions exhibits a particular sort of time inconsistency: present-bias, which matches some evidence on behavior. This generalizes a result in Jackson and Yariv (2014) that shows that utilitarian collective utility functions must be present-biased (as defined below).

For any profile $(\delta_1, u; \dots; \delta_n, u) \in \mathcal{U}^n$, a time-separable collective utility function takes the form

$$V[\delta_1, u; \dots; \delta_n, u](\mathbf{C}) = \sum_t \tilde{\delta}_t u(c_t),$$

where u is some collective instantaneous utility function and $(\tilde{\delta}_t)$ are some time dependent weights, which may or may not correspond to exponential discounting. Standard utilitarian aggregation of individual utilities (or one that puts different weights on different individuals) is a special case.

Let $\mathbf{C}[x, t]$ denote a consumption stream with $c_t = x$ and $c_{t'} = 0$ for $t' \neq t$.

Present-Biased Collective Utility Functions: A collective utility function is *present-biased* if:

- $V(\mathbf{C}[x, t]) \leq V(\mathbf{C}[y, t + k])$ *implies* $V(\mathbf{C}[x, t + 1]) \leq V(\mathbf{C}[y, t + k + 1])$ *for any x, y, and* $t \geq 0, k \geq 1$, *and*
- *For any* $t \geq 1$ *and* $k \geq 1$, *there exist* x *and* y *such that* $V(\mathbf{C}[x, 1]) > V(\mathbf{C}[y, k + 1])$ while $V(\mathbf{C}[x, t + 1]) < V(\mathbf{C}[y, t + k + 1])$.

Present-bias implies that immediate consumption entails more *im*patience than delayed consumption. The first part of the definition implies that if consumption of x at time t is inferior to consumption of y further delayed by k periods, the same should be true when moved to time $t + 1$ (i.e., after one period passes)—at least as much patience is exhibited with respect to delayed consumption as with respect to more immediate consumption. The second part of the definition requires that for some conceivable consumption streams reversals occur: at present, the immediate consumption is preferred, while delayed consumption is preferred when all choices are deferred (this description corresponds to that of Strotz 1955 and to the impulsiveness notion described in Ainslie 1975).

We now show that if one simply imposes a separability condition as well as a weak Pareto efficiency condition (unanimity), then a present-bias is implied. Indeed, much of the empirical evidence suggesting time-inconsistent behavior (e.g., Strotz 1955; Laibson 1997; della Vigna and Malmendier 2006, and references therein) has maintained the separable structure of preferences but found a time-dependent discount factor (hyperbolic, or quasi-hyperbolic). As it turns out, whenever collective preferences take such a form, but still satisfy unanimity, they *must be* equivalent to maximizing a weighted sum of agents' utility functions; which, in turn, generates a present bias.

For simplicity, the following result focuses on a case in which all agents, as well as the collective, share the same instantaneous utility function, thus highlighting the role of heterogeneity of time discounting.

As mentioned, a condition on a collective utility function that is useful in what follows is that of unanimity. It requires that if all agents prefer one stream to another, then the collective utility function should reflect that preference, a Pareto efficiency requirement.

Unanimity: *A collective utility function V satisfies unanimity if*:

- $V[U](\mathbf{C}) \geq V[U](\widehat{\mathbf{C}})$ whenever $\sum_t \delta_i^{t-1} u_i(c_t) \geq \sum_t \delta_i^{t-1} u_i(\widehat{c}_t)$ for all i, and
- the first inequality is strict whenever the second is strict for all i.

PROPOSITION 1: *For any profile* $(\delta_1, u; \ldots; \delta_n, u) \in \mathcal{U}^n$ *such that* $u_i = u$ *for all* i, *and* $\delta_k \neq \delta_j$ *for some* k, j, *a collective utility function of the form*

$$(2) \qquad V[\delta_1, u; \ldots; \delta_n, u](\mathbf{C}) = \sum_t \tilde{\delta}_t u(c_t)$$

satisfies unanimity if and only if there exist nonnegative weights $w_j \geq 0$ (*positive for at least one j*) *such that* $\tilde{\delta}_t = \sum_i w_i \delta_i^{t-1}$ *for each t, and so*

$$V[\delta_1, u; \ldots; \delta_n, u](\mathbf{C}) = \sum_i w_i U_i(\mathbf{C}).$$

Also, such a V is necessarily either dictatorial[16] *or present-biased.*

The proof works by showing that in order to respect unanimity, the collective discount factor must be a weighted sum of the discount factors of the agents, and so must correspond to a weighted utilitarian collective utility function. The last conclusion of the proposition then follows from Jackson and Yariv (2014), who show that a utilitarian collective utility function must either be dictatorial or present-biased.

The proposition encompasses many of the formulations of time-inconsistent preferences (e.g., hyperbolic, under which $\tilde{\delta}_t = \frac{1}{a + bt}$, or quasi-hyperbolic, corresponding to $\tilde{\delta}_1 = 1$ and $\tilde{\delta}_t = \beta \delta^{t-1}$ for all $t > 1$, etc.). As long as behavior has a separable structure and satisfies unanimity, the proposition shows that a present-bias is to be expected.

III. General Aggregation of Preferences

Although time-separable aggregation of utility functions exhibits a particular form of time inconsistency, it is conceivable that there are other forms of aggregation that are time-consistent. For example, would a Rawlsian method that maximizes the minimum utility be time-consistent? Would other collective utility functions that are based on order statistics or incorporate inequality aversion be time-consistent?

In this section, we show that there do not exist any non-dictatorial collective utility functions that are time-consistent in societies where there is some heterogeneity in agents' discount rates.

Two pieces of notation are useful.

Given $\mathbf{C} \in [0, 1]^\infty$ and $c_1 \in [0, 1]$, let $(c_1, \mathbf{C})$ denote the consumption stream $\mathbf{C}'$ such that $\mathbf{C}'_1 = c_1$ and $\mathbf{C}'_t = \mathbf{C}_{t-1}$ for $t > 1$.

In addition, given $\mathbf{C}, \mathbf{C}' \in [0, 1]^\infty$, let $(\mathbf{C}|_t \mathbf{C}')$ denote the stream that corresponds to $\mathbf{C}_t$ for $\tau \leq t$, and $\mathbf{C}'_t$ thereafter: $(\mathbf{C}|_t \mathbf{C}')_\tau = \mathbf{C}_\tau$ for $\tau \leq t$ and $(\mathbf{C}|_t \mathbf{C}')_\tau = \mathbf{C}'_\tau$ for $\tau > t$.

---

[16] That is, V is an affine transformation of a collective utility function for which $w_i > 0$ for just one i.

Time Consistency: *The utility function V is time-consistent if, for any society* $U = (\delta_1, u_1; \ldots \delta_n, u_n)$, *all streams* $\mathbf{C}, \bar{\mathbf{C}}, \widehat{\mathbf{C}}, \tilde{\mathbf{C}}$, *and times* $0 \le t < t' \le \infty$:[17]

- $V(\mathbf{C}) > V(\bar{\mathbf{C}})$ *if and only if* $V(c_1, \mathbf{C}) > V(c_1, \bar{\mathbf{C}})$ *for any* $c_1 \in [0, 1]$,
- $V(\mathbf{C}|_t\widehat{\mathbf{C}}|_{t'}\mathbf{C}) > V(\bar{\mathbf{C}}|_t\widehat{\mathbf{C}}|_{t'}\bar{\mathbf{C}})$ *if and only if* $V(\mathbf{C}|_t\tilde{\mathbf{C}}|_{t'}\mathbf{C}) > V(\bar{\mathbf{C}}|_t\tilde{\mathbf{C}}|_{t'}\bar{\mathbf{C}})$.

Time consistency essentially imposes two conditions: stationarity, in the sense that rankings of consumption streams do not depend on when they occur; and independence, in the sense that rankings of consumption streams do not depend on periods in which consumption levels coincide across the two consumption streams.[18] The first condition already embodies much of the flavor of the second condition. The fact that the ranking does not change when some consumption is placed in front of the sequence means that it is insensitive to what is placed in the first period (as long as it is the same in both streams). Indeed, using the first condition recursively t' times implies that $V(\widehat{\mathbf{C}}|_{t'}\mathbf{C}) > V(\widehat{\mathbf{C}}|_{t'}\bar{\mathbf{C}})$ if and only if $V(\tilde{\mathbf{C}}|_{t'}\mathbf{C}) > V(\tilde{\mathbf{C}}|_{t'}\bar{\mathbf{C}})$, which is much of the essence of the second condition.[19]

The following background theorem adapts a powerful result from Koopmans (1960), which implies that whenever V is sufficiently well-behaved, time consistency is tantamount to the maximization of a standard time-separable and discounted utility function.[20]

THEOREM 1: [*Koopmans* (*1960*)] *A collective utility function V is continuous, time-consistent, and satisfies unanimity only if for each* $U \in \mathcal{U}^n$, *there exists* $\delta \in (0, 1)$ *and a continuous and increasing* $u : [0, 1] \to \mathbb{R}$ *such that the preferences represented by* $V[U]$ *coincide with those represented by*

$$\sum_t \delta^{t-1} u(c_t).$$

The details behind the adaptation of Koopmans' (1960) results to our setting appear in the Appendix.[21] We remark that a variation on this result could be used

---

[17] For $t = 0$, $(\mathbf{C}|_t\widehat{\mathbf{C}}|_{t'}\mathbf{C})$ becomes $(\widehat{\mathbf{C}}|_{t'}\mathbf{C})$; and for $t' = \infty$, $(\mathbf{C}|_t\widehat{\mathbf{C}}|_{t'}\mathbf{C})$ becomes $(\mathbf{C}|_t\widehat{\mathbf{C}})$.

