

Collective Dynamic Choice: The Necessity of Time Inconsistency

Matthew O. Jackson* and Leeat Yariv^{†‡}

October 2010, This Draft: October 2011

Abstract

We examine collective decisions over streams of consumption. Individuals all consume the same stream and evaluate it according to time discounted and smooth utility functions. If individuals differ in their time discount factors, then the only way to aggregate their preferences while respecting unanimous preferences and time consistency is by appointing a dictator, even when all agents have exactly the same instantaneous utility function. Therefore, heterogeneous groups (or individual decision makers) who incorporate different temporal motives must be time inconsistent. Furthermore, we show that if preferences are aggregated in a utilitarian manner, they necessarily exhibit a specific form of time inconsistency: a present-bias, favoring short-term consumption and reversing choices as decisions are pushed further into the future. We also show that if preferences are aggregated by voting, the resulting choices will necessarily be intransitive even with strong restrictions placed on the feasible consumption streams. Finally, we performed lab experiments where a “social planner” makes a choice that affects the consumption stream of two other subjects who differ in their discount factors. We find that three quarters of the planners exhibit a present bias, while less than two percent are time consistent, with the remaining subjects exhibiting either future bias or situation-based inconsistencies. We estimate the collective utility functions of the planners, finding that about one third of subjects act as if they are utilitarians, while the remaining two thirds act as if they also weight inequality (negatively) together with total utility in making choices. Subjects who weight inequality more are those exhibiting mixed time inconsistencies, and those weighting it less exhibit present bias. JEL Classification Numbers: D72, D71, D03, D11, E24

Keywords: collective decisions, time inconsistency, collective utility functions, consumption plans, representative agents, voting, voting rules, majority voting, transitivity, hyperbolic discounting, present bias

*Department of Economics, Stanford University, the Santa Fe Institute, and CIFAR. <http://www.stanford.edu/~jacksonm> e-mail: jacksonm@stanford.edu

[†]Division of the Humanities and Social Sciences, Caltech. <http://www.hss.caltech.edu/~lyariv/index.htm> e-mail: lyariv@hss.caltech.edu

[‡]We thank Nageeb Ali, Sandro Ambuehl, James Andreoni, Kenneth Arrow, Mariagiovanna Baccara, Miguel Angel Ballester, Douglas Bernheim, Martin Browning, Christopher Chambers, Jeff Ely, Keith Ericson, Drew Fudenberg, Jerry Green, Olivier l’Haridon, Andrew Hertzberg, Julian Jamison, Lauren Merrill, Jochen Mierau, Massimo Morelli, Efe Ok, Antonio Rangel, Ariel Rubinstein, Erik Snowberg, and Tomasz Strzalecki for useful discussions and suggestions. We gratefully acknowledge financial support from the National Science Foundation (SES 0551014 and SES 0961481) and the Gordon and Betty Moore Foundation.

1 Introduction

1.1 Overview

Many important decisions over sequences of consumption or budget allocations are made by groups of decision makers, be it different individuals (in a household, a political committee, a district, a firm, etc.) or different motives within one individual. If the units composing the groups differ in their time preferences, a tension arises when making collective intertemporal choices. Our focus in this paper is on this tension, and how it distorts collective intertemporal decisions.

The idea that individuals may vary in their time preferences has strong empirical founding. Consider, for instance, heterosexual households. In most parts of the world, women have significantly higher life expectancies than men. For example, in the United States and the United Kingdom, current estimated life expectancies are 82 years for women and 78 for men. Similar patterns hold across the world, though the gender-specific expectancies vary. For example, the comparable statistics (for women and men, respectively) are 85 and 79 in France and Spain, 74 and 62 in Russia, 76 and 72 in China, 87 and 80 in Japan, 77 and 70 in Brazil, and 54 and 52 in South Africa.¹ A compounding effect arises through the age difference between men and women at marriage. In the United States, men are typically several years older than their wives. For instance, between 1947 and 2010, a groom was, on average, 2.3 years older than his bride (see Drefahl, 2010). In fact, Browning (2000) used Canadian data on married couples to estimate the combined effects of different life expectancies and age of marriage between spouses. His analysis suggests that the wife of a 65 year old man would have, on average, an approximately 50% longer expected survival horizon than her husband. When translated into discount factors, this would suggest that husbands and wives discount the value of savings at substantially different rates.²

Differences in time preferences are also likely to exist in many contexts other than household decisions, ranging from legislators representing different districts and making collective decisions over the size of a current budget, to board members of a company, who differ in ages and investment portfolios and make joint intertemporal allocations of company resources. Moreover, there is a volume of emerging neuroscientific evidence suggesting that some form of parallel processing and aggregation of motives, which respond differently to timed rewards, goes on in an individual's brain

¹These are expectancies from birth for children born from 2010 to 2015, see the *United Nations Statistics Division Social Indicators*, Updated in December 2010, based on data from the *World Population Prospects: The 2008 Revision* (CD-ROM Edition), supplemented by official national statistics published in the *United Nations Demographic Yearbook 2008*.

²See also Schaner (2010) for some evidence on differences in time preferences inside households in Western Kenya. We discuss further related literature below.

(e.g., see Hare, McClure, and Rangel, 2009, McClure et al., 2004, 2007, and Glimcher and Rustichini, 2007). The multitude of applications where different time preferences are aggregated when effectively choosing a time-stream of consumption makes clear the importance of understanding the outcomes that could be generated in such settings.

Nonetheless, at the foundation of much of classical economics is an assumption that firms, countries, and other organizations that are generally comprised of heterogeneous individuals act as a “representative agent”, and that all are *time consistent*. Operationally, individuals, as well as their representative, are assumed to evaluate each period’s ‘consumption’ using some instantaneous utility function and discount each period’s instantaneous utility in an exponentially decaying manner. This sort of formulation embodies an important form of time consistency: if one prefers \$10 today to \$15 tomorrow, she should also prefer \$10 in 100 days to \$15 in 101 days. Technically, this consistency requirement is useful since it leads to stationarity in dynamic models of decision-making (where actors can be either groups or individuals). While this sort of modeling enhances tractability, it may be inappropriate if aggregating heterogeneous time preferences does not lead to time consistent collective decisions.

We show, in fact, that very natural procedures for making collective dynamic decisions are inherently *time inconsistent*, even if underlying individuals are perfectly time consistent. Furthermore, aggregation rules that are frequently utilized, such as weighted utility maximization or majoritarian voting, generate some of the biases identified empirically in the context of time preferences. Utilitarian aggregation rules generate decisions exhibiting a present-bias, while majoritarian voting rules lead to intransitivities in decisions.

As motivation, consider the following simple example.³ Suppose two time-consistent individuals, Constantine and Patience, make collective decisions by maximizing their joint utilitarian welfare. They are making choices over joint (say, household) consumption streams. Constantine has a discount factor of 0.5, while Patience has a discount factor of 0.8. Both experience identical instantaneous utility. When Constantine and Patience are considering 10 utiles today relative to 15 utiles tomorrow, they are comparing a total collective utility of $10 + 10 = 20$ with $15(0.8 + 0.5) = 19.5$. They therefore jointly choose the immediate 10 utiles. Suppose now that Constantine and Patience compare 10 utiles at day $t \geq 1$ and 15 utiles at day $t + 1$. Now, the comparison is between a total collective utility of $10(0.8^t + 0.5^t)$ and $15(0.8^{t+1} + 0.5^{t+1})$. For instance, if $t = 1$, the comparison

³This example is in line with an observation that representative agents may not exhibit stationary discount rates, a point that has roots as early as Marglin (1963) and Feldstein (1964), and has been examined by others including Becker (1992) and Gollier and Zeckhauser (2005). We provide additional references below.

is between 13 and 13.35, and the delayed consumption leads to greater total utility. In fact, for any $t \geq 1$, the total utility would be larger from the delayed consumption of 15 utiles, but the immediate decision exhibits the reverse. If we observe the household's behavior, we would find it to be time-inconsistent: they prefer an immediate reward of 10 to a delayed reward of 15, but reverse their decision whenever rewards are both pushed into the future.

As we show, this example turns out to be quite general. Whenever alternatives are chosen to maximize a weighted sum of individual utility functions, corresponding to a utilitarian planner's objective function, collective decisions will exhibit a particular form of time inconsistency. Namely, whenever the group is heterogeneous, maximizing efficiency would lead to a present-bias: the group will behave not only as if it prefers early to later consumption, it will also appear more and more patient as decisions are postponed into the future. In fact, we show that for a uniform distribution of discount factors in the population, overall efficiency maximization translates into behavior that corresponds to hyperbolic discounting.

These observations have important implications from an econometric perspective. Estimating the preferences of a heterogeneous population as if it were homogeneous (e.g., estimating the preferences of a representative agent) boils down to estimating an *average* of preferences in the population, which is technically tantamount to estimating preferences that are derived from summing up agents' utilities. Our results then imply that an econometrician trying to assess time preferences by averaging population behavior may come to the conclusion that preferences are time-inconsistent and exhibit a present-bias whenever there is some heterogeneity within the population. Viewed in another way, the "representative agent" may appear time-inconsistent even if the population that it represents is perfectly time consistent.

One of our main results shows that the class of utilitarian aggregation rules is, in many ways, not special. Consider any group of agents who are each individually time consistent. Individuals have arbitrary and heterogeneous discount factors and arbitrary (and possibly heterogeneous) instantaneous utility functions that are well-behaved. A minimal restriction on any aggregation rule is that it respect unanimity, or Pareto efficiency. In other words, whenever everyone agrees that one consumption stream is superior to another, the collective decision exhibits the same preference as well.⁴ We illustrate that any time-consistent aggregation rule that satisfies this minimal restriction must be dictatorial, i.e., it must track the preferences of only one of the group's members. In other words, if an aggregation rule is non-dictatorial and respects unanimous choices, then it must be

⁴Clearly, if the group chooses consumption streams by maximizing the sum of participants' utility functions, consensual preferences are respected.

time inconsistent.

This result is quite general in terms of the format of consumption – consumption streams may entail elements that are private to particular agents, thus allowing outcomes to be the consequence of bargaining. Furthermore, since the result regards observed ultimate choices of consumption streams, it relates to 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 choices can be represented by a collective utility function, or even simply evaluated 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, our results imply that the wide array of contexts involving non-dictatorial choice mechanisms either necessitate commitment devices or will involve reversals over time.

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 will correspond to a weighted sum of the individual agents' utility functions. Therefore, combined with our previous results, whenever non-dictatorial, these rules will exhibit a present-bias.

Our last theoretical results examine other methods of aggregating preferences: general voting rules. Clearly, majority voting is another aggregation method that is widely used in economic and political contexts, and is qualitatively different from aggregation procedures relying on agents' cardinal utility functions (such as efficiency maximization or maximization of average net present values). One may conjecture that there would be a “median voter,” who might effectively appear as 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 preferences are not extreme in a well-defined sense, then for a rich set of consumption streams we can find two other consumption streams with which a voting cycle is formed.

To summarize our main theoretical results:

- If preferences are aggregated by some weighted averaging of individual utility functions, then the collective utility function is present-biased.
- If individual preferences are aggregated into some collective utility function that is non-dictatorial and respects unanimity, then that collective utility function must be time inconsistent.

- If preferences are aggregated via any voting rule that is non-dictatorial and respects unanimity, such as majority voting, then the resulting social welfare ordering will exhibit cycles, unless the set of consumption streams is severely restricted.

We complement our theoretical analysis with a series of laboratory experiments, in which one subject makes a choice that affects the consumption stream of two *other* subjects who differ in their discount factors. In other words, subjects acted as if they *social planners* aggregating the preferences of two other individuals. There are three main insights that emerge. First, as suggested by our main result, social planners are time inconsistent. In particular, we find that three quarters of the planners exhibit a present bias, while less than two percent are time consistent, with the remaining subjects exhibiting either future bias or situation-based inconsistencies. Second, for the most part, the sum of payoffs and their standard deviation explain individual choices, with subjects balancing utilitarian motives with egalitarian ones. Third, subjects can be classified into ‘types’ that correspond to how they weigh utilitarian and egalitarian motives. Specifically, we estimate the collective utility functions of the planners, finding that just under one third of subjects act as if they are pure utilitarians, while the remaining two thirds act as if they also weight inequality (negatively) together with total utility in making choices. Subjects who weight inequality the most are those exhibiting mixed time inconsistencies, and those weighting it slightly exhibit present bias, with most (about two-thirds of) subjects who act as if they weight inequality weighting it slightly but noticeably.⁵

1.2 Related Literature

From the perspective of individual decision-making, we note that 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 a lab, many individual decision makers appear to be time consistent (see Andreoni and Sprenger, 2010a, 2010b). 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, Choi, Laibson, and Madrian, 2008) to physical exercise (della Vigna and Malmendier, 2006). Our results provide a potential explanation for why individual decision makers may behave in an inconsistent or even intransitive manner when faced with decisions over streams of consumption: individuals may be thought of as making decisions by aggregating heterogeneous underlying “personalities” or “motives.” That is, a simple model of an individual as

⁵In particular, that two thirds of the subjects’ choices diverge from a pure utilitarian decision only when the standard deviation in payoffs is more than three times the difference in average payoffs across decisions.

aggregating a group of internal but diverse preferences, matches observed behaviors that cannot be matched with classical models.