[18] There is a large literature that interprets time consistency in terms of behavioral plans (see, e.g., Kydland and Prescott 1977). This approach views an agent as consistent whenever plans of action are not overturned over time. Whenever consumption streams are evaluated in the same way in each period (so that agents do not have dated utility functions), the concepts are similar.

[19] Nonetheless, the conditions do not imply one another formally. For example, consider

$$V(c_1, c_2, \ldots) = \sum_{t=1}^{\infty} e^{-\sum_{k=1}^{t-1} c_k} c_t = c_1 + e^{-c_1}c_2 + e^{-c_1-c_2}c_3 + \cdots$$

This utility satisfies the first condition, but not the second. On the other hand, consider $V(c_1, c_2, \ldots) = c_1$. It automatically satisfies the second condition, but not the first.

[20] We follow Koopmans (1960) and define continuity and differentiability using the sup metric $d(\mathbf{C}, \bar{\mathbf{C}}) = \sup_t |c_t - \bar{c}_t|$.

[21] A partial converse holds: $\sum_t \delta^{t-1} u(c_t)$ is continuous and time-consistent, but not necessarily unanimous unless the collective discount factor δ and collective instantaneous utility u correspond to δ_i, u_i for some agent i (that is, $\delta = \delta_i$ and u coincides with u_i up to an affine transformation).

to state our assumption on individuals' preferences as continuity, monotonic-ity, and time consistency (instead of explicitly assuming time-discounted utility functions).

Note that the theorem is an ordinal claim as V could be any strictly monotonic transformation of $\sum_t \delta^{t-1} u(c_t)$ and still satisfy the axioms.

We now state our first main result. If there is any heterogeneity in temporal pref-erences by way of differing discount factors, then the only well-behaved collective utility functions that are time-consistent and respect unanimity are dictatorial: they ignore the preferences of all but one agent (or a group of agents who share the same exact preferences).

THEOREM 2: *Consider a collective utility function V that is continu-ous, time-consistent, and satisfies unanimity so that by Theorem 1, for any $U = (\delta_1, u_1; \ldots; \delta_n, u_n) \in \mathcal{U}^n, V[U]$ represents the same preferences as $\sum_t \delta^{t-1} u(c_t)$ for some $\delta \in (0, 1)$ and a continuous and increasing $u : [0, 1] \to \mathbb{R}$. If u is twice continuously differentiable, then there exists a single $i \in \mathcal{N}$ such that $\delta = \delta_i$. Moreover, if $(\delta_1, u_1; \ldots; \delta_n, u_n)$ is such that if $\delta_j = \delta_k$ for some $j \neq k$, then u_j is an affine transformation of u_k, then u is an affine transformation of u_i and $V[U]$ is dictatorial.*

The theorem states that in order to be time-consistent and satisfy unanimity, the collective utility function must be a time-discounted sum of evaluations of the con-sumption stream, where the collective discount factor must be exactly that of some agent i. In fact, in that case, the collective utility function's instantaneous utility function u can only depend on the utility functions of the agents who have the same discount factor as i. Thus, if agents are distinguished by their discount factors, then the collective utility function must be dictatorial. Alternatively, if a collective utility function responds nontrivially to at least two agents with differing time preferences and also respects unanimity, then it must be time-inconsistent.

In view of common impossibility results á la Arrow, we stress the quantifiers of the theorem. In the setting of Theorem 2, for any *fixed* profile of time preferences, unanimity and time consistency imply that only one agent's preferences are paid attention to in determining the collective utility function. Note that this allows dif-ferent preference profiles to involve different dictators. Nonetheless, the important implication is that if more than one agent's preferences are paid attention to in a given scenario, then a society must be time-inconsistent in that scenario.

Instantaneous utility functions can be thought of as indirect utility functions of per period wealth that is then divided into various dimensions of private and public consumption. In fact, we return to a discussion of an extension of Theorem 2 to general multidimensional consumption vectors below. In some settings, consump-tion streams can also be thought of as resulting from bargaining among the indi-viduals comprising the group. In that case, whenever outcomes can be rationalized by a collective utility function, the theorem implies that the function cannot be simultaneously time-consistent, Pareto efficient, and non-dictatorial. In particular, if one is to design Pareto efficient bargaining protocols resulting in non-dictatorial and time-consistent choices, then commitment tools are necessary. The theorem

illustrates that when observing ultimate choices, they necessarily appear either time-inconsistent or dictatorial.

The Proof of Theorem 2 appears in the Appendix, and proceeds as follows. Theorem 1 establishes that an increasing and twice differentiable utility function that is time-consistent must be representable as a time-additive discounted sum of utility functions. There are then two things that remain to be shown: that the collective discount factor coincides with some agent's discount factor, and that the collective instantaneous utility coincides with the instantaneous utility of that agent (up to an affine transformation). To show that the collective discount factor has to match some agent's, we proceed by contradiction. Suppose that the collective discount factor does not correspond to any of the agents' discount factors. We show that this implies a violation of unanimity. This is very easy in some cases, for instance, in the case where the collective discount factor is strictly higher than all of the agents' discount factors, so that it reflects more "patience." We can then construct two consumption streams such that one entails more immediate consumption (and thereby preferred by all agents) and one entails delayed consumption that is higher overall (and thereby ranked higher by the collective utility). An analogous construction can be done if the collective discount factor is lower than all of the agents' discount factors. The difficult case is when the collective discount factor is inbetween the lowest and the highest of the agents' discount factors. The construction then works off the following key insight. Agents who are less patient than the collective would like to have consumption moved forward more than the collective would. Furthermore, they are willing to have some consumption moved from intermediate periods to both earlier and later periods. More patient individuals would like consumption to be moved back in time more than the collective. Moreover, they are willing to have some consumption moved forward as long as enough consumption is also moved to later periods. In the proof, we construct two streams involving consumption in three periods such that one has higher consumption in the first and third periods relative to the other by just the right amounts, so that all agents prefer the former consumption stream, while the collective utility function ranks it lower, in contradiction to unanimity.

A simple example illustrates the essence of how such a construction works. Consider a society of two agents, with $\delta_1 = 0$ and $\delta_2 = 1$, and a collective utility function that uses the average discount factor, $\delta_{avg} = 1/2$.[22] Suppose all agents have linear utility functions and that $\mathbf{C} = (x, x, x, 0, 0, \ldots)$ and $\mathbf{C}' = (x + \varepsilon, x - 6\varepsilon, x + 6\varepsilon, 0, 0, \ldots)$. Here, $U_1(\mathbf{C}) = x < U_1(\mathbf{C}') = x + \varepsilon$ and $U_2(\mathbf{C}) = 3x < U_2(\mathbf{C}') = 3x + \varepsilon$, so that both agents prefer $\mathbf{C}'$ to $\mathbf{C}$. However, $U_{\delta_{avg}}(\mathbf{C}) = 1.75x > U_{\delta_{avg}}(\mathbf{C}') = 1.75x - 0.5\varepsilon$.

The details of the proof provide a general recipe for finding such reversals if the collective discount factor does not match any of the agents' discount factors.[23] We note that this construction requires only three periods.

---

[22] We use extreme values of discount factors for illustrative purposes, but as the proof shows this can be done for any set of discount factors.
[23] Differentiability of the collective utility function allows us to assess the impacts of small changes in consumption streams, as constructed in the example above.

The final step in the proof establishes that the collective (instantaneous) utility function must also match the utility function of the agent whose discount factor it matches (up to an affine transformation). This is done by a similar construction to that above: if not, then one can find a change that appeals to all the more and less patient individuals, as well as to the agent who has the same discount factor as the collective (because of his or her different utility function), which again contradicts unanimity.[24]

The heterogeneity in discount factors is critical to the results, and so it is heterogeneity in time preferences that is the culprit in necessitating time inconsistency. To see this, note that if a society is composed of agents who share the same discount factor, $\delta_1 = \cdots = \delta_n \equiv \delta$, then there are many collective utility functions that are time-consistent and respect unanimity. In that case, for instance, the collective utility function defined by

$$V[\delta_1, u_1; \ldots, \delta_n, u_n](\mathbf{C}) = \sum_t \delta^{t-1} u(c_t), \quad \text{where } u(c_t) \equiv \frac{1}{n} \sum_i u_i(c_t)$$

is non-dictatorial, unanimous, and time-consistent.

We stress that the result requires more than two periods of consumption. With only two periods, all agents agree that more consumption at each date is better; so the only potential disagreement stems from one stream offering more current consumption and less future consumption than another. As an example, if all agents had the same utility function and only differed in terms of their patience, then using the average of the discount factors to discount the second period utility would satisfy unanimity and would induce a valid collective utility function satisfying the time consistency condition (when restricted to two periods).

To summarize, Theorem 2 illustrates the inevitability of time inconsistencies whenever consumption occurs over several periods and the population is heterogeneous in terms of temporal preferences. When time consistency is weakened to allow for discounted utility functions with time-varying discount factors, Proposition 1 implies that a present-bias is inevitable.

IV. Voting over Consumption Streams

Although we have shown that there is no time-consistent and unanimous manner of nontrivial aggregation of heterogeneous time preferences in the form of a collective utility function, we should also ask whether a society might come to collective decisions that are "rational" collectively, without necessarily being representable by a collective utility function.

In this section, we show that making such choices by any of a general class of voting schemes that respect unanimity will be time-consistent, but must be intransitive.

---

[24] This construction is a bit more involved and requires positive consumption in at least five periods. In particular, an analogous claim to that of the theorem would hold for a society of agents contemplating consumption streams over any finite number of at least five periods. It may be possible to lower it to as few as three periods, although the proof's details would necessarily differ.

In particular, we consider another common way by which groups make decisions collectively: by tallying which individuals prefer one option to another and mapping that set into a choice (for instance, by following majority rule or even weighted (non-anonymous) and/or supermajority voting rules).

Hypothetically, voting or some more general form of making collective binary choices might allow a representative or pivotal agent to be naturally determined. For instance, suppose that all the agents in a society have the same utility function u and differ only in their discount factors. If society operated under standard majority rule, would it decide according to the utility corresponding to the median discount factor? After all, when considering societies of voters over unidimensional sets of alternatives, and where voters have single-peaked preferences, the preferences of the median agent coincide with decisions made by simple majority voting. As it turns out, this is not the case when voting is over time streams of consumption. The median discount factor does not represent a society's voting behavior, nor does any particular discount factor. If any specific discount factor represented a society's voting behavior, then the society's voting behavior would have to be transitive. As we show below, for a rather wide class of voting rules, intransitivities are inherent, unless the set of potential consumption streams is severely limited.

Before presenting our next main result, we present an example illustrating the underlying forces that generate cycles in collective decisions.

Example (**Cycles in Collective Decisions**): Suppose a group is composed of 3 individuals sharing the same instantaneous utility function $u_i(c) = c$ for $i = 1, 2, 3$, but having different discount factors: $\delta_1 = 0, \delta_2 = \frac{1}{2}$, and $\delta_3 = 1$. Consider the following three consumption streams:

$$\mathbf{C} = (x, x, x, 0, 0, \ldots),$$

$$\mathbf{C}' = (x + \varepsilon, x - 6\varepsilon, x + 6\varepsilon, 0, 0, \ldots),$$

$$\mathbf{C}'' = (x + 2\varepsilon, x - 6\varepsilon, x + 3\varepsilon, 0, 0, \ldots),$$

for some $\varepsilon > 0$.

The most impatient individual is concerned only with period 1 consumption, so that

$$U_1(\mathbf{C}) = x < U_1(\mathbf{C}') = x + \varepsilon < U_1(\mathbf{C}'') = x + 2\varepsilon.$$

The moderately patient individual is concerned with earlier consumption and its distribution over time and displays preferences:

$$U_2(\mathbf{C}') = 1.75x - 0.5\varepsilon < U_2(\mathbf{C}'') = 1.75x - 0.25\varepsilon < U_2(\mathbf{C}) = 1.75x.$$

The most patient individual is concerned with the overall sum of consumptions and so

$$U_3(\mathbf{C}'') = 3x - \varepsilon < U_3(\mathbf{C}) = 3x < U_3(\mathbf{C}') = 3x + \varepsilon.$$

If these agents vote by majority rule, a cycle would emerge: Individuals 1 and 3 prefer $\mathbf{C}'$ to $\mathbf{C}$, individuals 1 and 2 prefer $\mathbf{C}''$ to $\mathbf{C}'$, and individuals 2 and 3 prefer $\mathbf{C}$ to $\mathbf{C}''$.

The example illustrates three dimensions that individuals may care about that are the basis for the cycle: immediate consumption, overall consumption, and the distribution of consumption across time. The latter dimension is particularly important when instantaneous utility functions are strictly concave, and we next show that, whenever this is the case, the type of disagreements generating cycles in the example are quite general, even when the set of alternatives is very restricted.

A. *Intransitivities of Binary Voting Rules*

We generalize the example above. We define a general class of binary voting rules that depend only on the *set* of agents who prefer either of two alternatives in determining the one chosen. This allows for non-anonymous and non-neutral voting rules, such as weighted voting rules and/or those that favor some alternatives over others. While simple majority rule, which is within this class, may be the most commonly used, this setup allows for a variety of other rules that are used in practice. For instance, it allows for the type of rules used in the executive council of the European Union, where countries have different weights, and approving some proposals (in lieu of maintaining the status quo) requires a total weight exceeding 50 percent of the overall weights cast.[25]

Formally, for any $\mathbf{C}, \mathbf{C}'$, let $p(\mathbf{C}, \mathbf{C}', U)$ denote the set of individuals who strictly prefer $\mathbf{C}$ to $\mathbf{C}'$:

$$p(\mathbf{C}, \mathbf{C}', U) = \{i | U_i(\mathbf{C}) > U_i(\mathbf{C}')\}.$$

Let a *Binary Voting Rule* be a collective social welfare ordering $R[U]$ that is complete[26] and depends only on the information in p. That is, if

$$p(\mathbf{C}, \mathbf{C}', U) = p(\mathbf{C}, \mathbf{C}', U') \quad \text{and} \quad p(\mathbf{C}', \mathbf{C}, U) = p(\mathbf{C}', \mathbf{C}, U')$$

then

$$\mathbf{C}R[U]\mathbf{C}' \quad \text{if and only if} \quad \mathbf{C}R(U')\mathbf{C}'.$$

This condition embodies a variation of Arrow's Independence of Irrelevant Alternatives, as it requires that comparisons only respond to the set of agents preferring one consumption stream to another. It also embodies an ordinality condition that is inherent in Arrow's setting, and related to various versions of "neutrality" conditions appearing in the *single profile* literature (see, e.g., Parks 1976 and Roberts

---

[25] See Barberà and Jackson (2006) for a discussion of the optimality of voting rules other than simple majority, as well as for additional references and background.

[26] So, for every U, C and C', either $CR[U]C'$ or $C'R[U]C$.

1980).[27] It states that the only information that matters in determining whether one consumption stream is preferred to another is information about which agents prefer each of the two alternatives being compared. Clearly, most of the commonly used voting rules satisfy this condition: majority rule, any weighted voting rule, even voting rules that favor certain alternatives (where, say, choosing certain consumption streams requires specific quorums).

The question is then whether allowing for ordinal collective preferences rather than cardinal ones can allow for a society to be represented by a transitive and time-consistent social welfare ordering.

Locally Non-Dictatorial Orderings: *A social welfare ordering R is* locally non-dictatorial *if*

$$\mathbf{C}R[U]\mathbf{C}' \quad whenever \quad |p(\mathbf{C}, \mathbf{C}', U)| \geq n - 1.$$

A social welfare ordering is locally non-dictatorial if whenever at least all but one agent prefer one consumption stream to another, then so does society. Therefore, locally, at any particular choice, there is no single agent who can force society's preference when all others have an opposing preference. Any (weighted) super-majoritarian voting rule short of unanimity is locally non-dictatorial.

Consider $x > 0$ and $\gamma > 0$ and define a set of consumption streams by

$$\mathbf{C}(x, \gamma) = \{\mathbf{C} \mid c_1 + c_2/\gamma + c_3/\gamma^2 = x, c_t = 0 \quad \forall t > 3\}.$$

These are three-period consumption streams in which an initial amount of the consumption good x is to be split over three periods and grows (or depreciates) at a gross rate $\gamma > 0$ between periods, so that a unit stored in one period becomes γ units in the following period. If $\gamma = 1$, then units are exactly stored across periods; if $\gamma < 1$, then there is depreciation; and if $\gamma > 1$, then there is a positive growth or return to investment across periods.