With regards to aggregating preferences with heterogeneous discount factors, it has long been recognized that there are difficulties. Marglin (1963) and Feldstein (1964) noted that it was difficult to derive an appropriate aggregate time independent discount rate for a planner facing a society of heterogeneous agents. More recent work in the context of household decisions, Bernheim (1999), Browning (2000), Mazzocco (2007), Xue (2008), Hertzberg (2010), Abdellaoui, l'Haridon, and Paraschiv (2010), and Schaner (2010), among others, have considered the implications of preference heterogeneity on intertemporal consumption decisions under particular aggregation protocols. For instance, they illustrate that households may have hyperbolic preferences (in Hertzberg, 2010) and the role of commitment devices in determining consumption patterns (in Mazzocco, 2007, and Schaner, 2010). Work by Weitzman (2001), Caplin and Leahy (2004), Blackorby, Bossert, and Donaldson (2005), Gollier and Zeckhauser (2005), Green and Hojman (2009), and Zuber (2010), examine variations on such planner aggregation issues in more detail.⁶ For example, Gollier and Zeckhauser (2005) show that a representative agent will have a time-varying discount factor if there is sufficient uncertainty and heterogeneity in the environment. In terms of time inconsistency, our results show that whenever agents differ in their discount factors, even in deterministic settings, where agents have identical instantaneous utility functions, time inconsistency must result in every non-dictatorial aggregation method. Thus, our results illustrate that some of the phenomena identified in the literature are not unique to the specific planner or representative agent formulations, but hold generally, and emerge even when there is only heterogeneity in discount factors.

Zuber (2010) is the closest to ours in this regard, showing that a planner can only aggregate agent preferences in a stationary and consistent manner if all agents have the same discount factor. That result is in a significantly different setting as 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. This is a quite substantial constraint that requires a different approach. 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.

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 con-

⁶See 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 (that technically is equivalent to a particular form of utilitarian aggregation in our setting) as the foundations for hyperbolic discounting,

sistent one) is analogous to a subjective expected utility function: time periods are 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 (technically, Mongin’s result would impose a sort of ‘continuity’ requirement on our time dimension⁷). 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 regards 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 multi-dimensional (“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 how, 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. Moreover, beyond showing that there are issues with voting cycles, we also examine collective utility functions and show the general impossibility of time consistency, which is quite different from any of the above mentioned papers, and our results concerning the intransitivity of voting rules apply to a very wide class of procedures, containing majority rule as a special case.

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 literatures documenting time inconsistencies and intransitivities are too vast to cover here.⁸ The idea that individuals might be usefully thought of as having some internally inconsistent preferences appears in a variety of places, possibly the most related of

⁷He also requires some additional conditions to derive his results.

⁸Some seminal references include Hershman (1961) and Thaler (1981) for time inconsistencies and Tversky (1969) for intransitivities. 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, 2000). Intransitivities have been observed in bees, as in Shafir (1994) and jays, as in Waite (2001), where the jays exhibit intransitivities 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).

which is the recent paper by Green and Hojman (2009), that provides a general revealed-preference welfare bound analysis allowing for such possibilities.⁹ A contribution of the current paper is the insight that viewing individuals as nondegenerate collectives leads *necessarily* to behaviors exhibiting time inconsistency and/or intransitivities in ways that are in line with empirical observations on a variety of dimensions.

In terms of our experimental results, our design is related to that of Engelmann and Strobel (2004), who consider allocation decisions in which agents effectively made decisions ‘behind the veil of ignorance.’ Agents were ultimately assigned a random role in the relevant group and paid according to what their chosen allocation had specified for that role. There are two main differences between our experiments and Engelmann and Strobel’s. First, we tailored our allocations to capture different aspects of time preferences (present-bias, future-bias, etc.). Second, our design attempts to identify preferences of social planners, rather than general attitudes toward bargaining outcomes, and so in our setting, experimental planners make choices that truly affect only others.¹⁰

2 The Setting

2.1 Agents and Consumption Streams

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

A stream of consumption is denoted by $C = (c_1, c_2, \dots)$, where each $c_t \in [0, 1]$.¹¹

Once a consumption stream is decided upon, all agents have utility functions that are functions of that same stream.

In Section 8.3 we discuss extensions to settings where different individuals consume different streams.

In terms of interpretations, it is not critical that the consumption be common per se, but rather that individuals making collective choices each be able to evaluate their personal utility based on

⁹Some 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 (2008), Ambrus and Rozen (2009), and Cherepanov, Feddersen, and Sandroni (2009). 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.

¹⁰When agents are randomly paid according to a role, a pure risk-neutral expected utility maximizer would optimally choose alternatives maximizing the sum of payoffs. Indeed, utilitarian motives appear quite strongly in the Engelmann and Strobel experiments. Our design does not directly tie utilitarian motives and experimental payoffs.

¹¹The assumption that timed consumption is uniformly bounded will assure that net present values are always well-defined (for well-behaved instantaneous utility functions). The assumption that it is bounded by 0 and 1 is without loss of generality.

the collective decision. This would apply to a variety of examples, e.g., ones in which some entity (a government, a household, etc.) decides upon the allocation of some budget across different time periods. What is presumed is that agents can predict their individual resulting utilities conditional on a given budget being spent in a given period. It need not necessarily be that the budget be spent on public goods or some common consumption, only that the resulting utilities are predictable. Our focus is thus on the collective decision over allocations across periods taking any bargaining within periods as a given. Of course, in the interpretation of multiple motives within a single person, the consumption truly is common.

Note that finite-horizon problems can be considered by examining strings that have only finitely many positive entries.

2.2 Individual Agents

We consider settings where agents maximize a time additive discounted utility function. That is, agent i has a discount factor $\delta_i \in (0, 1)$ and an increasing and twice continuously differentiable instantaneous utility function $u_i : [0, 1] \rightarrow \mathbb{R}$ such that a stream $C = (c_1, c_2, \dots)$ is evaluated as¹²

$$U_i(C) = \sum_t \delta_i^t u_i(c_t). \quad (1)$$

Let \mathcal{U} denote the set of possible preferences (δ_i, u_i) satisfying the conditions above.¹³

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

It is natural to have heterogeneity in preferences across individuals in many applications. Members of the Senate may be at different phases in their term or represent constituents with different short versus long term needs, individuals in management teams may be of different ages or have different objectives, a person may balance different temporal motives (that effectively act as different agents), households may consist of individuals with different tastes and responsibilities, etc. We note that in many such contexts it is highly conceivable that agents would have heterogeneous discount factors and/or utility functions.

¹²This specification of time discounted additively separable utility clearly precludes certain sorts of complementarities in consumption across periods. Nonetheless, as it has become so standard in the literature, partly because of its time consistency properties, and in part because some of the complementarities might be secondary, it is still important to understand its properties.

¹³For notational convenience, we ignore the normalization factors (of $\{1 - \delta_i\}_i$) that are of no consequence to our results.

We also remark that we examine the case where all primitive 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 for these underlying agents would lead to such conclusions in the aggregate a fortiori.

2.3 Collective Decisions

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

2.3.1 Collective Utility Functions

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

A collective utility function can be thought of as providing a “planner’s” utility function for a society. Examples include taking a weighted average of the agents’ utility functions ($V[U](C) = \sum_i w_i U_i(C)$) or considering the minimum of agents’ utilities ($V[U](C) = \min_i U_i(C)$).

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

2.3.2 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 society by $R(U)$ for $U = (\delta_1, u_1, \dots, \delta_n, u_n) \in \mathcal{U}^n$. In some cases the social welfare orderings will be complete and reflexive, but that need not be the case.

The induced strict preference relation $P(U)$ is defined as usual by

$$\begin{aligned} C P(U) C' &\text{ if} \\ C R(U) C' &\text{ and not } C' R(U) C. \end{aligned}$$

One prominent example of such a preference relation is the case in which $CP(U)C'$ if a majority of individuals find C preferred to C' , which corresponds to the standard majority rule.

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

3 Utilitarian Aggregation of Preferences and Present-bias

We begin by examining a particular class of collective utility functions, which is perhaps the most natural and prominent: that of *weighted utilitarian functions*. We show that this class of utility functions exhibits a particular sort of time inconsistency, one that matches some evidence on behavior.

A collective utility function is a *weighted utilitarian function* if there exist weights $w_i \in [0, 1]$ such that $\sum_i w_i = 1$ and

$$V(C) = \sum_i w_i U_i(C).$$

The weights w_i can correspond to the fraction of the population with each particular discount factor, or to factors used to trace out the Pareto frontier. Furthermore, if $w_i \propto \frac{1-\delta_i}{n}$, then $V(C)$ corresponds to the average (normalized) discounted utilities across the population. Thus, from an econometric perspective, it is $V(C)$ that is often an object of estimation.

Utilitarian functions can generate very familiar ‘non-standard’ preferences, as the following two examples illustrate.

Example 1 (Hyperbolic Discounting) Consider a society with a continuum of agents¹⁴ with $\delta_i = i$, where i is uniformly distributed on $[0, 1]$. All agents share an identical instantaneous utility function, $u_i = u$.

In this simple case, the resulting utilitarian collective utility function assessed at consumption stream C is:

$$V(C) = \sum_t \int \delta_i^t u(c_t) di = \sum_t \int i^t u(c_t) di$$

or

$$V(C) = \sum_t \frac{u(c_t)}{1+t}.$$

¹⁴This moves outside of our finite model, but is easily approximated in the finite model.

Thus, a society with uniformly distributed discount factors generates a utilitarian collective utility function that is purely hyperbolic.^{15,16}

Example 2 (Fixed Costs of Delay) Consider a society of two types of agents. One type of agent is completely impetuous: $\delta_1 = 0$ and $u_1(c_t) = a$ for any $c_t > 0$.¹⁷ This type of agent simply wants immediate gratification, and is insensitive to the amount of immediate consumption. The other type of agent is standard, with a discount factor $\delta_2 > 0$ and instantaneous utility u_2 , an increasing function. Let λ be the proportion of impetuous agents. The resulting utilitarian collective function assessed at consumption stream C is:

$$V(C) = \sum_t \lambda \delta_1^t u_1(c_t) + (1 - \lambda) \delta_2^t u_2(c_t)$$

or

$$V(C) = \begin{cases} \lambda a + (1 - \lambda) \sum_t \delta_2^t u_2(c_t) & \text{if } c_1 > 0 \\ (1 - \lambda) \sum_t \delta_2^t u_2(c_t) & \text{otherwise} \end{cases} .$$

This society exhibits a collective utility function that has a fixed cost of delaying any immediate consumption, but exhibits exponential discounting thereafter, matching the model and array of experiments presented by Benhabib, Bisin, and Schotter (2010).

The formulation of hyperbolic discounting is a prominent one that is often used to capture a *present-bias*, such that a decision maker has a different assessment of early rewards relative to later ones, placing greater (relative) weight on the former. The collective utility function entailing fixed costs of delay corresponds to even more extreme contrast between the present period and the future. As we now illustrate, it turns out that a present-bias is inherent to utilitarian aggregation of heterogeneous time preferences.

Let $C[x, t]$ 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:*

- $V(C[x, t]) \leq V(C[y, t + k])$ implies $V(C[x, t + 1]) \leq V(C[y, t + k + 1])$ for any x, y , and $t \geq 0$, $k \geq 1$, and

¹⁵We note that we did not normalize each agent's utility (with the factors $\{1 - \delta_i\}$). In terms of normalized discounted utility, this formulation effectively puts greater weight on agents that are more patient. Normalizing provides a variation on the functional form, but with similar induced behavior.

¹⁶This example is reminiscent of a model analyzed by Sozou (1998), who illustrated how uncertainty over exponential hazard rates can generate hyperbolic discounting. Relatedly, Dasgupta and Maskin (2004) noted that uncertainty about future rewards may translate into an apparent present-bias. Xue (2008) shows that aggregation in a two person model can exhibit some quasi-hyperbolic features.

¹⁷Again, this sits at the extreme of our model as we assume that $\delta_i > 0$ and u_i is increasing, but is easily approximated within the model.

- For any $t \geq 1$ and $k \geq 1$, there exist x and y such that $V(C[x, 1]) > V(C[y, k + 1])$ while $V(C[x, t + 1]) < V(C[y, t + k + 1])$.

Present bias indicates that decisions made using the collective utility function correspond to more impatience as the relevant consumption becomes more immediate. The first part of the definition states that if one level of consumption at some time $t + k$ is preferred to another at an earlier time t , then the same preference ordering holds when both consumptions are postponed, so that future preferences are at least as patient. The second part of the definition indicates that there exist some choices that reverse themselves over time: if offered today the immediate consumption is preferred, while if deferred to some point in the future the choice corresponds to more patience (a description going back to Strotz, 1955 and corresponding to the impulsiveness suggested by Ainslie, 1975).

The following proposition provides the formal claim that any non-trivial weighted utilitarian function exhibits present-bias.