By restricting the set of alternatives to this simple set, we make it more difficult to find voting cycles since the set of admissible consumption streams are quite constrained.

Our main result in this section is that any binary voting rule that is locally non-dictatorial is intransitive in a very strong sense, even when restricting attention to this set of consumption streams over a short horizon.

We say that a strictly concave utility function ($u_i = u$ for all i) is *non-extreme* (relative to (x, γ)) if there is at least one discount factor $\delta \in (0, 1)$ for which there is a maximizing consumption stream $\mathbf{C}^*(u) = \arg\max_{\mathbf{C} \in \mathbf{C}(x, \gamma)} \sum_t \delta^{t-1} u(c_t)$ that is interior.[28] Let $\mathcal{U}^* \subset \mathcal{U}^n$ denote profiles of preferences for which $u_i = u$ for all i, for some non-extreme and strictly concave utility function u.

---

[27] Note that "neutrality" is a bit of a misnomer originating from that literature, since it entails something very different from the usual usage of neutrality that refers to an insensitivity to the labeling of alternatives.

[28] Note that this is unique given the differentiability and strict concavity of the utility function. The interiority is satisfied, for example, if $\delta/\gamma \leq 1$ and $u'(1/2) < u'(0)(\delta/\gamma)^2$ or if $\delta/\gamma \geq 1$ and $u'(1/2)(\delta/\gamma)^2 < u'(0)$.

Recalling that $P[U]$ corresponds to the strict preference induced by $R[U]$, the following holds.

THEOREM 3: *Consider consumption streams in* $\mathbf{C}(x,\gamma)$ *for any* $x > 0$ *and* $\gamma > 0$. *If a binary voting rule R is locally non-dictatorial, then there exists a profile* $U \in \mathcal{U}^*$ *and consumption streams for which R is intransitive. In fact, for any neighborhood* $B \subset \mathbf{C}\ (x,\gamma)$ *of* $\mathbf{C}^*(u)$ *(as defined above), there exists* $\mathbf{C}' \in B$ *and* $\mathbf{C}'' \in B$ *such that* $\mathbf{C}^*(u)R[U]\mathbf{C}'R[U]\mathbf{C}''P[U]\mathbf{C}^*(u)$.

The theorem states that for any utility function such that a most preferred consumption stream for some possible time preference is interior, and any local neighborhood of that consumption stream, one can find a profile of discount factors for the agents such that there is a voting cycle within that neighborhood. Although intransitivities are within a set of alternatives pertaining to consumption smoothing over only three periods, under the theorem's conditions, intransitivities also arise for consumption smoothing problems with longer horizons, or when consumption streams are less restricted.

The conclusions of the theorem depend on the strict concavity of the utility function. If, for example, all agents have the same linear utility function then, under these restrictions on consumption, the agent with the median discount factor's favorite allocation (generally, to have all consumption either at date 1 or 3 depending on γ), is a Condorcet winning alternative and cycles disappear. Indeed, in the example above, the three consumption streams did not correspond to a consumption smoothing problem.

The theorem extends without strict concavity if we loosen the restrictions on the consumption streams. We show this for majority rule, but the proof makes it clear that it extends easily to other voting rules.

Society makes choices using *simple majority rule* if $\mathbf{C}$ is (weakly) preferred to $\widehat{\mathbf{C}}$ whenever at least half the society weakly prefers $\mathbf{C}$ to $\widehat{\mathbf{C}}$:

$$\mathbf{C}R[U]\widehat{\mathbf{C}} \text{ if } \left| \left\{ i : U_i(\mathbf{C}) \geq U_i(\widehat{\mathbf{C}}) \right\} \right| \geq n/2.$$

PROPOSITION 2: *If* $n \geq 3$, $u_i = u$ *for all* i, *where u is continuous and strictly increasing, P is defined by majority rule, and the largest group of agents having identical discount factors is smaller than* $n/2$, *then* $P[U]$ *is intransitive.*

The proof of this proposition illustrates that any agent can be made the "pivotal voter." To glean some intuition for the workings of the proof, let us show how to identify consumption streams $\mathbf{C}$ and $\mathbf{C}'$, such that individuals $\mathcal{N}_1$ prefer consumption stream $\mathbf{C}$ to $\mathbf{C}'$ and individuals $\mathcal{N}_2$ prefer $\mathbf{C}'$ to $\mathbf{C}$. There is a corresponding system of linear inequalities. Namely,

$$\sum_t \delta_i^{t-1}[u(c_t) - u(c_t')] > 0 \text{ for } i \in \mathcal{N}_1$$

$$\sum_t \delta_i^{t-1}[u(c_t) - u(c_t')] < 0 \text{ for } i \in \mathcal{N}_2.$$

Now, if discount factors are all different, then whenever the range of u is sufficiently rich, the linear independence of $\left\{\left(1, \delta_i, \delta_i^2, \ldots\right)\right\}_i$ guarantees a solution.

The proof, much like the intuition, uses the richness of the set of consumption plans (and the resulting richness of the instantaneous utility function's range). Restricting consumption streams to those corresponding to consumption smoothing problems does not avoid intransitivities—as long as utility functions are strictly concave, Theorem 3 guarantees that interior solutions of some agents are associated with intransitivities.

Note that the voting cycles captured in Proposition 2 are driven by the linear independence of the vectors of coefficients of the sequence of discount factors $\left\{\left(1, \delta_i, \delta_i^2, \ldots\right)\right\}_i$. In fact, as long as there is enough dependence between these vectors (relative to the potential instantaneous utility functions and admissible consumption streams), voting does not entail intransitivies.

B. *Well-Ordered Consumption Streams*

One way to introduce such dependence is to require alternatives to be *well-ordered*. Consumption streams $\mathbf{C}$ and $\widehat{\mathbf{C}}$ are well-ordered relative to a society with discount factors $\delta_1, \ldots, \delta_n$ and a utility function $u_i = u$ for all i if $u(c_t) - u(\widehat{c}_t)$ is monotone in t (either nonincreasing, or nondecreasing).

Well-ordering provides a strong linkage between the preferences of individuals. Intuitively, suppose that $\mathbf{C}$ and $\widehat{\mathbf{C}}$ are well-ordered and that, say, $u(c_t) - u(\widehat{c}_t)$ is increasing. Calculating differences in net present values between $\mathbf{C}$ and $\widehat{\mathbf{C}}$ involve sums of the form

$$\sum_t \delta^{t-1}\left(u(c_t) - u(\widehat{c}_t)\right).$$

As we increase δ, more weight is put on elements further in the sequence $\left\{u(c_t) - u(\widehat{c}_t)\right\}_t$ and so whenever an agent with a discount factor of δ evaluates this expression as positive (so that $\mathbf{C}$ is preferred to $\widehat{\mathbf{C}}$) so does any agent with a higher discount factor. In particular, there is a natural ordering of agents according to their discount factors. Consequently, restricting the set of consumption streams so that agents are well-ordered rules out voting cycles.

A detailed analysis of well-ordered consumption streams appears in the supplementary online Appendix.

V. General Consumption Patterns

Up to this point, we considered situations in which each agent's consumption was common and one-dimensional. In some contexts it may be useful for a planner to be able to evaluate consumption streams that involve combinations of public and private consumptions. Similarly, in many applications in which agents use voting mechanisms, alternatives impact several dimensions of consumption. As it turns out, our results extend directly to settings in which consumption is multidimensional.

Specifically, denote a stream of consumption by $\mathbf{C} = (c_1, c_2, \ldots)$, where each $c_t \in [0, 1]^\ell$ for some positive integer ℓ.

A special case is where $\ell = 1$, which reduces to our original formulation. Another special case is where $\ell = n$ and each agent has a private consumption stream. The general case allows for any combination, where some streams are evaluated by some subsets of agents, and others are private. This allows for agents to be bargaining over streams of consumption, as their utility function need only increase in their own consumption. In particular, an agent could be made worse off by an increase in another agent's consumption, or even by an increase in a public element of consumption, if it comes at the expense of a decrease in the agent's private consumption.

A. *Aggregation and Time Inconsistency*

The definition of a collective utility function, time consistency, and unanimity are exactly as before (as none of them depended on the dimensionality of c_t).

We require each u_i to be twice continuously differentiable, nondecreasing overall, and increasing in at least one dimension of c_t.[29]

The following is a variation of Theorem 2 above.[30]

THEOREM 4: *Consider a collective utility function V that is continuous, time-consistent, and satisfies unanimity. Then, for any $U = (\delta_1, u_1; \ldots; \delta_n, u_n)$ as above, $V[U]$ represents the same preferences as $\sum_t \delta^{t-1} u(c_t)$ for some $\delta \in (0, 1)$ and a continuous and increasing $u : [0, 1] \to \Re$. If u is twice continuously differentiable, then at a profile of preferences such that agents all have different discount factors, V is dictatorial.*

The proof of Theorem 4 can be achieved using techniques similar to those used in the proof of Theorem 2 and is therefore omitted.[31]

B. *Voting with General Consumption Patterns*

We now consider majoritarian voting rules when agents' consumption is multidimensional. Let us say that utility functions are *weakly similar* if there exists u such that for every i, there exist nonnegative scalars a_i and b_i such that $u_i(x\mathbf{1}) = a_i + b_i u(x\mathbf{1})$ for all scalars x.

Note that weak similarity allows for different agents to care about different dimensions of the consumption stream, and to embody separate private consumptions. It

---

[29] The utility function u_i is nondecreasing overall if whenever $c \geq c'$ (so that $c_j \geq c'_j$ for every j), $u_i(c) \geq u_i(c')$.

[30] We define continuity and differentiability using the sup metric with the Euclidean norm: $d(C, \bar{C}) = \sup_t \|c_t - \bar{c}_t\|$.

[31] The result of Koopmans (1960) still applies and step 1 of the proof works simply by considering consumption streams where all elements of c_t are identical (effectively corresponding to the common consumption case). Step 2 is then slightly more involved, working with gradients instead of derivatives, but analogous. The detailed proof appears in the supplementary online Appendix.

only imposes restrictions when all elements of the consumption stream are moved together in unison.

PROPOSITION 3: *If* $n \geq 3$ *and* $(u_1; \ldots; u_n)$ *are weakly similar through some u that is continuous and strictly increasing, then if P is defined by majority rule, and* $(\delta_1; \ldots; \delta_n)$ *is such that the largest group of agents having identical discount factors is smaller than a majority, then* $P[\delta_1, u_1; \ldots; \delta_n, u_n]$ *is intransitive.*

We omit the proof of this proposition since it follows directly from the intransitivity on the restricted domain of consumptions in which each dimension has the same consumption level (but consumption may vary across dates). A similar generalization of Theorem 3 also holds in this setting.

VI. Concluding Remarks

The main message of this paper is that the aggregation of heterogeneous preferences over consumption streams in a non-dictatorial manner that respects unanimity is bound to exhibit time inconsistencies or intransitivities.[32] This insight is relevant for decisions that are made by groups of individuals as well as ones made by one person juggling an assortment of temporal motives.

The results are important for policymaking when heterogeneous temporal preferences are present in the population. Marglin (1963) and Feldstein (1964) suggested that choosing a sensible representative agent may involve nonstationary discount rates, and recent work has examined the implications of time inconsistency in the population on optimal policies (see, e.g., Amador, Werning, and Angeletos 2006). The results in this paper indicate that such considerations are unavoidable. Policymakers who trade-off different temporal preferences of individuals in any nontrivial way de facto face a time-inconsistent representative agent. In fact, if policymakers care about some proxy of (utilitarian) efficiency, they face a present-biased representative agent. In addition, the results suggest that even when estimated preferences pertaining to groups (say, households) exhibit time inconsistencies, they may arise from individual preferences, potentially time-consistent, that differ from the collective's. Therefore, welfare maximization requires a careful analysis with the primitive preferences taken into account, and not simply substituted by a nonexistent representative agent.

The results also open the door for considering specific bargaining protocols in groups with heterogeneous time preferences (such as households trading off consumption and savings within a budget constraint, political committees deciding on investments over time while being restricted in resources, etc.). Whenever such protocols allow for outcomes rationalizable through a collective utility function, our results suggest that the function will either be time-inconsistent or engage the

---

[32] Effectively there is a tradeoff. A cardinal method will produce transitive choices since it necessarily induces an ordering. However, it will generate time inconsistencies if non-dictatorial. Instead, an ordinal method derived from some voting procedure will produce time-consistent choices (since the agents are time-consistent and will not change their votes over time). However, it will be prone to intransitivities.

preferences of only one individual. The precise characterization of outcomes generated by such protocols is likely to require new tools and techniques.[33]

<p style="text-align:center">APPENDIX: PROOFS</p>

PROOF OF PROPOSITION 1:

We first show the claim in the first sentence of the proposition. It is clear that if there exist nonnegative weights w_i such that $V = \sum_i w_i U_i$, then unanimity is satisfied. Let us show the converse. We show that there exists a vector $w \in \Delta_{\mathcal{N}} = \{ w | w_i \in [0,1] \ \forall i, \ \sum_i w_i = 1 \}$ such that $\tilde{\delta}_t = \lambda \sum_i w_i \delta_i^{t-1}$ for some $\lambda > 0$. We show that if this is not the case, then we contradict unanimity.

Let

$$\mathcal{X} = \left\{ x \in \mathbb{R}_+^\infty : \exists w \in \Delta_{\mathcal{N}}, \lambda \geq 0 \quad \text{s.t.} \quad x = \lambda \sum_i w_i \big(1, \delta_i, \delta_i^2, \dots \big) \right\},$$

and

$$\mathcal{X}^T = \left\{ x \in \mathbb{R}_+^T : \exists w \in \Delta_{\mathcal{N}}, \lambda \geq 0 \quad \text{s.t.} \quad x = \lambda \sum_i w_i \big(1, \delta_i, \delta_i^2, \dots, \delta_i^T \big) \right\}.$$

Suppose that $(\tilde{\delta}_t)_t \notin \mathcal{X}$. It follows that $(\tilde{\delta}_1, \dots, \tilde{\delta}_T) \notin \mathcal{X}^T$ for some $T < \infty$. Note that $\mathcal{X}^T$ is a closed and convex set containing the 0 vector. By the Separating Hyperplane Theorem, there exists a vector $y \in \Re^T$ such that $y \cdot x \geq 0$ for all $x \in \mathcal{X}^T$ and $y \cdot (\tilde{\delta}_1, \dots, \tilde{\delta}_T) < 0$.

To obtain strict inequalities on the relevant vectors of discount factors, we slightly perturb y. In particular, normalizing y given the finite dimension (by simply dividing all entries by $2\max_i|y_i|$), we can take $y \in [-1/2, 1/2]^T$. For $\varepsilon < 1/2$, let $y(\varepsilon) = (y_1 + \varepsilon, y_2, y_3, \dots, y_T)$, so that $y(\varepsilon) \in [-1, 1]^T$. It follows that for small enough ε,

$$(3) \qquad \sum_{t \leq T} \delta_i^{t-1} y(\varepsilon)_t > 0 \quad \text{for} \quad i = 1, \dots, n,$$

$$\sum_{t \leq T} \tilde{\delta}_t y(\varepsilon)_t < 0.$$

---

[33] Indeed, one difficulty that arises in such settings is that even when considering an underlying problem of wealth distribution at each period, effectively a per period zero sum game, the overall time discounted game is not zero sum whenever individual discount factors are heterogeneous since agents can trade consumption across time (a point noted in the repeated games literature, see Lehrer and Pauzner 1999 and Lehrer and Yariv 1999).

Given that u is increasing and continuous, its range is $[0, 1]$. Therefore, there exist sequences $\mathbf{C} = (c_1, \ldots c_T, 0, 0, \ldots)$ and $\mathbf{C}' = (c'_1, \ldots c'_T, 0, 0, \ldots)$ such that $u(c_t) - u(c'_t) = y(\varepsilon)_t$. For these $\mathbf{C}$ and $\mathbf{C}'$,

$$(4) \qquad \sum_{t \leq T} \delta_i^{t-1}[u(c_t) - u(c'_t)] > 0 \quad \text{for} \quad i = 1, \ldots, n,$$

$$\sum_{t \leq T} \tilde{\delta}_t[u(c_t) - u(c'_t)] < 0.$$

Thus, all individual agents prefer consumption stream $\mathbf{C}$ to $\mathbf{C}'$, while the collective prefers $\mathbf{C}'$ to $\mathbf{C}$, violating unanimity. It follows that our supposition was incorrect and $(\tilde{\delta}_t)_t \in \mathcal{X}$. Note that since $(\tilde{\delta}_t)_t \neq 0$, it follows that $\lambda > 0$.