PROPOSITION 1 *If $u_i = u$ for all i , where u is continuous and strictly increasing, and V is a weighted utilitarian function with weights $w_i > 0$ and $w_j > 0$ for at least two agents i and j such that $\delta_i \neq \delta_j$, then V is present-biased.*

The intuition behind this result is straightforward. Suppose the group of n agents is characterized by the sequence of discount factors $\delta_1 < \delta_2 < \dots < \delta_n$. The effective discount factor of period t consumption in the collective utility is given by $w_1\delta_1^{t-1} + w_2\delta_2^{t-1} + \dots + w_n\delta_n^{t-1}$. For simplicity, assume that all the weights are positive, $w_i > 0$ for all i . As t increases,

$$\frac{w_1\delta_1^t + w_2\delta_2^t + \dots + w_n\delta_n^t}{w_1\delta_1^{t-1} + w_2\delta_2^{t-1} + \dots + w_n\delta_n^{t-1}} \rightarrow \delta_n.$$

Thus, as t grows agents with greater patience gain more and more implicit gravity in determining the rates of substitution across time, and the collective utility exhibits more and more patience. Since over time collective utility exhibits more patience, there is a present-bias. As noted in the proof in the appendix the above is also clearly true for any increasing transformation of such a V as such transformations of a utility function do not change the induced preference ordering.¹⁸

¹⁸The analysis would carry over to an even more general class of collective utility functions that take as inputs the net present values assessed by the group of individuals. Technically, for any functional $F[U_1, \dots, U_n]$, the corresponding marginal rate of substitution between periods t and $t + 1$ is given by $\frac{\sum_i \delta_i^t u'(c_{t+1}) \frac{\partial F}{\partial U_i}}{\sum_i \delta_i^{t-1} u'(c_t) \frac{\partial F}{\partial U_i}}$. As long as the sensitivity of F to different agents is not too extreme, this will also converge to weighting only the most patient agents as t grows.

The immediate implication of the proposition is that whenever a heterogeneous group of individuals or a collection of temporal motives within one individual determine choices by selecting a Pareto efficient alternative (maximizing a weighted sum of individual utility functions), a present-bias ensues.

There are additional necessary conditions that one can derive from averaging a profile of discount factors. Indeed, the result is relevant for considering representative consumers as well as for econometric estimations of preferences. As consistent with what has been noted by a number of authors following Marglin (1963) and Feldstein (1964), a weighted utilitarian function would appear to have a time-dependent exponential discount factor, where the value of the implicit discount factor at time t is given by¹⁹

$$\widehat{\delta}_t \equiv \left[\sum_i w_i \delta_i^t \right]^{1/t}.$$

Suppose $t_1 > t_2$. Given that x^{t_2/t_1} is concave, it follows from Jensen's inequality that:

$$\widehat{\delta}_{t_2} = \left[\sum_i w_i \delta_i^{t_2} \right]^{1/t_2} = \left[\sum_i w_i (\delta_i^{t_1})^{t_2/t_1} \right]^{1/t_2} \leq \left[\sum_i w_i (\delta_i^{t_1}) \right]^{1/t_1} = \widehat{\delta}_{t_1}.$$

Therefore, in concordance with the present-bias that is innate to averaging of preferences, effective exponential discount factors increase with time.²⁰ This is also consistent with the experimentally observed *sub-additivity* (see Read, 2001) in which discounting over a delay is greater when the delay is divided into subintervals than when it is left undivided.

3.1 Separability and Utilitarianism

The above results presume that the collective utility function is utilitarian in order to derive a present bias. In fact, utilitarianism is implied by weaker conditions. In fact if one simply imposes a separability condition as well as a very weak efficiency condition (unanimity), then utilitarianism, and hence 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.

¹⁹In this case, estimating an average discount rate from a distribution G on $[0, 1]$ while not accounting for heterogeneity may lead an econometrician to assess the t' th moment of the distribution, $\mathbb{E}_G(\delta^t)$, rather than the expected discount to the t' th power, $[\mathbb{E}_G(\delta)]^t$.

²⁰From an econometric perspective, notice that a homogeneous population of time inconsistent agents with quasi-hyperbolic $(\beta - \delta)$ preferences would generate a similar comparative statics.

Again, for simplicity the following result focuses on a case such that all agents, as well as the collective, share the same instantaneous utility, which has a rich range.

A condition on a collective utility function that is useful in what follows is a minimal sort of efficiency property: unanimity. It requires that if all agents prefer one stream to another, then the collective utility function should reflect that preference.

Unanimity *A collective utility function V satisfies unanimity if $V[U](C) \geq V[U](\hat{C})$ whenever*

$$\sum_t \delta_t^i u_i(c_t) \geq \sum_t \delta_t^i u_i(\hat{c}_t) \quad \text{for all } i,$$

and where the first inequality is strict whenever the second is strict for all i .

PROPOSITION 2 *For any profile $(\delta_1, u; \dots; \delta_n, u) \in \mathcal{U}^n$ such that for some k, j , $\delta_k \neq \delta_j$ and $\text{Im } u = \mathbb{R}$, a collective utility function of the form*

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

satisfied unanimity if and only if there exists weights $w_j \geq 0$, where the inequality is strict for at least one j , such that

$$V[\delta_1, u; \dots; \delta_n, u](C) = \sum_i w_i U_i(C).$$

In particular, the collective decisions are either dictatorial or present-biased.

The richness of the range allows us to invoke unanimity and gain the utilitarian presentation of social preferences. Proposition 1 then implies present-bias of the social planner whenever not dictatorial. The proposition encompasses many of the formulations of time inconsistent preferences (e.g., hyperbolic, under which $\delta_t = \frac{1}{a+bt}$, or quasi-hyperbolic, corresponding to $\delta_1 = 1$ and $\delta_t = \beta\delta^{t-1}$ for all $t > 1$, etc.). We note again that, in terms of individual behavior, the collective formulation may simply stand for the evaluation of consumption streams by an individual who balances several motives in their own mind. As long as behavior has a separable structure and satisfies unanimity, the proposition shows that present-bias is to be expected.

4 General Aggregation of Utility Functions

Although utilitarian aggregation of utility functions exhibits time inconsistency, and of a particular form, it is conceivable that there are other forms of aggregation that will be time consistency.

For example, would maximizing the minimum utility be time consistent? Would some measure of inequality serve as a collective utility function that was well-behaved?

In this section, we address the question of whether there exists any collective utility functions that are time consistent in societies where there is some heterogeneity in agents' discount rates.

Our definition of “standard utility,” or *time consistency*, relates to conditions from Koopmans (1960).

We use two pieces of notation.

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

Given $C, C' \in [0, 1]^\infty$, let $(C|_t C')$ denote the stream that consists of consumption C_t up to time t and then C'_t thereafter. So, $(C|_t C')_\tau = C_\tau$ for $\tau \leq t$ and $(C|_t C')_\tau = C'_\tau$ for $\tau > t$.

Time Consistency *The utility function V is time consistent if, for any society $U = (\delta_1, u_1; \dots, \delta_n, u_n)$, for all streams $C, \bar{C}, \hat{C}, \tilde{C}$, and times $0 \leq t < t' \leq \infty$:*

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

Time consistency essentially imposes two types of 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 consumptions are the same between the two consumption streams.²¹ It is important to note that 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 $V(\hat{C}|_{t'} C) > V(\hat{C}|_{t'} \bar{C})$ if and only if $V(\tilde{C}|_{t'} C) > V(\tilde{C}|_{t'} \bar{C})$, which is much of the essence of the second condition.

Results from Koopmans (1960) imply that, whenever V is sufficiently well-behaved, time consistency is tantamount to the maximization of a standard discounted utility function.²²

²¹There 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.

²²We define continuity and differentiability using the sup metric $d(C, \bar{C}) = \sup_t |c_t - \bar{c}_t|$.

THEOREM 1 [*Koopmans (1960)*] *A continuous utility V is time consistent if and only if there exist a discount factor $\delta \in [0, 1]$ and a continuous u such that,*

$$V(C) = \sum_t \delta^t u(c_t) \text{ for all } C.$$

The precise adaptation of Koopmans' (1960) results to our setting appears in the appendix. We note that Theorem 1 implies that our assumptions on individuals' preferences could be equivalently presented as continuity and time consistency.²³

We now state our first main result. If there is some fundamental heterogeneity in temporal preferences by way of differing discount factors, then the only well-behaved collective utility functions that are both 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). Formally,

THEOREM 2 *A collective utility function is unanimous, twice continuously differentiable,²⁴ and time consistent only if there exists i and an increasing and twice continuously differentiable u such that*

$$V[\delta_1, u_1; \dots; \delta_n, u_n](C) = \sum_t \delta_i^{t-1} u(c_t). \quad (3)$$

Moreover, a collective utility function is time consistent and unanimous at a profile $(\delta_1, u_1; \dots; \delta_n, u_n) \in \mathcal{U}^n$ for which $\delta_j = \delta_k$ implies that u_j is an affine transformation of u_k for any j, k if and only if it is dictatorial.²⁵

The theorem states that in order to be time consistent and unanimous, the collective utility function must be a time discounted sum of evaluations of the consumption 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 , and so if agents are differentiated by their discount factors then the collective utility function must be dictatorial. Alternatively, if a collective utility function responds non-trivially to at least two agents with differing time preferences and also respects unanimity, then it must be time inconsistent.

²³We present individual preferences using specific discount factors and instantaneous utility functions in order to highlight the effects of heterogeneity in time preferences, as captured by differences in discount factors.

²⁴Differentiability of the collective utility is defined using the sup metric $d(C, \bar{C}) = \sup_t |c_t - \bar{c}_t|$.

²⁵That is, u is an affine transformation of u_i for the individual i corresponding to (3).

In view of common impossibility results a-là 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 different preference profiles to involve different dictators. Nonetheless, the important implication is that if more than one agent's preferences are paid attention to at a time, then a society must be time inconsistent.

Note that instantaneous utility functions can be thought of as indirect utility functions of per-period wealth that is then divided to various private and public consumptions. In fact we show that this result extends to general multi-dimensional consumption vectors in Section 8.3. In some settings, consumption streams can also be thought of as resulting from bargaining among the individuals comprising the group. In that case, whenever outcomes can be rationalized by a collective utility function,²⁶ the theorem implies that the function cannot be simultaneously time consistent, Pareto efficient, and non-dictatorial. In particular, if one is to design bargaining protocols resulting in non-dictatorial choices, the use of commitment tools may be necessitated. In terms of commitment to the consumption choices themselves, the setup we consider is one in which overall choices of consumption streams are observed. This could fit a setting in which individuals commit to their consumption streams at the outset. It could also correspond to settings in which some decisions are overturned over time, as long as they respect unanimity. In that case, the theorem illustrates that when observing ultimate choices, they will 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' dis-

²⁶Even in settings where consumption is bargained over, and not rationalizable, we still are interested in evaluating it from a planner's perspective. Our results then imply that such a planner (presuming Pareto optimality) cannot take more than one agent's preferences into account and still be time consistent. Thus, a planner will reverse its views of what is optimal over time.

count 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 more difficult case is when the collective discount factor is in between 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 ranks it lower, in contradiction to unanimity.

A simple example illustrates the essence of how such a construction works.²⁷ 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 $\delta_{avg} = 1/2$. Suppose all agents have linear utility functions and that $C = (x, x, x, 0, 0, \dots)$ and $C' = (x + \varepsilon, x - 6\varepsilon, x + 6\varepsilon, 0, 0, \dots)$. Here, $U_1(C) = x < U_1(C') = x + \varepsilon$ and $U_2(C) = 3x < U_2(C') = 3x + \varepsilon$ so that both agents prefer C' to C . However, $U_{\delta_{avg}}(C) = 1.75x > U_{\delta_{avg}}(C') = 1.75x - .5\varepsilon$.

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

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. 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.²⁸

²⁷We use extreme values of discount factors for illustrative purposes, but as the proof shows this can be done for any set of discount factors.

²⁸This construction is a bit more involved and requires positive consumption in at least 5 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 5 periods. It may be possible to lower it to as few as three periods, although the

The heterogeneity in discount factors is critical to the results, and so it is heterogeneity in time preference 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, $\delta_1 = \dots = \delta_n \equiv \delta$, 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

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

is non-dictatorial, unanimous, and time consistent.

We stress that the result does require more than two periods of consumption. With only two periods, say the first and second 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 be a valid collective utility function also satisfying the time consistency condition (when restricted to two periods). The time consistency condition does not have much bite in such a setting and is more easily satisfied. Similarly, even when consumption can take place in arbitrarily far away periods, but consumption streams involve only a one-shot consumption, aggregation becomes less challenging. In that case, the collective compares levels of consumption at different dates. Again, using the average discount factor for the collective utility would satisfy unanimity and correspond to time consistent preferences.

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 2 implies that present-bias is to be expected.

5 Voting over Consumption Streams

Although we have shown that there is no time-consistent and unanimous manner of non-trivial aggregation of heterogeneous time preferences in the form of a collective utility function, we might also consider whether a society might come to collective decisions that are “rational” collectively, without necessarily being represented by a collective utility function.

In this section we show that any manner of making such choices that is time consistent and respects unanimity must be intransitive. In particular we consider another common way by which

proof’s details would necessarily differ.

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 some possibly weighted, non-anonymous, and/or super-majority voting rule).

Hypothetically, voting or some more general form of making binary choices might allow a representative or pivotal agent to be naturally determined. Indeed, 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 (the fairly common) simple majority rule, would it be deciding 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 are the ones that emerge from simple majority voting. As it turns out, however, 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 3 (Cycles in Collective Decisions) Consider a society composed of three individuals sharing the same instantaneous utility function $u_i(c) \equiv u(c) = c$, but having different discount factors: $\delta_1 = 0$, $\delta_2 = \frac{1}{2}$, and $\delta_3 = 1$. Consider the following three consumption streams:

$$\begin{aligned} C &= (x, x, x, 0, 0, \dots), \\ C' &= (x + \varepsilon, x - 6\varepsilon, x + 6\varepsilon, 0, 0, \dots), \\ C'' &= (x + 2\varepsilon, x - 6\varepsilon, x + 3\varepsilon, 0, 0, \dots), \end{aligned}$$

for some $\varepsilon > 0$.