The second sentence of the proposition follows from Jackson and Yariv (2014), who show that weighted utilitarian collective utility functions that put positive weight on at least two individuals with differing discount factors are present-biased.[34] ∎

PROOF OF THEOREM 1:

We apply a theorem by Koopmans (1960, section 14). First note that by the fact that each u_i is increasing on $[0, 1]$ and V satisfies unanimity, his postulates 2 and 5 are satisfied. Next, his postulate 1 follows from continuity of V under the metric $d(\mathbf{C}, \mathbf{C}') = \sup_t |c_t - c'_t|$. Finally, time consistency implies his postulates 3, 3', and 4. Thus, there exists $0 < \delta < 1$ and a continuous and increasing[35] u, such that $\sum_t \delta^{t-1} u(c_t)$ for all $\mathbf{C}$ represents the same preferences as $V[U]$. ∎

PROOF OF THEOREM 2:

Let $W(\mathbf{C}) = \sum_t \delta^{t-1} u(c_t)$ represent the preferences induced by $V[U]$.

Without loss of generality, normalize u so that $u(0) = 0$ and $u(1) = 1$. Since the same is already true for each u_i, any agents who have utility functions that are affine transformations of each other have identical utility functions. Similarly for any agent whose utility functions are affine transformations of the social planner's.

Step 1: There exists i such that $\delta = \delta_i$.

Proof of Step 1: Suppose to the contrary that $\delta \neq \delta_i$ for all i.
For any $0 < x < 1$, consider $\mathbf{C} = (x, x, \ldots)$ and

$$\mathbf{C}^\varepsilon = \left(x + \varepsilon(1 - \gamma), x - 2\frac{\varepsilon}{\delta}, x + \frac{\varepsilon}{\delta^2}, x, x \ldots \right),$$

where $\varepsilon > 0$.

---

[34] Note that in that result, $\sum_i w_i = 1$, but this can be rescaled via λ.
[35] The increasing nature of u follows from the unanimity and increasing u_is.

From Taylor's approximation, for any i:

$$U_i(\mathbf{C}^\varepsilon) = U_i(\mathbf{C}) + \varepsilon u_i'(x)\left[\left(1 - \frac{\delta_i}{\delta}\right)^2 - \gamma\right] + O(\varepsilon^2).$$

Select γ so that

$$0 < \gamma < \min_i\left(1 - \frac{\delta_i}{\delta}\right)^2,$$

which is possible given our supposition that $\delta \neq \delta_i$ for all i. It follows that

$$U_i(\mathbf{C}^\varepsilon) > U_i(\mathbf{C})$$

for all i and sufficiently small ε. Note, however, that $W(\mathbf{C}^\varepsilon)$ is approximately

$$W(\mathbf{C}^\varepsilon) = W(\mathbf{C}) - \gamma\varepsilon u'(x) + O(\varepsilon^2),$$

and so for ε small enough, unanimity is violated, which is a contradiction.
Therefore, there exists i such that $\delta = \delta_i$.

Step 2: If any agents who have the same discount factor also have the same utility function, then $u = u_i$, where i is an agent with discount factor of $\delta_i = \delta$.

Proof of Step 2: Suppose the contrary, so that $\delta_i = \delta$ and yet $u \neq u_i$ (so under our normalization, these are not affine transformations of each other). Then, there exists $0 < x < 1, 0 < y < 1$, and $\alpha > 0$ such that $\frac{u_i'(x)}{u_i'(y)} > \alpha > \frac{u'(x)}{u'(y)}$.
Set $\mathbf{C} = (x, x, x, y, x, x\ldots)$ and

$$\mathbf{C}^\varepsilon = \left(x + \varepsilon, x - 2\frac{\varepsilon}{\delta}, x + \frac{\varepsilon}{\delta^2}, y - \alpha\gamma\varepsilon, x + \frac{\gamma\varepsilon}{\delta}, x, x\ldots\right),$$

for $\varepsilon > 0$.
As before, for any j such that $\delta_j \neq \delta_i$,

$$U_j(\mathbf{C}^\varepsilon) = U_j(\mathbf{C}) + \varepsilon\left[\left(1 - \frac{\delta_j}{\delta}\right)^2 + \frac{\gamma\delta_j^4}{\delta}\right]u_j'(x) - \alpha\gamma\varepsilon\delta_j^3 u_j'(y) + O(\varepsilon^2).$$

Since $\delta_j \neq \delta$, for sufficiently small ε and $\gamma = \sqrt{\varepsilon}$, $U_j(\mathbf{C}^\varepsilon) > U_j(\mathbf{C})$.
By a similar argument,

$$W(\mathbf{C}^\varepsilon) = W(\mathbf{C}) - \delta^3\gamma\varepsilon[\alpha u'(y) - u'(x)] + O(\varepsilon^2),$$

while $U_i(\mathbf{C}^\varepsilon)$ can be written as

$$U_i(\mathbf{C}^\varepsilon) = U_i(\mathbf{C}) - \delta^3\gamma\varepsilon[\alpha u_i'(y) - u_i'(x)] + O(\varepsilon^2).$$

For sufficiently small ε and $\gamma = \sqrt{\varepsilon}$, it follows that $W(\mathbf{C}^\varepsilon) < W(\mathbf{C})$ and $U_i(\mathbf{C}^\varepsilon) > U_i(\mathbf{C})$. Given that we have already established that $U_j(\mathbf{C}^\varepsilon) > U_j(\mathbf{C})$ for all $j \neq i$ for who $\delta_j \neq \delta_i$ (and the others have the same preferences as i), this violates unanimity. Therefore, our supposition was incorrect and $u = u_i$. ∎

PROOF OF THEOREM 3:

Suppose that $\mathbf{C}^*$ is interior and optimal for δ^*, so that

$$u'(c_1^*) \ = \ u'(c_2^*)\delta^*\gamma \ = \ u'(c_3^*)(\delta^*\gamma)^2.$$

Let $\mathbf{C}[a, b] \ = \ (c_1^* + a/\gamma, c_2^* - a - b, c_3^* + \gamma b, 0, \ldots)$.

For small a and b, $U_i(\mathbf{C}[a, b])$ is approximately $\big($up to expressions of $O(a^2)$, $O(b^2)\big)$,

$$(5) \qquad\qquad U_i(\mathbf{C}^*) + \frac{u'(c_1^*)}{\gamma}\left[a\left(1 - \frac{\delta_i}{\delta^*}\right) - b\frac{\delta_i}{\delta^*}\left(1 - \frac{\delta_i}{\delta^*}\right)\right].$$

CLAIM: *There exist small a', b' and a'', b'' such that $\mathbf{C}[a', b'], \mathbf{C}[a'', b''] \in B$, and for any linear ordering over $\mathbf{C}^*, \mathbf{C}[a', b']$, and $\mathbf{C}[a'', b'']$, there is a discount factor that implies that linear ordering of preferences. So, there is a set Δ of six discount factors that result in any ordering over these three alternatives.*

Restricting attention to $\mathbf{C}^*, \mathbf{C}[a', b'], \mathbf{C}[a'', b'']$, and preference profiles with discount factors in Δ, we can apply Arrow's theorem to conclude that any binary voting rule that is unanimous and non-dictatorial on this domain of preferences and relative to these consumption streams (both these conditions implied by the local no-dictator condition) is intransitive over these three alternatives and the domain of preferences composed of profiles where agents all have discount factors in Δ. The proposition then follows directly.

PROOF OF CLAIM:

Find x such that $1 \ < \ x \ < \ 1/\delta^*$ and set $[a', b'] \ = \ [x\varepsilon, \varepsilon]$ and $[a'', b''] \ = \ [\varepsilon, x\varepsilon]$ for some small ε (small enough to ensure $\mathbf{C}[a', b']$ and $\mathbf{C}[a'', b'']$ are interior and satisfy the restrictions below).

First, let us find δ's that entail $\mathbf{C}^*$ being ranked as most preferred and the other two consumption streams second and third, in either order. By the continuity of preferences, for δ near enough to δ^*, $\mathbf{C}^*$ is most preferred and $\mathbf{C}[a'', b'']$ and $\mathbf{C}[a', b']$ lead to (nearly) the same, but lower, utility. For $\delta_i \ < \ \delta^*$ the relative utility of $\mathbf{C}[a', b']$ rises relative to $\mathbf{C}[a'', b'']$ since $\left(1 - \frac{\delta_i}{\delta^*}\right)$ is positive (see equation (5)), while the reverse happens as $\delta_i \ > \ \delta^*$ since $\left(1 - \frac{\delta_i}{\delta^*}\right)$ is then negative. We choose δ_1, δ_2 such that $\delta_1 < \delta^* < \delta_2$, and both are close enough to δ^* to preserve the top ranking of $\mathbf{C}^*$ relative to the other two consumption streams, but imply a different ranking of $\mathbf{C}[a'', b'']$ and $\mathbf{C}[a', b']$.