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

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

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

$$U_2(C') = 1.75x - .5\varepsilon < U_2(C'') = 1.75x - .25\varepsilon < U_2(C) = 1.75x.$$

The most patient individual is concerned with the overall sum of utility functions and we have

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

If these agents were voting using majority voting, a cycle would emerge: Individuals 1 and 3 prefer C' to C , individuals 1 and 2 prefer C'' to C' , and individuals 2 and 3 prefer C to C'' .

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

5.1 Intransitivities of Majoritarian Rules

We now present results regarding majority rule since it is of particular interest both on empirical grounds, as it is observed in a wide array of applications ranging from political elections to judicial decisions, as well as on theoretical grounds, since May's (1952) theorem guarantees that majority voting is the only anonymous, neutral, and monotone choice function between two alternatives.^{29,30}

In the Supplementary Appendix we show that the same results extend to a much more general class of voting systems that also include supermajority and various weighted and non-neutral voting rules. The focus on majority rule makes the presentation simple.

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

$$CR(\delta_1, u_1; \dots; \delta_n, u_n)\widehat{C} \text{ if } \left| \left\{ i : U_i(C) \geq U_i(\widehat{C}) \right\} \right| \geq n/2.$$

In order to isolate the effects of time preference heterogeneity, we assume throughout the analysis that follows that instantaneous utility functions are identical for all agents, so that $u_i \equiv u$, where u is continuous and strictly increasing, for all i . Now, if a majority of individuals share the same discount factor, majority rule would effectively follow their choice and the existing heterogeneity

²⁹Anonymity requires that all voters be treated the same; neutrality imposes that the two alternatives be treated identically (i.e., reversing each preference reverses the collective preference); monotonicity demands that changing preferences in favor of one alternative makes it more likely to be chosen.

³⁰For consumption and saving problems, Boylan and McKelvey (1995) note the intransitivities that may arise. In this section we highlight the more general attributes of environments that generate or prohibit the emergence of voting cycles.

would play no role. We therefore concentrate on the case in which there is no majority of individuals with identical discount factors.

Without any restrictions on consumption streams, it turns out that our initial example extends directly. Even without strict concavity (as in Boylan and McKelvey, 1995) nor any restrictions on the optimal alternative for each individual, intransitivity is inherent.

PROPOSITION 3 *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 a majority, then $P(\delta_1, u; \dots; \delta_n, u)$ is intransitive.*

The proof of this proposition illustrates that any agent can be made the “pivotal voter.” To glean some intuition to the workings of the proof, let us show how to identify consumption streams C and C' , such that individuals N_1 prefer consumption stream C to C' and individuals N_2 prefer C' to C . There is a corresponding system of linear inequalities. Namely,

$$\begin{aligned} \sum_t \delta_i^{t-1} [u(c_t) - u(c'_t)] &> 0 \quad \text{for } i \in N_1 \\ \sum_t \delta_i^{t-1} [u(c_t) - u(c'_t)] &< 0 \quad \text{for } i \in N_2. \end{aligned}$$

Now, if discount factors are all different, then whenever the range of u is sufficiently rich, the linear independence of $\{(1, \delta_i, \delta_i^2, \dots)\}_i$ 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 3 are driven by the linear independence of the vectors of coefficients of the sequence of discount factors $\{(1, \delta_i, \delta_i^2, \dots)\}_i$. In fact, as long as there is enough dependence between these vectors (relative to the potential instantaneous utility functions and admissible consumptions), majority voting does not entail intransitivities. In the appendix we provide the precise condition that characterizes such settings.

6 Collective Time Preferences in the Laboratory

We conducted a set of experiments designed to elicit social preferences over joint consumption streams in the lab. In particular, we hoped to gain insight as to whether experimental social

planners will exhibit the time inconsistency predicted by Theorem 2 and also suggested by some of the empirical work identifying time inconsistencies. Furthermore, these allow us to analyze subjects' collective preferences.

6.1 A Methodological Point on Experiments with Time Preference

There are several challenges inherent in the elicitation of individual time preferences in a lab. First, explicitly delaying payments to subjects does not necessarily map into delayed changes in their *consumption* levels. Instead, the delays may simply affect their cash flow. The precise mapping into consumption levels is the outcome of a much more complicated and unobserved decision process. Indeed, delayed payments can be substituted for by shifting existing wealth across time at the subject's interest rate on savings. Therefore, measured discount factors may simply reflect the interest rates and/or borrowing constraints that subjects face.

Second, agents may be uncertain about their preferences for cash in the future. Therefore, elicitation of time preferences using delayed payments may include risk and uncertainty as confounds. This is certainly the case for subject pools composed of college students, who frequently face cash constraints and nontrivial uncertainty about their expenses and income across weeks or months.³¹

Third, in the context of our design, the possibility that some subjects are individually time inconsistent (be it as a consequence of balancing different motives in their mind, or due to a fundamental preference) poses an additional challenge for our main goal of the experiments. Our focus is on how aggregation plays into time inconsistency, and so we wish to begin with time consistent underlying preferences and measure what forms of inconsistency emerge from an aggregation decision. If underlying preferences are already time inconsistent, it would be much more difficult, if not impossible, to identify the role of aggregation. Thus, we wish to induce and control the underlying individual preferences. Effectively, imposing time-consistent individual preferences stacks the cards against our theoretical hypothesis that aggregation in and of itself can generate time inconsistencies.

In addition, subjects' organic time preferences are unpredictable from the experimenter's perspective and could potentially be rather narrow for the time horizons that are feasible in a lab setting (even with delayed payments). This would limit our ability to test aggregation of an assortment of different time preferences, which are very relevant for economic decisions that have longer

³¹In addition to this, subjects may also have some uncertainty about how they will be paid, although that sort of uncertainty can be mitigated through careful design (for instance, see Andreoni and Sprenger (2010a) and Benhabib, Bisin, and Schotter (2010)).

horizons. Thus, having control of subjects' individual base preferences is important for a variety of reasons.

These issues led us to develop a new elicitation technique, aimed at inducing and controlling individual discount factors, while maintaining the nature of the trade-offs featured in groups' collective intertemporal decisions.

Our design used a combination of two elicitation methods, which we now describe in broad terms (a detailed description follows). In the first, to represent streams of consumption over three different periods we used combinations of tokens of three different colors. Instead of a subject receiving some amount of consumption today, some amount t periods from now, and some amount t' periods from now, instead subject would get some amount of blue tokens, some amount of red tokens, and some amount of grey tokens. We then induced a discount factor by having different exchange rates for tokens into cash, depending on their color. As an example, a subject with *induced* discount factor of δ would receive one cent for each blue token, δ cents for each red token, and δ^2 cents for each grey token. For instance, a subject assigned with a discount factor of $\delta = .9$ would get 1 cent for each blue token, .9 cents for each red token, and .81 cents for each grey token. If, for example, such a subjects were faced with a choice between a stream of tokens $C = (105, 0, 0)$ of blue, red, and grey tokens respectively (so 105 blue tokens and none of the other colors), and another of $C' = (0, 160, 0)$ (so 160 red tokens), then C would be worth 105 experimental cents to the subject and C' would be worth $.9 * 160 = 144$ cents. In particular, C' would offer a greater payoff.

While this does not explicitly involve timing, it involves exactly the same calculations that we assume agents make in standard economic models, and in the experiments we can then focus on how subjects *aggregate* these preferences. This allows us to completely isolate the effects of aggregation while mimicking and controlling subjects' underlying 'time' preferences. In particular, it allows us to gain insights on the impact of wide range of discount factor combinations that, using organic time preferences of subjects, would require us to either run very large experiments (to get many combinations of discount factors), or introduce many time horizons to create substantial heterogeneities (that may be economically important); and these would still face the difficulties mentioned earlier of eliciting actual time preferences.

Under this sort of construction, the choice between C and C' becomes quite clear as the payoff values to subjects are then simple multiplications of the induced discount factors times the relevant numbers of blue, red, and grey tokens. In order to further focus on the aggregation, we also

presented the subjects with the net present values rather than the tokens and discount factors, and then checked that they made the same choices under the two different presentations. This allows us to avoid choices being driven by simple calculation errors. In particular, to check that these two different presentations were functionally equivalent we ran, as a control, several complementary sessions in which choices were presented as explicit ‘time’ (token) streams with discount factors. Subjects made choices that were statistically indistinguishable across the two types of experimental framings. A precise description of these comparisons appears in the supplementary appendix.

6.2 Details of the Design

Subjects made a series of decisions over two alternative consumption streams, such as the streams C and C’ described above. Every subject made a decision in each round. The, each determining a different payoff profile for a group of three subjects, call them Member 1, 2, and 3. The choice affected someone in the lab other than themselves. Namely, at the end of each round, we randomly matched subjects in pairs. Suppose subject A was paired with subject B. We randomly assigned a role for each as Member 1, 2, or 3. Subject A, assigned the role of Member x , would then be paid according to what Member x would be provided according to the present value Member x would receive according to subject B’s allocation choice. Similarly, subject B, assigned the role of Member y , would be paid according to Member y ’s present value under subject A’s selected allocation. Partners were not transparent to subjects and were randomly re-assigned at the end of each round. Cumulative payoffs were reported after each round.

Subjects interacted only through the computerized interface. Two practice rounds were followed by a number of rounds of actual choice (actual choices are listed in Appendix B).³² Each of the 38 choices subjects faced was comprised of three payoffs, to the three potential Members. When comparing two such alternatives, there are three dimensions of the alternatives that may conceivably play an important role in choice: the sum of payoffs, the distribution of payoffs, and the marginal differences between alternatives for each of the subjects. We designed the set of payoffs to vary in these three dimensions.

The experiments were conducted at the California Social Sciences Experimental Laboratory (CASSEL) at UCLA with 60 subjects participating in six separate sessions. Subjects were paid the sum of payoffs throughout the rounds, averaging \$39, in addition to a show-up fee of \$5.

³²The actual ordering of the choices that subjects faced was randomized to eliminate any systematic framing effects. Our full instructions are available at: <http://...>

6.3 Time Inconsistency

Several choices subjects faced corresponded to choices between two consumption streams and their delayed version. For instance, suppose group member 1 was assigned a discount factor of .2 and group member 2 a discount factor of .9. First, consider a choice between the streams $C = (105, 0, 0)$ and $C' = (0, 160, 0)$. The present discounted values of C and C' for these two members would be $U_1(C) = U_2(C) = 105$, and $U_1(C') = 32$ and $U_2(C') = 144$. The delayed decision is then between streams $C'' = (0, 105, 0)$ versus $C''' = (0, 0, 160)$, which induce present discounted utilities of $U_1(C'') = 21$ and $U_2(C'') = 95$, and $U_1(C''') = 6$ and $U_2(C''') = 130$. In each group we also had a third member whose payoff was often the same across the two choices, in this case 80 for all four choices.³³ Thus, in terms of profiles of present discounted utilities corresponding to the streams these become $(105, 105, 80)$ for C and $(32, 144, 80)$ for C' ; and $(21, 95, 80)$ for C'' and $(6, 130, 80)$ for C''' . In fact, this pair of choices corresponded to choices 1 and 2 in our experiments. A subject would be present-biased just with respect to these two choices if he or she picked C over C' and then C''' over C'' (as more than a third of the subjects did). Similarly, choices 3-6 corresponded to similar pairs in which we considered different initial wealth for the members 1 and 2 that are affected by the planner's selection (see Appendix B).

These choices allow us to distinguish time consistent planners (who make corresponding selections in the original and delayed choice problems), present-biased planners (who sometimes choose the more immediate consumption initially, but then the more delayed consumption when both choices are delayed), future-biased planners (picking delayed consumptions initially, but more immediate rewards when the both choices are delayed), and planners who do not fall squarely into any of the above categories.

Figure 1 depicts the distribution of individuals based on this classification. A large majority, 75% of individuals, are present-biased, while less than 2% (only one subject) were time consistent, and the remaining either future-biased or showing mixed inconsistencies. These results are very unlikely to arise from subjects randomly making selections. Indeed, Figure 1 also depicts the expected distribution of types were individuals selecting each choice in a pair with a probability of 0.5, which is significantly different from the observed distribution with any reasonable levels of statistical confidence.³⁴

³³The third member was included since we also ran some experiments related to groups of three voting over choices, to check on intransitivities. The intransitivities appeared quite prevalently and almost exactly in line with the theory, and so for the sake of space we do not report them here; but we kept with three members per group in the “planner” treatments in order to keep the interface similar across the treatments.

³⁴Furthermore, the probability of achieving at least 75% present-biased individuals under random choice is of the

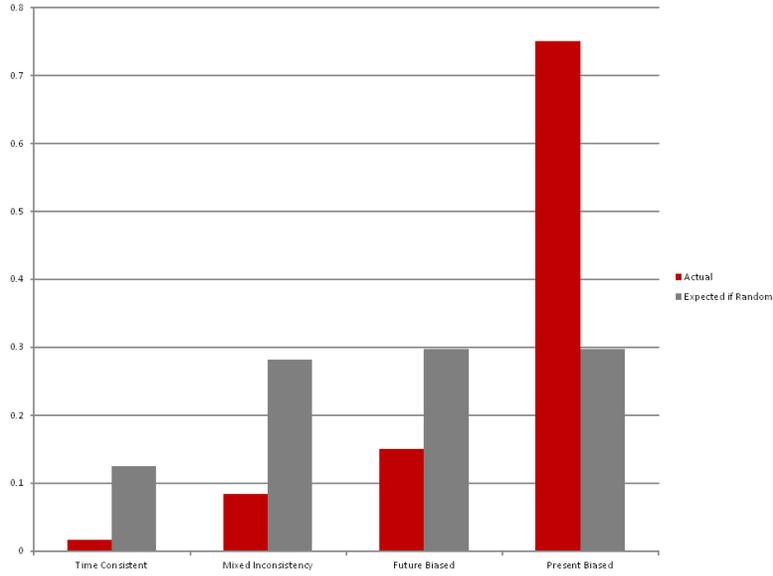


Figure 1: Frequency of Individuals of Different Time Consistency Types.

6.4 Social Planners’ Collective Utility Functions

As mentioned, the choices we offered subjects were designed to distinguish the ways in which planners made choices, including pairs of consumption streams that varied in the sum of payoffs, the distribution of payoffs, and the marginal differences between alternatives.

Specifically, we designed the choice problems to distinguish between the three most prominent ways from the welfare literature in which planners might evaluate alternatives: based on the sum of utilities across agents (utilitarianism), based on maximizing the minimum utility (maximin or a “Rawlsian” approach), or based on some weighting of the distribution of payoffs (some form of egalitarianism), say the minimization of the standard deviation of payoffs.³⁵

Figure 2 illustrates how the choices faced by subjects distinguish these motives. Figure 2 contains four choices that differed in terms of which selection a subject would make if they were using a different sort of collective utility function. the type of objective that would lead to a selection of either alternative and the corresponding experimental frequencies with which the first alternative was selected.³⁶

magnitude of 10^{-28} .

³⁵The analysis that follows does not change qualitatively if we consider egalitarian motives that depend on the variance or absolute differences of payoffs.

³⁶These choices correspond to choices 1, 2, 6, and 12 in our design, see Appendix B.

		Profile of Payoffs	Types of Planner Preferences Preferring	% Choosing
Choice 1	C	(105, 105, 80)	M, I, U	75%
	C'	(32,144,80)		25%
Choice 2	C	(21, 95, 80)	M, I U	54%
	C'	(6, 130, 80)		46%
Choice 3	C	(61, 55, 80)	M I, U	21%
	C'	(46, 90, 80)		79%
Choice 4	C	(21,95,80)	M, I, U	17%
	C'	(26,100,80)		83%
		(U1, U2, U3)	M=Maximin I=Inequality U=Utilitarian	

Figure 2: Choices Made that Distinguish Between Three Prominent Collective Utility Functions

The table illustrates that when all three objectives lead to the same alternatives (the first and fourth choice pairs), then most of the subjects followed those prescriptions (75 and 83 percent, respectively). In the case where the prescriptions under utilitarianism and egalitarianism diverge (Choice 2), there is a substantial split in the decisions of the subjects. However, the choice that differed only in the prescription of maximin (Choice 3) compared to utilitarianism and inequality aversion, the pattern looks (statistically) indistinguishable from the case where all three lead to the same alternative (Choice 4). Indeed, as the bottom two choices of Figure 2 illustrate, the choices in which both utilitarian and egalitarian motives would push agents toward selecting the second alternative, but maximin motives would generate diverse decisions, differed in only four percent of subjects.³⁷ Thus, in what follows, we consider collective utility functions that depend on the sum and distribution of payoffs.

A family of collective utility functions that incorporates utilitarianism and inequality aversion is one that evaluates an alternative according to:

$$a * \text{Sum of Payoffs} - (1 - a) * \text{Standard Deviation of Payoffs}, \quad (4)$$

³⁷Aggregate discrete-choice regression analysis in which selections were explained by the sum of payoffs, the variance of payoffs, and the maximin of payoffs, led to similar observations regarding the relative importance of the sum and variance of payoffs.

where $a \in [0, 1]$. Thus, a serves as a weight of how much a planner cares about utilitarian concerns relative to distributional ones.³⁸

For each subject and for each specification of the parameter a , we can calculate the fraction of decisions that would be implied by the maximization of (4) and coincide with that subject's observed selections. We call that fraction the *individual score*. For each a the overall *score* is the average over all individual subject scores.

We allow for heterogeneity in the a 's that best describe subjects. In particular, we considered different numbers of "types" to describe the agents. A type is characterized by a range of parameters a - that all lead to the same predictions (given the finite number of alternatives, not all a 's are distinguished). So, if we just allow for one type, then we find the single a that best fits all of the decisions by all of the agents. If we allow for two types, then we look for two different a 's that best describe the decisions of all of the agents, when each agent must be assigned to one of the two different a 's and so forth. So, for each number of types k , we look for a partition of the subject population into k groups and a value of a for each group in the partition that maximized overall scores. Figure 3 presents the values of the best fitting (ranges of) a 's for each number of types k , and the corresponding average score for individuals of each type.

In terms of preference parameters, assuming all subjects have the same collective utility function leads to an estimated parameter a of 0.75. However, allowing for more than one type of collective utility function among the subjects leads to a substantial fraction of individuals (between 15% – 25%) who put very little weight on utilitarian concerns and care mostly about the variation in payoffs. Regardless of the number of types fitted, more than three quarters of the population has a taste parameter a that is .7 or higher. Thus, although most subjects are best fit by having some weight on inequality, many weight it in a minor way; although a minority of subjects appear to weight inequality much more substantially. Nonetheless, we note that the types with lower utilitarian incentives were characterized with a lower overall score, suggesting that our model of planners' objectives provides a better fit for individuals that place greater weight on utilitarian motives.³⁹

In Figure 4, we present the best fitting score for each type, as dependent on the number of

³⁸As mentioned before, the results are quantitatively similar when using the sum of absolute differences between payoffs instead of the standard deviation of payoffs, and both offer slightly better fits than those derived by using variance instead of standard deviation, which is more nonlinear in the level of inequality. The full comparisons appear in the supplementary appendix.

³⁹The score is still substantially higher than 50% that would be the expected score generated by random selections, or 0% that would be generated were the model fully ill-specified. It is also possible that some subjects had objectives that accounted for multiple choices made, rather than viewing each choice on the margin. Although much of the variation in behavior is captured with objective functions that look at each choice in isolation, that possibility is an important consideration for further research.

Number of Types		Type 1	Type 2	Type 3	Type 4	Type 5	Type 6
1	Best fit 'a' range % of this type Score	0.75 (100%) 0.72					
2	Best fit 'a' range % of this type Score	0-0.25 (25%) 0.67	0.95-1 (75%) 0.80				
3	Best fit 'a' range % of this type Score	0-0.25 (21%) 0.66	0.75 (40%) 0.79	0.95-1 (39%) 0.86			
4	Best fit 'a' range % of this type Score	0-0.25 (17%) 0.66	0.7 (18%) 0.82	0.75 (27%) 0.81	0.95-1 (38%) 0.86		
5	Best fit 'a' range % of this type Score	0-0.25 (15%) 0.6449	0.55 (7%) 0.7968	0.7 (14%) 0.838	0.75 (26%) 0.814	0.95-1 (38%) 0.8653	
6	Best fit 'a' range % of this type Score	0-0.25 (15%) 0.64	0.55 (7%) 0.80	0.7 (13%) 0.84	0.75 (25%) 0.82	0.85 (17%) 0.88	0.95-1 (23%) 0.85

Figure 3: Table of Best Fitting Utilitarian and Inequality Preferences By Number of Types

types. When the possible number of types is 60 (the number of subjects), each subject could have an individual a , and so the derived score is the maximal feasible one under this model and is 82%. As can be seen in the figure, allowing for 5 or 6 types yields scores that are not substantially lower. On the other extreme, assuming the population is completely homogeneous, leads to a lower score of approximately 72%.

To conclude, our experimental results are in line with Theorem 2 in that, with almost no exceptions (fifty nine out of sixty) experimental planners are time-inconsistent, with three quarters of them exhibiting choices that are present-biased. Furthermore, planners in these experiments generally act in ways that are consistent with a combination of utilitarian and equality-sensitive motives, with a large fraction putting a rather substantial weight on the utilitarian attributes of alternatives. Recall that Proposition 1 implied that whenever utilitarian motives guide decisions, behavior will be present-biased. In that sense, the preference characterization of our experimental social planners is in line with the prevalent present-bias observed in the lab.

7 Concluding Remarks

A main message of this paper is that the aggregation of heterogeneous temporal preferences in a non-dictatorial manner that respects unanimity is bound to exhibit time inconsistencies or intran-

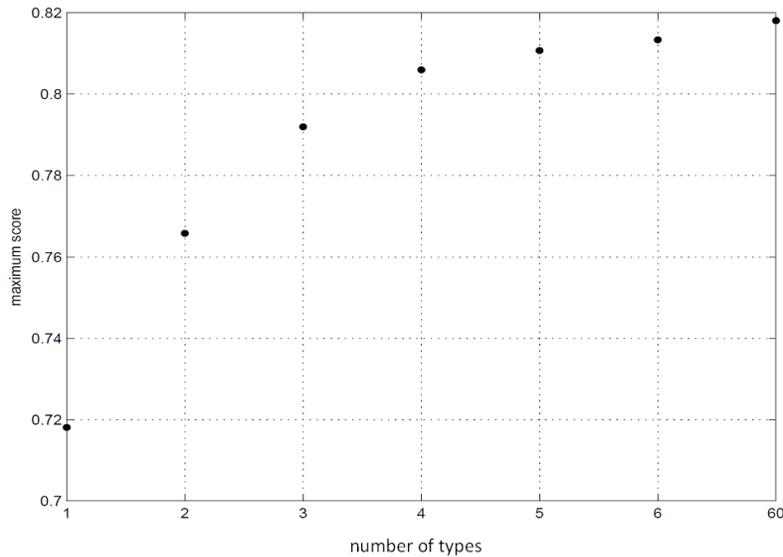


Figure 4: Scores as a Function of the Number of Types.

sitivities. 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.

Beyond the general theorems, we also analyzed two classes of specific aggregation methods that are commonly utilized. One class, utilitarian aggregation rules that operate according to the weighted sum of utility functions, is transitive but time inconsistent. In fact, whenever there is some heterogeneity in the population, utilitarian aggregation exhibits present-bias, where choices appear to correspond to increasing patience over time. This type of time inconsistency matches a large body of experimental and empirical work on time preferences. The other specific class of methods, binary voting rules, in which choices are based only on the number of individuals that prefer one alternative relative to another, are time consistent but intransitive. Indeed, such rules exhibit strict voting cycles even when restricted to sets of alternatives that correspond to consumption smoothing problems over very short horizons.

The results are potentially important for policy making when heterogeneous temporal preferences are present in the population. Marglin (1963) and Feldstein (1964) suggested that choosing a sensible representative agent may involve non-stationary discount rates, and recent work has been trying to look at 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 consid-

erations are unavoidable. Policy makers who trade-off different temporal preferences of individuals in any non-trivial way will de-facto be facing a time inconsistent representative agent. In fact, if policy makers care about some proxy of (utilitarian) efficiency, they will be facing 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 that are different from the collective's, potentially time consistent, and so welfare maximization requires a careful analysis with the primitive preferences taken into account, and not simply substituted by a non-existent 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 preferences of only one individual. The precise characterization of outcomes generated by such protocols is likely to require new tools and techniques.⁴⁰

Finally, we note that our results extend to much more general classes of consumption streams, as shown in the supplementary appendix. There, we also identify restrictions on consumption streams that allow for transitive and time consistent aggregation: namely a well-ordering condition that requires that any two admissible streams only differ with regards to shifting consumption in one direction.

References

- [1] Abdellaoui, M., O. l'Haridon, C. Paraschiv (2010) "Individual vs Collective Behavior: An Experimental Investigation of Risk and Time Preferences in Couples," mimeo: GREGHEC-CNRS
- [2] Ainslie, G. (1975), "Specious Reward: A Behavioral Theory of Impulsiveness and Impulse Control," *Psychological Bulletin*, Volume 82(3), pages 463-496.
- [3] Ambrus, A. and K. Rozen (2009), "Rationalizing Choice with Multi-self Models," mimeo.
- [4] Amador, M., I. Werning, and G. M. Angeletos (2006) "Commitment versus Flexibility," *Econometrica*, Volume 74(2), pages 365-396.
- [5] Andreoni, J. and C. Sprenger (2010a), "Estimating Time Preferences from Convex Budget Sets," mimeo.
- [6] Andreoni, J. and C. Sprenger (2010b), "Risk Preferences are Not Time Preferences," mimeo.

⁴⁰Indeed, 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, for instance, by Lehrer and Pauzner, 1999).