Next, let us find δs that induce $\mathbf{C}[a',b']$ as the most preferred of the three streams, and rank the other two consumption streams in an arbitrary order. Consider $\delta' = \delta^*/x$. It follows from (5) that under $\delta' < \delta^*$, $\mathbf{C}[a',b']$ is most preferred and is nearly, except for second-order effects for small enough ε, indifferent between the other two consumption streams. Moreover, we can find δ_3 slightly smaller than δ' inducing $\mathbf{C}^*$ to be preferred to $\mathbf{C}[a'',b'']$, and δ_4 slightly larger than δ' inducing the reverse, both while still maintaining $\mathbf{C}[a',b']$ as the most preferred.

By an analogous argument, we can find δ_5 implying $\mathbf{C}[a'',b'']$ being preferred to $\mathbf{C}^*$ that is preferred to $\mathbf{C}[a',b']$, and δ_6 implying $\mathbf{C}[a'',b'']$ being preferred to $\mathbf{C}[a',b']$ that is preferred to $\mathbf{C}^*$. The specification of $\Delta = \{\delta_1, \ldots, \delta_6\}$ completes the proof. ∎

PROOF OF PROPOSITION 2:

By the suppositions in the proposition, and by ordering agents in nondecreasing order of discount factors, we end up with groups $\mathcal{S}_1, \ldots, \mathcal{S}_K$ such that the groups collect the agents with identical discount factors.

Let

$$
\mathbf{D} = \begin{pmatrix}
1 & \delta_1 & \delta_1^2 & \cdots & \delta_1^{K-1} \\
1 & \delta_2 & \delta_2^2 & \cdots & \delta_2^{K-1} \\
\vdots & \vdots & \vdots & & \vdots \\
1 & \delta_K & \delta_K^2 & \cdots & \delta_K^{K-1}
\end{pmatrix},
$$

where the labeling is such that each of the discount factors $\delta_1, \ldots, \delta_K$ is distinct. Since $\delta_i \neq \delta_j$ for all i,j, the matrix $\mathbf{D}$ is invertible. In particular, the system $\mathbf{Dx} = \mathbf{a}$ has a solution for any vector $\mathbf{a} \in \mathbb{R}^K$.

Find k_1, k_2, and k_3 partitioning the agents according to their discount factors, such that agents with the lowest k_1 discount factors form one group, the next k_2 discount factors form another group, and the last k_3 discount factors form the third group, and such that any two groups form a strict majority. For $b > 0$, consider the following vectors:

$$
\mathbf{a}^1 = \begin{pmatrix} \left.\begin{matrix} -b \\ \vdots \\ -b \end{matrix}\right\}k_1 \\ \left.\begin{matrix} b/2 \\ \vdots \\ b/2 \end{matrix}\right\}k_2 \\ \left.\begin{matrix} b/2 \\ \vdots \\ b/2 \end{matrix}\right\}k_3 \end{pmatrix}, \quad
\mathbf{a}^2 = \begin{pmatrix} \left.\begin{matrix} b/2 \\ \vdots \\ b/2 \end{matrix}\right\}k_1 \\ \left.\begin{matrix} -b \\ \vdots \\ -b \end{matrix}\right\}k_2 \\ \left.\begin{matrix} b/2 \\ \vdots \\ b/2 \end{matrix}\right\}k_3 \end{pmatrix}, \quad
\mathbf{a}^3 = \begin{pmatrix} \left.\begin{matrix} b/2 \\ \vdots \\ b/2 \end{matrix}\right\}k_1 \\ \left.\begin{matrix} b/2 \\ \vdots \\ b/2 \end{matrix}\right\}k_2 \\ \left.\begin{matrix} -b \\ \vdots \\ -b \end{matrix}\right\}k_3 \end{pmatrix}
$$

and let $\mathbf{x}^1, \ldots, \mathbf{x}^3$ be defined so that $\mathbf{D}\mathbf{x}^i = \mathbf{a}^i$ for all $i = 1, \ldots, 3$.

Notice that $\sum_{i=1}^3 \mathbf{a}^i = 0$, so that $\sum_{i=1}^3 \mathbf{x}^i = 0$ (given the invertibility of $\mathbf{D}$ and the fact that $\mathbf{D}\sum_{i=1}^3 \mathbf{x}^i = \sum_{i=1}^3 \mathbf{a}^i = 0$).

Define now the sequence of consumption streams $\mathbf{C}^1, \ldots, \mathbf{C}^3$, as follows:

$$\mathbf{C}^1 = (1/2, \ldots, 1/2, 0, \ldots)$$

$$\mathbf{C}^2 = \left(u^{-1}\big(u(c_1^1) - x_1^1\big), \ldots, u^{-1}\big(u(c_k^1) - x_k^1\big), \ldots, u^{-1}\big(u(c_K^1) - x_K^1\big), 0, 0, \ldots \right)$$

$$\mathbf{C}^3 = \left(u^{-1}\big(u(c_1^2) - x_1^2\big), \ldots, u^{-1}\big(u(c_k^2) - x_k^2\big), \ldots, u^{-1}\big(u(c_K^2) - x_K^2\big), 0, 0, \ldots \right).$$

Given the fact that u is increasing and continuous, this can be done for small enough b.

Note that, by construction, for any $t \leq n$, $u(c_t^j) - u(c_t^{j+1}) = x_t^j$ for $j = 1, \ldots, 3$. Since $\sum_{i=1}^3 \mathbf{x}^i = \mathbf{0}$, it follows that $u(c_t^3) - u(c_t^1) = x_t^1$.

In particular,

$$\mathbf{C}^1 P[U] \mathbf{C}^2 P[U] \mathbf{C}^3 P[U] \mathbf{C}^1,$$

which is what we wanted to show. ∎

REFERENCES

Abdellaoui, Mohammed, Olivier l'Haridon, Corina Paraschiv. 2010. "Individual vs. Collective Behavior: An Experimental Investigation of Risk and Time Preferences in Couples." http://www.hec.fr/content/download/50938/454640/version/1/file/CR944+O.+L%27Haridon.pdf.

Ainslie, George. 1975. "Specious Reward: A Behavioral Theory of Impulsiveness and Impulse Control." *Psychological Bulletin* 82 (4): 463–96.

Amador, Manuel, Iván Werning, and George-Marios Angeletos. 2006. "Commitment vs. Flexibility." *Econometrica* 74 (2): 365–96.

Ambrus, Attila, and Kareen Rozen. 2009. "Rationalizing Choice with Multi-Self Models." Unpublished.

Andreoni, James, and Charles Sprenger. 2012a. "Estimating Time Preferences from Convex Budget Sets." *American Economic Review* 102 (7): 3333–56.

Andreoni, James, and Charles Sprenger. 2012b. "Risk Preferences are Not Time Preferences." *American Economic Review* 102 (7): 3357–76.

Arrow, Kenneth J. 1950. "A Difficulty in the Concept of Social Welfare." *Journal of Political Economy* 58 (4): 328–46.

Arrow, Kenneth J. 1963. *Social Choice and Individual Values*. New York: John Wiley and Sons.

Barberà, Salvador, and Matthew O. Jackson. 2006. "On the Weights of Nations: Assigning Voting Power to Heterogeneous Voters." *Journal of Political Economy* 114 (2): 317–39.

Benabou, Roland, and Jean Tirole. 2006. "Incentives and Prosocial Behavior." *American Economic Review* 96 (5): 1652–78.

Bernheim, B. Douglas. 1999. "Comment on 'Family Bargaining and Retirement Behavior.'" In *Behavioral Economics and Retirement Policy*, edited by Henry J. Aaron, 273–81. Washington, DC: Brookings Institution Press.

Bernheim, B. Douglas, and Antonio Rangel. 2004. "Addiction and Cue-Triggered Decision Processes." *American Economic Review* 94 (5): 1558–90.

Beshears, John, James J. Choi, David I. Laibson, and Brigitte C. Madrian. 2008. "The Importance of Default Options for Retirement Saving Outcomes: Evidence from the United States." In *Lessons*

from Pension Reform in the Americas, edited by Stephen J. Kay and Tapen Sinha. Oxford: Oxford University Press.

Blackorby, Charles, Walter Bossert, and David J. Donaldson, eds. 2005. "Temporal Consistency." In *Population Issues in Social Choice Theory, Welfare Economics, and Ethics*, 272–85. Cambridge: Cambridge University Press.

Boylan, Richard T., and Richard D. McKelvey. 1995. "Voting over Economic Plans." *American Economic Review* 85 (4): 860–71.

Brocas, Isabelle, and Juan D. Carrillo. 2008. "The Brain as a Hierarchical Organization." *American Economic Review* 98 (4): 1312–46.

Browning, Martin. 2000. "The Saving Behaviour of a Two-person Household." *Scandinavian Journal of Economics* 102 (2): 235–51.

Caplin, Andrew, and John Leahy. 2004. "The Social Discount Rate." *Journal of Political Economy* 112 (6): 1257–68.

Chapman, Gretchen B., Noel T. Brewer, Elliot J. Coups, Susan Brownlee, Howard Leventhal, and Elaine A. Levanthal. 2001. "Value for the Future and Preventive Health Behavior." *Journal of Experimental Psychology: Applied* 7 (3): 235–50.

Chapman, Gretchen B., and Arthur S. Elstein. 1995. "Valuing the Future: Temporal Discounting of Health and Money." *Medical Decision Making* 15 (4): 373–86.

Cherepanov, Vadim, Timothy Feddersen, and Alvaro Sandroni. 2009. "Rationalization." Unpublished.

Condorcet, Marquis de. 1785. "Essay on the Application of Mathematics to the Theory of Decision-Making." In *Condorcet: Selected Writings*, edited by Keith M. Baker, 33–70. Indianapolis: Bobbs-Merrill, 1976.

Della Vigna, Stefano and Ulrike Malmendier. 2006. "Paying Not to Go to the Gym." *American Economic Review* 96 (3): 694–719.

Drefahl, Sven. 2010. "How Does the Age Gap Between Partners Affect Their Survival?" *Demography* 47 (2): 313–26.

Evren, Ozgur, and Efe A. Ok. 2011. "On the Multi-utility Representation of Preference Relations." *Journal of Mathematical Economics* 47 (4): 554–63.

Farmer, J. Doyne, and John Geanakoplos. 2009. "Hyperbolic Discounting Is Rational: Valuing the Far Future with Uncertain Discount Rates." http://cowles.econ.yale.edu/P/cd/d17a/d1719.pdf.

Feldstein, M. S. 1964. "The Social Time Preference Discount Rate in Cost Benefit Analysis." *Economic Journal* 74 (294): 360–79.

Frederick, Shane, George Loewenstein, and Ted O'Donoghue. 2002. "Time Discounting and Time Preference: A Critical Review." *Journal of Economic Literature* 40 (2): 351–401.

Fudenberg, Drew, and David K. Levine. 2006. "A Dual Self Model of Impulse Control." *American Economic Review* 96 (5): 1449–76.

Galperti, Simone, and Bruno Strulovici. 2014. "Forward-Looking Behavior Revisited: A Foundation of Time Inconsistency." https://economics.uchicago.edu/pdf/Strulovici_052714.pdf.

Glassman, Matthew Eric, and Amber Hope Wilhelm. 2013. *Congressional Careers: Service Tenure and Patterns of Member Service*, 1789–2013. Congressional Research Service (CRS) Report for Congress. Washington, DC January.

Glimcher, Paul W., and Aldo Rustichini. 2004. "Neuroeconomics: The Consilience of Brain and Decision." *Science* 306 (5695): 447–52.

Gollier, Christian, and Richard Zeckhauser. 2005. "Aggregation of Different Time Preferences." *Journal of Political Economy* 113 (4): 878–96.

Green, Jerry R., and Daniel A. Hojman. 2009. "Choice, Rationality and Welfare Measurement." http://www.hks.harvard.edu/fs/dhojman/papers/CRW.pdf.

Hare, Todd A., Colin F. Camerer, and Antonio Rangel. 2009. "Self-Control in Decision-Making Involves Modulation of the vmPFC Valuation System." *Science* 324 (5927): 646–48.

Herrnstein, Richard J. 1961. "Relative and absolute strength of response as a function of frequency of reinforcement." *Journal of the Experimental Analysis of Behavior* 4 (3): 267–72.

Hertzberg, Andrew. 2010. "Exponential Individuals, Hyperbolic Households." https://www0.gsb.columbia.edu/mygsb/faculty/research/pubfiles/5005/HH_Paper_Oct_2010.pdf.

Hori, Takeo, and Koichi Futagami. 2012. "A Non-Unitary Discount Rate Model." Unpublished.

Jackson, Matthew O., and Leeat Yariv. 2014. "Present Bias and Collective Dynamic Choice in the Lab." *American Economic Review* 104 (12): 4184–4204.

Jamison, Dean T., and Julian C. Jamison. 2007. "The Amount and Speed of Discounting." http://cear.gsu.edu/files/2013/03/Seminar_2010_0416_Jamison.pdf.

Koopmans, Tjalling C. 1960. "Stationary Ordinal Utility and Impatience." *Econometrica* 28 (2): 287–309.

Kydland, Finn E., and Edward C. Prescott. 1977. "Rules Rather than Discretion: The Inconsistency of Optimal Plans." *Journal of Political Economy* 85 (3): 473–92.

Laibson, David. 1997. "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics* 112 (2): 443–77.

Lehrer, Ehud, and Ady Pauzner. 1999. "Repeated Games with Differential Time Preferences." *Econometrica* 67 (2): 393–412.

Lehrer, Ehud, and Leeat Yariv. 1999. "Repeated Games With Incomplete Information On One Side: The Case Of Different Discounting Factors." *Mathematics of Operations Research* 24 (1): 204–19.

Lengwiler, Yvan. 2005. "Heterogeneous Patience and the Term Structure of Real Interest Rates." *American Economic Review* 95 (3): 890–96.

Marglin, Stephen A. 1963. "The Social Rate of Discount and the Optimal Rate of Investment." *Quarterly Journal of Economics* 77 (1): 95–111.

Mazzocco, Maurizio. 2007. "Household Intertemporal Behaviour: A Collective Characterization and a Test of Commitment." *Review of Economic Studies* 74 (3): 857–95.

McClure, Samuel M., Keith M. Ericson, David I. Laibson, George Loewenstein, and Jonathan D. Cohen. 2007. "Time Discounting for Primary Rewards." *Journal of Neuroscience* 27 (21): 5796–5804.

McClure, Samuel M., David I. Laibson, George Loewenstein, and J. D. Cohen. 2004. "Separate Neural Systems Value Immediate and Delayed Monetary Rewards." *Science* 306 (5695): 503–07.

McKelvey, Richard D. 1976. "Intransitivities in multidimensional voting models and some implications for agenda control." *Journal of Economic Theory* 12 (3): 472–82.

McKelvey, Richard D. 1979. "General Conditions for Global Intransitivities in Formal Voting Models." *Econometrica* 47: 1085–1112.

Mongin, Philippe. 1995. "Consistent Bayesian Aggregation." *Journal of Economic Theory* 66 (2): 313–51.

Mongin, Philippe. 1998. "The Paradox of the Bayesian Experts and State-dependent Utility Theory." *Journal of Mathematical Economics* 29: 331–61.

O'Donoghue, Ted, and Matthew Rabin. 1999. "Doing it Now or Later." *American Economic Review* 89 (1): 103–24.

Parks, Robert P. 1976. "An Impossibility Theorem for Fixed Preferences: A Dictatorial Bergson-Samuelson Welfare Function." *Review of Economic Studies* 43 (3): 447–50.

Plott, Charles R. 1967. "A Notion of Equilibrium and Its Possibility under Majority Rule." *American Economic Review* 57 (4): 787–806.

Rachlin, Howard. 2004. *The Science of Self-Control.* Cambridge: Harvard University Press.

Roberts, Kevin W. S. 1980. "Social Choice Theory: The Single-profile and Multi-profile Approaches." *Review of Economic Studies* 47 (2): 441–50.

Schaner, Simone G. 2015. "Do Opposites Detract? Intrahousehold Preference Heterogeneity and Inefficient Strategic Savings." *American Economic Journal: Applied Economics* 7 (2): 135–74.

Shafir, Sharoni. 1994. "Intransitivity of preferences in honey bees: Support for 'comparative' evaluation of foraging options." *Animal Behavior* 48 (1): 55–67.

Strotz, R. H. 1955. "Myopia and Inconsistency in Dynamic Utility Maximization." *Review of Economic Studies* 23 (3): 165–80.

Thaler, Richard H. 1981. "Some Empirical Evidence On Dynamic Inconsistency." *Economics Letters* 8: 201–07.

Thaler, Richard H., and H. M. Shefrin. 1981. "An Economic Theory of Self-Control." *Journal of Political Economy* 89 (2): 392–406.

Tversky, A. 1969. "Intransitivity of Preference." *Psychological Review* 76 (1): 31–48.

Waite, Thomas A. 2001. "Intransitive preferences in hoarding gray jays (Perisoreus canadensis)." *Behavioral Ecology and Sociobiology* 50 (2): 116–21.

Weitzman, Martin L. 2001. "Gamma Discounting." *American Economic Review* 91 (1): 260–71.

Zuber, Stéphane. 2011. "The aggregation of preferences: Can we ignore the past?" *Theory and Decision* 70 (1): 367–84.