- [7] Arrow, K. J. (1950), "A Difficulty in the Concept of Social Welfare," *The Journal of Political Economy*, Volume 58(4), pages 328-346.
- [8] Arrow, K. J. (1963), *Social Choice and Individual Values*, Second Edition, Cowles Foundation for Research in Economics: Yale University.
- [9] Barbera, S. and M.O. Jackson (2006) "On the Weights of Nations: Assigning Voting Power to Heterogeneous Voters," *Journal of Political Economy*, Vol 114, No. 2, pp 317-339.
- [10] Benhabib, J., A. Bisin, and A. Schotter (2010), "Present-Bias, Quasi-Hyperbolic Discounting, and Fixed Costs," *Games and Economic Behavior*, forthcoming.
- [11] Bernheim, B. D. (1999) "Comment on 'Family Bargaining and Retirement Behavior'," in Henry Aaron (ed.), *Behavioral Economics and Retirement Policy*, Brookings Institution Press, 273 - 281.
- [12] Bernheim, B. D. and A. Rangel (2004), "Addiction and Cue-Triggered Decision Processes," *The American Economic Review*, Volume 94(5), pages 1558-1590.
- [13] Beshears, J., J. Choi, D. Laibson, and B. Madrian (2008), "The Importance of Default Options for Retirement Saving Outcomes: Evidence from the United States," *Lessons from Pension Reform in the Americas*. Oxford: Oxford University Press. Kay, S. J. and T. Sinha, editors.
- [14] Blackorby, C., W. Bossert, and D. Donaldson (2005), "Temporal Consistency," In *Population Issues in Social Choice Theory, Welfare Economics, and Ethics* (pages 272-285). Cambridge: Cambridge University Press.
- [15] Boylan, R. T. and R. D. McKelvey (1995), "Voting over Economic Plans," *The American Economic Review*, Volume 85(4), pages 860-871.
- [16] Brocas, I. and J. Carrillo (2008), "The Brain as a Hierarchical Organization," *The American Economic Review*, Volume 98(4), pages 1312-1346.
- [17] Browning, M. (2000), "The Saving Behaviour of a Two-person Household," *Scandinavian Journal of Economics*, Volume 102(2), pages 235-251.
- [18] Caplin, A. and J. Leahy (2004), "The Social Discount Rate," *The Journal of Political Economy*, Volume 112(6), pages 1257-1268.
- [19] Cherepanov, V., T. Feddersen, and A. Sandroni (2009), "Rationalization," mimeo.
- [20] Condorcet, Marquis de. (1785), *Essai sur l'application de l'analyse a la probabillite des decisions rendues a la probabillite des voix*, Paris: De l'imprimerie royale, translated in 1976 to "Essay on the Application of Mathematics to the Theory of Decision-Making." in Condorcet: Selected Writings, Indianapolis, IN: Bobbs-Merrill. K. M. Baker, editor.
- [21] Dasgupta, P. and E. Maskin (2004), "Uncertainty and Hyperbolic Discounting," mimeo.
- [22] Della Vigna, S. and U. Malmendier (2006), "Paying Not To go To The Gym," *The American Economic Review*, Volume 96, pages 694-719.
- [23] Drefahl, S. (2010) "How Does the Age Gap Between Partners Affect Their Survival?," *Demography*, Volume 47(2), pages 313-326.
- [24] Engelmann, D. and M. Strobel (2004) "Inequality Aversion, Efficiency, and Maximin Preferences in Simple Distribution Experiments," *The American Economic Review*, Volume 94(4), pages 857-869.
- [25] Evreny, O. and E. A. Ok (2008), "On the Multi-Utility Representation of Preference Relations," mimeo.
- [26] Farmer, J. D. and J. Geanakoplos (2009). "Hyperbolic Discounting is Rational: Valuing the Far Future with Uncertain Discount Rates," mimeo.

- [27] Feldstein, M. S. (1964), "The Social Time Preference Discount Rate in Cost Benefit Analysis," *The Economic Journal*, Volume 74(294), pages 360-379.
- [28] Frederick, S., G. Loewenstein, and T. O'Donoghue (2002), "Time Discounting and Time Preference: A Critical Review," *Journal of Economic Literature*, Volume 40(2), pages 351-401.
- [29] Fudenberg, D. and D. K. Levine (2006), "A Dual Self Model of Impulse Control," *The American Economic Review*, Volume 96, pages 1449-1476.
- [30] Gollier, C. and R. Zeckhauser (2005), "Aggregation of Different Time Preferences," *The Journal of Political Economy*, Volume 113(4), pages 878-896.
- [31] Green, J. R. and D. A. Hojman (2009), "Choice, Rationality and Welfare Measurement," mimeo.
- [32] Hare, T. A., C. F. Camerer, and A. Rangel (2009), "Self-Control in Decision-Making Involves Modulation of the vmPFC Valuation System," *Science*, Volume 324, pages 646-648.
- [33] Herrnstein, R. J. (1961), "Relative and absolute strength of response as a function of frequency of reinforcement," *Journal of the Experimental Analysis of Behavior*, Volume 4(3), pages 267-272.
- [34] Hertzberg, A. (2010), "Exponential Individuals, Hyperbolic Households," mimeo.
- [35] Jamison, D.T. and J. Jamison (2007), "The Amount and Speed of Discounting," Disease Control Priorities Project Working Paper No. 4.
- [36] Koopmans, T. C. (1960), "Stationary ordinal utility and impatience," *Econometrica*, Volume 28(2), pages 287-309.
- [37] Kydland, F. E. and E. C. Prescott (1977), "Rules Rather than Discretion: The Inconsistency of Optimal Plans," *The Journal of Political Economy*, Volume 85(3), pages 473-492.
- [38] Laibson, D. (1997), "Golden Eggs and Hyperbolic Discounting," *The Quarterly Journal of Economics*, Volume 62, pages 443-477.
- [39] Lehrer, E. and A. Pauzner (1999), "Repeated Games with Differential Time Preferences," *Econometrica*, Volume 67, pages 393-412.
- [40] Marglin, S. (1963), "The Social Rate of Discount and The Optimal Rate of Investment," *The Quarterly Journal of Economics*, Volume 77(1), pages 95-111.
- [41] May, K. O. (1952), "A set of independent necessary and sufficient conditions for simple majority decisions," *Econometrica*, Volume 20(4), pages 680-684.
- [42] Mazzocco, M. (2007) "Household Intertemporal Behaviour: A Collective Characterization and a Test of Commitment," *The Review of Economic Studies*, Volume 74, pages 857-895.
- [43] McClure, S. M., D. I. Laibson, G. Loewenstein, and J. D. Cohen (2004), "Separate Neural Systems Value Immediate and Delayed Monetary Rewards," *Science*, Volume 306, pages 503-507.
- [44] McClure, S. M., K.M. Ericson, D. I. Laibson, G. Loewenstein, and J. D. Cohen (2007), "Time Discounting for Primary Rewards," *The Journal of Neuroscience*, Volume 27(21), pages 5796-5804.
- [45] McKelvey, R. D. (1976), "Intransitivities in Multidimensional Voting Models and some Implications for Agenda Control," *Journal of Economic Theory*, Volume 12, pages 472-482.
- [46] McKelvey, R. D. (1979), "General Conditions for Global Intransitivities in Formal Voting Models," *Econometrica*, Volume 47, pages 1085-1112.
- [47] Mongin, P. (1995) "Consistent Bayesian Aggregation," *Journal of Economic Theory*, Volume 66(2), pages 313-351.

- [48] O’Donoghue, T. and M. Rabin, (1999), “Doing it Now or Later,” *The American Economic Review*, Volume 89(1), pages 103-124.
- [49] Parks, R. P. (1976), “An Impossibility Theorem for Fixed Preferences: A Dictatorial Bergson-Samuelson Welfare Function,” *The Review of Economic Studies*, Volume 43(3), pages 447-450.
- [50] Plott, C. (1967), “A Notion of Equilibrium and its Possibility under Majority Rule,” *The American Economic Review*, Volume 57, pages 787–806.
- [51] Rachlin, H. (2000), *The Science of Self-Control*, Cambridge, MA: Harvard University Press.
- [52] Read, D. (2001), “Is Time Discounting Hyperbolic or Subadditive?,” *Journal of Risk and Uncertainty*, Volume 23(1), pages 5-32.
- [53] Roberts, K. W. S. (1980), “Social Choice Theory: The Single-Profile and Multi-Profile Approaches,” *The Review of Economic Studies*, Volume 47(2), pages 441-450.
- [54] Schaner, S. G. (2010), “Intrahousehold Preference Heterogeneity, Commitment, and Strategic Savings: Theory and Evidence from Western Kenya,” mimeo.
- [55] Shafir, S. (1994), “Intransitivity of Preferences in Honey Bees: Support for ‘Comparative’ Evaluation of Foraging Options,” *Animal Behavior*, Volume 48, pages 55-67.
- [56] Sozou, P. D. (1998), “Hyperbolic Discounting and Uncertain Hazard Rates,” *Proc. R. Soc. Lond. B*, Volume 265, pages 2015-2020.
- [57] Strotz, R. H. (1955), “Myopia and Inconsistency in Dynamic Utility Maximization,” *The Review of Economic Studies*, Volume 23(3), pages 165-180.
- [58] Thaler, R. H. (1981), “Some Empirical Evidence on Dynamic Inconsistency,” *Economics Letters*, 201-207.
- [59] Thaler, R. H. and H. M. Shefrin (1981), “An Economic Theory of Self-Control,” *The Journal of Political Economy*, Volume 89(2), pages 392 - 406.
- [60] Tversky, A. (1969), “Intransitivity of Preference,” *Psychological Review*, Volume (76), pages 31-48.
- [61] Waite, T. A. (2001), “Intransitive preferences in hoarding gray jays (*Perisoreus canadensis*),” *Behavioral Ecology and Sociobiology*, Volume (50), pages 116-121.
- [62] Weitzman, M. L. (2001) “Gamma Discounting,” *The American Economic Review*, Volume 91(1), pages 260-271.
- [63] Xue, L. (2008) “The Bargaining Within,” *Economics Letters* 101, pages 145-147.
- [64] Zuber, S. (2010), “The aggregation of preferences: can we ignore the past?” mimeo.

8 Appendix A – Proofs

Without loss of generality, we normalize all utility functions so that $u_i(0) = 0$ for all i .

We prove a stronger version of Proposition 1.

A collective utility function is a *generalized weighted utilitarian function* if there exist weights $w_i \in [0, 1]$ such that $\sum_i w_i = 1$ and

$$V(C) = F \left(\sum_i w_i U_i(C) \right),$$

where F is an increasing function.

PROPOSITION 1* *If $u_i = u$ for all i where u is continuous and strictly increasing, and V is a generalized weighted utilitarian function with weights $w_i > 0$ and $w_j > 0$ for at least two agents i and j such that $\delta_i \neq \delta_j$, then V is present-biased.*

Proof of Proposition 1*: By the definition of a generalized weighted utilitarian function and the conditions of the proposition, we can write

$$V(C[x, t]) = F \left(\sum_i w_i \delta_i^{t-1} u(x) \right).$$

Thus, $V(C[x, t]) \geq V(C[y, t+k])$ implies that

$$\sum_i w_i \delta_i^{t-1} u(x) \geq \sum_i w_i \delta_i^{t+k-1} u(y).$$

If $u(x) = 0$ then $u(y) = 0$ and from monotonicity, $x = y = 0$ and all time evaluations are identical. Suppose then that $u(x) > 0$. The above inequality can be written as

$$\frac{\sum_i w_i \delta_i^{t-1}}{\sum_i w_i \delta_i^{t+k-1}} \geq \frac{u(y)}{u(x)}.$$

Note that

$$\frac{\sum_i w_i \delta_i^{t-1}}{\sum_i w_i \delta_i^{t+k-1}}$$

is decreasing in t since $\delta_i \in (0, 1)$ for all i . In order to illustrate the first requirement of present-bias it is enough to show that it is strictly decreasing for $k = 1$. Therefore, we now show that

$$\frac{\sum_i w_i \delta_i^{t-1}}{\sum_i w_i \delta_i^t}$$

is strictly decreasing in t .

Note that the derivative of this expression is

$$\frac{(\sum_i w_i \ln(\delta_i) \delta_i^{t-1}) (\sum_i w_i \delta_i^t) - (\sum_i w_i \ln(\delta_i) \delta_i^t) (\sum_i w_i \delta_i^{t-1})}{(\sum_i w_i \delta_i^t)^2}.$$

The numerator can be written as

$$\sum_i \sum_j \ln(\delta_i) w_i w_j \delta_i^{t-1} \delta_j^{t-1} (\delta_j - \delta_i),$$

which we can rewrite as

$$\sum_{i,j:i < j} w_i w_j \delta_i^{t-1} \delta_j^{t-1} (\ln(\delta_i) - \ln(\delta_j)) (\delta_j - \delta_i),$$

and each of the expressions in this summand is positive whenever $w_i w_j \delta_i^{t-1} \delta_j^{t-1} \neq 0$ and $\delta_j \neq \delta_i$, and is zero otherwise.

This implies that

$$\frac{\sum_i w_i \delta_i^{t-2}}{\sum_i w_i \delta_i^{t-1}} \geq \frac{u(y)}{u(x)},$$

or that $V(C[x, t-1]) \geq V(C[y, t])$ and so the first part of the definition of present-bias is satisfied.

Next, given that u is continuous and increasing, and $u(0) = 0$ (recall that all utility functions are so normalized), we can find x and y such that

$$u(x) = \sum_i w_i \delta_i^k u(y).$$

Given that $w_i > 0$ and $w_j > 0$ for some $\delta_i \neq \delta_j$, it follows easily that

$$\frac{\sum_i w_i \delta_i^t}{\sum_i w_i \delta_i^{t+1}}$$

is strictly decreasing in t . Therefore, iterative application of the above implies that

$$\sum_i w_i \delta_i^t u(x) < \sum_i w_i \delta_i^{t+k} u(y)$$

for all $t \geq 1$. By continuity of u , we can then find some ε such that

$$u(x + \varepsilon) > \sum_i w_i \delta_i^k u(y)$$

and

$$\sum_i w_i \delta_i^t u(x + \varepsilon) < \sum_i w_i \delta_i^{t+k} u(y)$$

for all $t \geq 1$. This establishes the second part of the definition of present-bias. ■

Proof of Proposition 2: Suppose $\delta_1, \dots, \delta_K$ is the set of distinct discount factors among $\delta_1, \dots, \delta_n$. It follows that $\{(1, \delta_i, \delta_i^2, \dots)\}_{i=1}^K$ are linearly independent. We now show that $\{(1, \delta_i, \delta_i^2, \dots)\}_{i=1}^K \cup \{(\tilde{\delta}_1, \tilde{\delta}_2, \tilde{\delta}_3, \dots)\}$ are linearly dependent. Indeed, suppose the contrary and consider the following set of inequalities.

$$\begin{aligned} \sum_t \delta_i^{t-1} [u(c_t) - u(c'_t)] &> 0 \quad \text{for } i = 1, \dots, K, \\ \sum_t \tilde{\delta}_t [u(c_t) - u(c'_t)] &< 0 \end{aligned} \tag{5}$$

From linear independence of the discount vectors and richness of the domain of u , it follows that there are solutions C and C' to this system. By definition, all individual agents prefer consumption

stream C to C' , while the collective prefers C' to C , violating unanimity. There must therefore be weights w_i such that:

$$V[\delta_1, u; \dots; \delta_n, u](C) = \sum_i w_i U_i(C).$$

Furthermore, the existence of a solution to the first K inequalities in system (5) suggest that $w_i \geq 0$ with at least one strict inequality, as desired. ■

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 $d(C, C') = \sup_t |c_t - c'_t|$. Finally, time consistency implies his postulates 3, 3', and 4. Thus, there exists $0 < \delta < 1$ and a continuous u , such that $V(C) = \sum_t \delta^t u(c_t)$ for all C . ■

Proof of Theorem 2:

From Theorem 1, $V[\delta_1, u_1; \dots; \delta_n, u_n](C) = \sum_t \delta^t u(c_t)$ for all C . By unanimity, it follows that u is increasing and, by assumption, it is twice continuously differentiable.

Without loss of generality, let us normalize u so that $u(0) = 0$ and $u(1) = 1$, and do the same with each u_i , so that any agents who have utility functions that are affine transformations of each other now have identical utility functions.

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

Proof of Step 1: Suppose to the contrary.

For any $0 < x < 1$, consider

$$C = (x, x, \dots)$$

and

$$C^\varepsilon = (x + \varepsilon(1 - \gamma), x - 2\frac{\varepsilon}{\delta}, x + \frac{\varepsilon}{\delta^2}, x, x \dots),$$

where $\varepsilon > 0$.

From Taylor's approximation, for any i :

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

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

$$U_i(C^\varepsilon) > U_i(C)$$

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

$$V(C^\varepsilon) = V(C) - \gamma\varepsilon u'(x) + O(\varepsilon^2),$$

and so for ε small enough, unanimity is violated, in 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

$$u = u_i \text{ where } i \text{ is an agent with discount factor of } \delta_i = \delta.$$

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

$$C = (x, x, x, y, x, x \dots)$$

and

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

for $\varepsilon > 0$.

As before, for any j such that $\delta_j \neq \delta_i$,

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

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

By a similar argument,

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

while $U_i(C^\varepsilon)$ can be written as:

$$U_i(C^\varepsilon) = U_i(C) - \delta^3\gamma\varepsilon [\alpha u'_i(y) - u'_i(x)] + O(\varepsilon^2).$$

For sufficiently small ε and $\gamma = \sqrt{\varepsilon}$ it follows that $V(C^\varepsilon) < V(C)$ and $U_i(C^\varepsilon) > U_i(C)$. This violates unanimity. Therefore, our supposition was incorrect and $u = u_i$. ■

Proof of Proposition 3: By the suppositions in the proposition, and ordering agents in nondecreasing order of discount factors, we end up with groups S_1, \dots, S_K such that the groups collect the agents with identical discount factors.

Let

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

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

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:

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

and let x^1, \dots, x^3 be defined so that $Dx^i = a^i$ for all $i = 1, \dots, 3$.

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

Define now the sequence of consumption streams C^1, \dots, C^3 , as follows:

$$\begin{aligned} C^1 &= (1/2, \dots, 1/2, 0, \dots) \\ C^2 &= (u^{-1}(u(c_1^1) - x_1^1), \dots, u^{-1}(u(c_k^1) - x_k^1), \dots, u^{-1}(u(c_K^1) - x_K^1), 0, 0, \dots) \\ C^3 &= (u^{-1}(u(c_1^2) - x_1^2), \dots, u^{-1}(u(c_k^2) - x_k^2), \dots, u^{-1}(u(c_K^2) - x_K^2), 0, 0, \dots). \end{aligned}$$

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 $j = 1, \dots, 3$. Since $\sum_{i=1}^3 x^i = \mathbf{0}$, it follows that $u(c_t^3) - u(c_t^1) = x_t^1$.

In particular,

$$C^1PC^2PC^3PC^1,$$

which is what we wanted to show. ■

Appendix B – Experimental Choices

Choice	Alternative 1	Alternative 2	Alternative 1 Chosen
1	(105,105,80)	(32,144,80)	84%
2	(21,95,80)	(6,130,80)	47%
3	(125,85,80)	(52,124,80)	78%
4	(41,75,80)	(26,110,80)	27%
5	(145,65,80)	(72,104,80)	83%
6	(61,55,80)	(46,90,80)	34%
7	(165,45,80)	(92,84,80)	38%
8	(81,35,80)	(66,70,80)	33%
9	(125,85,80)	(32,144,80)	83%
10	(21,95,80)	(16,120,80)	53%
11	(165,45,80)	(32,144,80)	41%
12	(21,95,80)	(26,100,80)	17%
13	(185,25,80)	(32,144,80)	82%
14	(21,95,80)	(46,90,80)	25%
15	(250,170,80)	(104,248,80)	87%
16	(82,150,80)	(52,220,80)	29%
17	(290,130,80)	(144,208,80)	80%
18	(122,110,80)	(92,180,80)	37%
19	(53,95,80)	(130,40,80)	37%
20	(53,74,80)	(40,78,80)	75%
21	(105,105,105)	(30, 210,105)	62%
22	(105,105,105)	(70,170,105)	58%
23	(105,105,105)	(110,130,105)	33%
24	(105,105,105)	(29,245,105)	46%
25	(105,105,105)	(77,197,105)	21%
26	(105,105,105)	(125,149,105)	21%
27	(125,85,125)	(52,124,52)	83%
28	(125,85,85)	(52,124,124)	29%
29	(41,75,41)	(26,110,26)	67%
30	(41,75,75)	(26,110,110)	21%
31	(145,65,65)	(72,104,104)	21%
32	(61,55,61)	(46,90,46)	42%
33	(105,105,150)	(32,144,150)	62%
34	(21,95,150)	(6,130,150)	25%
35	(125,85,150)	(52,124,150)	92%
36	(41,75,150)	(26,110,150)	25%
37	(145,65,150)	(72,104,150)	75%
38	(61,55,150)	(46,90,150)	21%

Figure 5: The list of choices (in terms of profiles present discounted utility for each of the two streams) and the percentages of choices made by subjects.

Supplementary Appendix

NOT FOR PUBLICATION

8.1 Intransitivities of Binary Voting Rules

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 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 fifty percent of the overall weights cast.⁴¹

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

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

Let a *Binary Voting Rule* be a collective social welfare ordering $R(U)$ that is complete⁴² and depends only on the information in p , so that if

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

then

$$CR(U)C' \text{ if and only if } CR(U')C'.$$

This condition is a variation on 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, 1980).⁴³ 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

⁴¹See Barbera and Jackson (2006) for a discussion of the optimality of voting rules other than simple majority, as well as additional references and background.

⁴²So, for every U, C and C' , either $CR(U)C'$ or $C'R(U)C$.

⁴³Note that "neutrality" is a bit of a misnomer originating from that literature, since it embodies something much different than the usual usage of neutrality that refers to an insensitivity to the labeling of alternatives.

(where, say, choosing certain consumption streams entails specific quorums). Unlike the utilitarian approach taken in the previous section, these rules are intrinsically ordinal, in that decisions depend only on which agents prefer one consumption stream to another and not the magnitudes of utilities involved.

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

Locally Non-dictatorial Orderings *A social welfare ordering R is locally non-dictatorial whenever*

$$CR(U)C' \text{ if } |p(C, C', U)| \geq n - 1.$$

That is, 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. Notice that any supermajoritarian voting rule short of unanimity is locally non-dictatorial.

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

$$\mathcal{C}(x, \gamma) = \{C | c_0 + c_1/\gamma + c_2/\gamma^2 = x, c_t = 0 \ \forall t > 2\}$$

These are three-period consumption streams where 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 regarding 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 smoothing streams over a short horizon.

We say that a strictly concave utility function ($u_i = u$ for all i) is *non-extreme* if there is at least one discount factor $\delta \in (0, 1)$ for which there is a maximizing consumption stream $C^* = \operatorname{argmax}_{C \in \mathcal{C}(x, \gamma)} \sum_t \delta^t u(c_t)$ that is interior.⁴⁴

⁴⁴Note 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 $u'(1/2)(\delta/\gamma)^2 < u'(0)$.

THEOREM 3 *Consider consumption streams in $\mathcal{C}(x, \gamma)$ for any $x > 0$ and $\gamma > 0$, and suppose that agents all have the same non-extreme and strictly concave utility function u . If a binary voting rule R is locally non-dictatorial, then there exists a profile U and consumption streams for which R is intransitive. In fact, for any neighborhood $B \subset \mathcal{C}(x, \gamma)$ of C^* (as defined above), there exists $C' \in B$ and $C'' \in B$ such that $C^*R(U)C'R(U)C''P(U)C^*$.*

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. Notice that intransitivities are within a rather restricted set of alternatives pertaining to consumption smoothing over only three periods. Clearly, then, under the theorem's conditions, intransitivities also arise for consumption smoothing problems with longer horizons, or when consumption streams are less restricted.

We stress that 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), will be a Condorcet winning alternative and cycles disappear. Indeed, in the example above, the three consumption streams did not correspond to a consumption smoothing problem; We return to this point below.

8.2 Well-Ordered Alternatives

One situation in which there is enough dependence to rule out the potential for intransitivities entails consumption streams that can be ordered so that differences between any two streams of instantaneous utilities are monotonic. Formally,

Well-ordered Alternatives *Consumption streams C and \widehat{C} are well-ordered relative to a society with discount factors $\delta_1, \dots, \delta_n$ 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 C and \widehat{C} are well-ordered and that, say, $u(c_t) - u(\widehat{c}_t)$ is increasing. When consider the differences in net present values between C and \widehat{C} , we consider sums of the form

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

As we increase δ , more weight is put on elements further in the sequence $\{u(c_t) - u(\widehat{c}_t)\}_t$ and so whenever an agent with a discount factor of δ evaluates this expression as positive, so that C is preferred to \widehat{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.

Before describing the structure well-ordering of alternatives imposes, we introduce several standard restrictions on the voting rules we consider.

A binary voting rule R is *neutral* if for any (C, C', U) and $(\widehat{C}, \widehat{C}', \widehat{U})$, if

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

then

$$CR(U)C' \text{ if and only if } \widehat{C}R(\widehat{U})\widehat{C}'.$$

A binary voting rule R is *monotone* if for any (C, C', U) and \widehat{U} , if

$$p(C, C', U) \subset p(C, C', \widehat{U}) \text{ and } p(C', C, \widehat{U}) \subset p(C', C, U)$$

then

$$CP(U)C' \text{ implies } CP(\widehat{U})C'.$$

Any weighted majority rule is a neutral and monotone binary voting rule.

When alternatives are well-ordered, neutrality and monotonicity are enough to rule out intransitivities in the strict societal voting rule. The intuition is the following. Suppose that $CP(U)C'P(U)C''$ and, for simplicity, assume that $u(c_t) - u(c'_t)$ and $u(c'_t) - u(c''_t)$ are both increasing, so that their sum, $u(c_t) - u(c''_t)$, is also increasing. As discussed above, this type of well-ordering guarantees that if an agent with discount factor δ prefers C to C' , so would any individual who is more patient. Denote the threshold discount factor by δ' . Similarly, the threshold discount factor for preferring C' to C'' can be denoted by δ'' . Now, since individuals themselves are transitive, any agent preferring C to C' and C' to C'' also prefers C to C'' . It follows that the threshold discount factor for preferring C to C'' must be (weakly) lower than $\max\{\delta', \delta''\}$. So that in the society, the set of agents who prefer C to C'' is a super-set of either the set who prefer C to C' or the set who prefer C' to C'' . Neutrality and monotonicity then guarantee that $CP(U)C''$.

The following proposition formally summarizes the above discussion.⁴⁵

⁴⁵We use the notion of *quasi-transitivity*: $CP(U)C'P(U)C''$ implies $CP(U)C''$. This notion is a restricted form of full transitivity, which allows for P to also be replaced by R .

PROPOSITION 4

1. If C and \widehat{C} are well-ordered relative to a society with discount factors $\delta_1 \leq \dots \leq \delta_n$ and a utility function $u_i = u$ for all i , then if $i \leq j$ both (weakly) prefer C to \widehat{C} , then so does any k such that $i \leq k \leq j$.
2. If C is a set of consumption streams such that any pair of consumption streams in C is well-ordered relative to a society with discount factors $\delta_1 \leq \dots \leq \delta_n$ and a utility function $u_i = u$ for all i , then any neutral and monotone voting rule is quasi-transitive (so its strict ranking is transitive and therefore acyclic) over C .

Quasi-transitivity rules out cycles in the strict relation, but is a weaker condition than full transitivity of R as it does not deal with cases of social indifference. The restriction to quasi-transitivity is necessary unless the set of alternatives is confined enough so that the individuals themselves are not indifferent between alternatives. The following example illustrates intransitivities involving societal indifferences.

Example 4 (Intransitivities with Indifferences) Suppose society is composed of two agents and the voting rule is defined so that $CP(U)C'$ if both agents prefer C to C' and otherwise C and C' are declared to be indifferent (and $CR(U)C'R(U)C$). Suppose the agents have a linear instantaneous utility function, $u(c) \equiv c$, and discount factors $\delta_1 = 3/4$ and $\delta_2 = 1/2$. Consider now three consumption streams:

$$\begin{aligned} C &= (.9, 1, 0, 0, \dots), \\ C' &= (1, .8, 0, 0, \dots), \quad \text{and} \\ C'' &= (.85, 1, 0, 0, \dots). \end{aligned}$$

Agent 1, with discount factor $\delta_1 = 3/4$, prefers C to C' and is indifferent between C' and C'' . Agent 2, with discount factor $\delta_2 = 1/2$, is indifferent between C and C' , and prefers C' to C'' . Thus, socially $C''R(U)C'R(U)CP(U)C''$, contradicting transitivity.

8.3 More General Consumption Patterns

Up to this point, we considered situations in which each agents' 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. In addition, specifically accounting for multiple dimensions of consumption may be crucial for particular application of

intertemporal decision-making. In this section we extend our underlying framework to allow for more general consumption streams.

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

This more general formulation of a vector of consumptions allows for the evaluation of all sorts of combinations of agents' private and public consumptions. A special case is where $\ell = 1$, which reduces to the previous formulation.

The definition of a collective utility function, time consistency, and unanimity are exactly as before as none 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 .⁴⁶

Note that this still allows for agents to be bargaining over streams of consumption, as their utility function need only increase in their own consumption. In particular, these assumptions allow for an agent to be made worse off by an increase in another agent's consumption, or even a public element of consumption, if that comes at the expense of a decrease in the agent's private consumption. Therefore, this more general setting embodies general combinations of private and public consumption.

The following is a variation of Theorem 2.

THEOREM 4 *A collective utility function is unanimous, twice continuously differentiable, and time consistent at a profile of preferences such that agents all have different discount factors if and only if it is dictatorial.*

Theorem 4 generalizes Theorem 2, which follows from the proof of Theorem 4 in the special case where $\ell = 1$. The proof of Theorem 4 is slightly more complicated, but is analogous to the proof of Theorem 2.

Similarly, there are analogs of our intransitivity of voting results in this setting. For example, consider the following extension of Proposition 3.

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

⁴⁶The utility function u_i is nondecreasing overall if whenever $c \geq c'$ (so that $c_j \geq c'_j$ for every j), $u_i(c) \geq u_i(c')$.

PROPOSITION 5 *If $n \geq 3$ and $(u_1; \dots; u_n)$ are weakly similar through some u that is continuous and strictly increasing, then if P is defined by majority rule, and the largest group of agents having identical discount factors is smaller than a majority, then $P(\delta_1, u_1; \dots; \delta_n, u_n)$ is intransitive.*

We omit the proof of Proposition 5 since this follows more directly from our earlier results. In particular, it follows directly from the intransitivity on the restricted domain of consumptions, where each dimension has the same consumption level (but consumption may vary across dates). Similarly, a variant of Theorem 3 also holds in this setting.

8.4 Proofs of the Results in the Supplementary Appendix

Proof of Theorem 3: Suppose that C^* is interior and optimal for δ^* , so that

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

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

For small a and b , $U_i(C[a, b])$ is approximately (up to elements of $O(a^2)$, $O(b^2)$),

$$U_i(C^*) + \frac{u'(c_1^*)}{\gamma} \left[a \left(1 - \frac{\delta_i}{\delta^*} \right) - b \frac{\delta_i}{\delta^*} \left(1 - \frac{\delta_i}{\delta^*} \right) \right]. \quad (6)$$

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

Restricting attention to C^* , $C[a', b']$, $C[a'', b'']$, 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 nondictatorial 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 $C[a', b']$ and $C[a'', b'']$ are interior and restricted further below).

First, let us find δ 's that entail 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^* , C^* is most preferred and $C[a'', b'']$ and $C[a', b']$ lead to (nearly) the same, but lower, utility. For $\delta_i < \delta^*$ the relative utility of $C[a', b']$ rises relative to $C[a'', b'']$ since $\left(1 - \frac{\delta_i}{\delta^*}\right)$ is positive (see 6), while the reverse happens as $\delta_i > \delta^*$ as $\left(1 - \frac{\delta_i}{\delta^*}\right)$ 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 C^* relative to the other two consumption streams but imply a different ranking of $C[a'', b'']$ and $C[a', b']$.

Next, let us find δ 's that induce $C[a', b']$ as the most preferred of the three streams, and rank the other two consumption streams ordered in an arbitrary order. Consider $\delta' = \delta^*/x$. It follows from (6) that under $\delta' < \delta^*$, $C[a', b']$ 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' induces C^* to be preferred to $C[a'', b'']$, and δ_4 slightly larger than δ' inducing the reverse, both while still maintaining $C[a', b']$ as the most preferred.

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

Proof of Proposition 4:

1. Consider the case where $u(c_t) - u(\hat{c}_t)$ is nonincreasing, as the other case is similar, and suppose that there is some $t' > 0$ such that $u(c_t) - u(\hat{c}_t) > 0$ for $t < t'$ and $u(c_t) - u(\hat{c}_t) \leq 0$ for $t \geq t'$, as otherwise preferences are trivial.

Consider any j such that

$$\sum_t \delta_j^t (u(c_t) - u(\hat{c}_t)) \geq 0. \quad (7)$$

It follows that if $\delta_k \leq \delta_j$, then

$$\delta_k^t \frac{\delta_j^{t'}}{\delta_k^{t'}} \geq \delta_j^t$$

for $t < t'$ and

$$\delta_k^t \frac{\delta_j^{t'}}{\delta_k^{t'}} \leq \delta_j^t$$

for $t \geq t'$.

Therefore, from (7), the above inequalities, and the definition of t' , it follows that

$$\sum_t \delta_k^t \frac{\delta_j^{t'}}{\delta_k^{t'}} (u(c_t) - u(\hat{c}_t)) \geq 0.$$

This implies that

$$\sum_t \delta_k^t (u(c_t) - u(\hat{c}_t)) \geq 0,$$

and so, if an agent j prefers C to \widehat{C} , then so do all agents with lower discount factors. Analogously, if an agent j prefers \widehat{C} to C , then so do all agents with higher discount factors.

2. From part 1 it follows that the set of agents preferring one stream to another (when well-ordered) is a connected interval which contains either agent 1 or n if it is nonempty. Consider streams $C, C', C'' \in \mathcal{C}$ such that $CP(U)C'$, and $CP(U)C''$. We now show that $CP(U)C''$.

Suppose first that both $u(c_t) - u(c'_t)$ and $u(c'_t) - u(c''_t)$ are increasing. Then, there exist δ' and δ'' such that any agent with discount $\delta > \delta'$ prefers C to C' and agent with discount $\delta > \delta''$ prefers C' to C'' (with indifference at δ' and δ'' , where $\delta' \in (0, 1]$ and $\delta'' \in (0, 1]$, respectively).

Now, $u(c_t) - u(c''_t) = (u(c_t) - u(c'_t)) + (u(c'_t) - u(c''_t))$ is increasing. Furthermore, for any δ ,

$$\sum_t \delta^t (u(c_t) - u(c''_t)) = \sum_t \delta^t (u(c_t) - u(c'_t)) + \sum_t \delta^t (u(c'_t) - u(c''_t)). \quad (8)$$

Therefore, there exists $\delta^* \leq \max\{\delta', \delta''\}$ such that any agent with discount $\delta > \delta^*$ prefers C to C'' (with indifference at δ^* if $\delta^* \in (0, 1]$). Suppose $\delta^* \leq \delta'$, then

$$p(C, C', U) \subset p(C, C'', U) \text{ and } p(C'', C, U) \subset p(C', C, U).$$

From neutrality and monotonicity it follows that $CP(U)C''$. A similar argument follows if $\delta \leq \delta''$. The case in which $u(c_t) - u(c'_t)$ and $u(c'_t) - u(c''_t)$ are decreasing also follows analogously.

Assume then that $u(c_t) - u(c'_t)$ is increasing, $u(c'_t) - u(c''_t)$ is decreasing, and $u(c_t) - u(c''_t)$ is increasing. As before, there exist δ' and δ'' such that any agent with discount $\delta > \delta'$ prefers C to C' and agent with discount $\delta < \delta''$ prefers C' to C'' (with indifference at δ' and δ'' , whenever $\delta' \in (0, 1]$ and $\delta'' \in [0, 1)$, respectively).

We now show that $\delta' \leq \delta''$. Indeed, suppose the contrary, so that $\delta' > \delta''$. In that case,

$$p(C, C', U) \subset p(C'', C', U) \text{ and } p(C', C'', U) \subset p(C', C, U),$$

and so $CP(U)C'$ and $C'P(U)C''$ would be inconsistent with neutrality and unanimity.

It follows that $\delta' \leq \delta''$. Note that for any $\delta' \leq \delta \leq \delta''$, C is weakly preferred to C' and C' is weakly preferred to C'' . From (8), C is weakly preferred to C'' . Since $u(c_t) - u(c''_t)$ is increasing, there exists $\delta^* \leq \delta' \leq \delta''$ such that any agent with $\delta > \delta^*$ prefers C to C'' (with indifference at $\delta = \delta^*$ if $\delta^* \in (0, 1]$). The claim then follows as before. All other cases are analogous. ■

Proof of Theorem 4:

Again, we apply a theorem by Koopmans (1960, Section 14). First note that by the fact that each u_i is nondecreasing and increasing in some dimension on $[0, 1]^\ell$ and V satisfies unanimity, his postulates 2 and 5 are satisfied. Next, his postulate 1 follows from continuity (in fact, differentiability) of V under the metric $d(C, C') = \sup_t |c_t - c'_t|$. Finally, time consistency implies his postulates 3, 3', and 4. Thus, there exists $0 < \delta < 1$ and a continuous u such that $V[\delta_1, u_1; \dots; \delta_n, u_n](C) = \sum_t \delta^t u(c_t)$ for all C .

By unanimity, it follows that u is nondecreasing and is increasing in some arguments and, by assumption, it is twice differentiable.

Without loss of generality, let us normalize u so that $u(0) = 0$ and $u(1) = 1$ (where the inputs are now vectors).

Assume $\delta_i \neq \delta_j$ for all $i \neq j$.

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

Proof of Step 1: This proof proceeds as in the proof of Theorem 2, simply restricting attention to the domain of consumptions such that $c_{t,k} = c_{t,k'}$ for all k, k' , which then reduces the domain to one dimension at each time. The previous proof then goes through (writing utility functions as functions of one dimension and noting that they are all strictly increasing when restricted to this domain of consumptions) and then implies the desired result. Therefore, there exists i such that $\delta = \delta_i$.

Step 2: Given that all agents have different discount factors, it follows that u is an affine transformation of 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 is not an affine transformation of u_i

It follows that there exist $x \in [0, 1]^\ell$, $y \in [0, 1]^\ell$, $z_1 \in \mathbb{R}^\ell$, and $z_2 \in \mathbb{R}^\ell$, such that for all sufficiently small a : $x + az_1 \in [0, 1]^\ell$, $y + az_2 \in [0, 1]^\ell$, and

$$u(x + az_1) + \delta u(y + az_2) > u(x) + \delta u(y)$$

while

$$u_i(x + az_1) + \delta u_i(y + az_2) < u_i(x) + \delta u_i(y).$$

Let $w = (.5, .5, \dots, .5)$ be an ℓ dimensional vector with all entries being .5.

Set

$$C = (w, w, w, x, y, 0, 0 \dots)$$

and

$$C^{\varepsilon, a} = (w(1 + \varepsilon), w(1 - 2\frac{\varepsilon}{\delta}), w(1 + \frac{\varepsilon}{\delta^2}), x - az_1, y + az_2, 0, 0, \dots)$$

for $\varepsilon > 0$.

As before, for any j such that $\delta_j \neq \delta_i$,

$$U_j(C^{\varepsilon, a}) - U_j(C) = \varepsilon \left[\left(1 - \frac{\delta_j}{\delta}\right)^2 \right] \frac{du_j}{dh}(w) + \delta_j^3 [u_j(x + az_1) + \delta_j u_j(y + az_2) - u_j(x) - \delta_j u_j(y)] + O(\varepsilon^2),$$

where $\frac{du_j}{dh}(w)$ stands for the total derivative of u_j evaluated at w , which is strictly positive given that preferences are non-decreasing overall and increasing in at least one dimension.

Since $\delta_j \neq \delta$, for sufficiently small ε and $a = \varepsilon^{3/2}$, $U_j(C^{\varepsilon, a}) > U_j(C)$.

By a similar argument,

$$V(C^\varepsilon) - V(C) = \delta^3 [u(x + az_1) + \delta u(y + az_2) - u(x) - \delta u(y)] + O(\varepsilon^2),$$

while $U_i(C^{\varepsilon, a})$ can be written as:

$$U_i(C^{\varepsilon, a}) - U_i(C) = \delta^3 [u_i(x + az_1) + \delta u_i(y + az_2) - u_i(x) - \delta u_i(y)] + O(\varepsilon^2).$$

For sufficiently small ε , and $a = \varepsilon^{3/2}$ it follows that $V(C^{\varepsilon, a}) < V(C)$ and $U_i(C^{\varepsilon, a}) > U_i(C)$. This violates unanimity. Therefore, our supposition was incorrect and $u = u_i$. ■