NBER WORKING PAPER SERIES

THE SOCIAL DISCOUNT RATE

Andrew Caplin
John Leahy

Working Paper 7983
http://www.nber.org/papers/w7983

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2000

We thank Andy Abel, Marco Basseto, Susanto Basu, Gary Becker, Ed Glaeser, Jeffrey Campbell, Jean-Jacques Laffont, David Laibson, Robert Lucas Jr., Casey Mulligan, Barry Nalebuff, Mario Rizzo, Nancy Stokey, Robert Solow, and Lars Svensson for helpful comments and discussions. Caplin thanks the C.V. Starr Center at NYU and Leahy thanks the National Science Foundation for financial support. The views expressed in this paper are those of the authors and not necessarily those of the National Bureau of Economic Research.

The Social Discount Rate
Andrew Caplin and John Leahy
NBER Working Paper No. 7983
October 2000
JEL No. D60, D90, E61

ABSTRACT

In welfare theory it is standard to pick the consumption stream that maximizes the welfare of the representative agent. We argue against this position, and show that a benevolent social planner will generally place a greater weight on future consumption than does the representative agent. Our analysis has immediate implications for public policy: agents discount the future too much and the government should promote future oriented policies.

Andrew Caplin
Department of Economics
New York University
269 Mercer Street
New York, NY 10003
and NBER
caplina@fasecon.econ.nyu.edu

John Leahy
Department of Economics
Boston University
270 Bay State Road
Boston, MA 02215
and NBER
jleahy@bu.edu

1 Introduction

When analyzing public policy in dynamic models with a representative agent, it is standard practice to pick the consumption stream that maximizes the representative agent's utility function. Since this practice is so universal, it is almost unnecessary to cite examples. One can simply thumb through any textbook treatment of growth theory, such as Barro and Sali-i-Martin [1995] or Romer [1996], or any analysis of optimal monetary or fiscal policy, such as Lucas and Stokey [1983] or Chari and Kehoe [1999]. The implicit assumption in all of these treatments is that the appropriate social discount rate for future consumption is the same as the private discount rate.

Occasionally there have been those who have voiced discomfort with this prescription, especially those seeking to orient public policy more towards the future. Ramsey [1928], while acknowledging that private agents discount the future, argued that it was "ethically indefensible" for the government to do so. Pigou [1929], Allais [1947], and Solow [1974] have expressed similar views. For the most part, however, these views have either been ignored or dismissed as paternalistic by the majority of economists. Arrow and Kurz [1970] capture the feelings of the profession when they write:

> It is hard to see why the revealed preference of individuals should be disregarded in the realm of time, where it is accepted, broadly speaking, in evaluating current commodity flows. (p. 12)

In this paper we argue against current practice, and in favor of the more future-oriented vision. To immunize ourselves from charges of paternalism, we retain a standard utilitarian approach to evaluating social welfare. We show that social welfare corresponds to the agent's utility only under very special conditions. Instead, a benevolent social planner will generally place a greater weight on future consumption.

The main premise upon which we construct our argument is the observation that agents discount the past, as well as the future. If individuals discount both the past and the future, preferences change over time. The present period weighs more heavily in current decisions than in the decisions of any prior or subsequent period. This means that agents will regret their past consumption choices. We label this form of regret *retrospective time inconsistency*, since it involves disagreement over choices that were made in

the past, as opposed to the standard form of time inconsistency which relates to future plans.

While our approach is not in keeping with current practice, it is not without precedent. In its essentials, our viewpoint was anticipated by Pigou [1929] in his classic work on welfare economics:

> The existence of preference for the present over equally certain future pleasures does not imply that any economic dissatisfaction would be suffered if future pleasures were substituted at full value for present ones. The non-satisfaction this year of a man's preference to consume this year rather than next year is balanced by the satisfaction of his preference next year to consume next year rather than this year. (p. 25)

Time varying tastes pose a challenge for welfare economics. We adopt the vision, familiar from the literature on time inconsistent preferences, that the agent at each point in time is a separate individual. We then follow Pigou [1929] and Schelling [1984] in arguing that all of these temporally distinct selves should be of concern to the social planner. The identification of an optimal policy turns on how the planner should weigh the perspectives of various periods. Any disagreement among the temporal selves will have to be resolved by some form of compromise. We apply the logic of static welfare theory to the dynamic case and formalize the compromise using a Bergsonian social welfare function.

In this framework the currently standard practice of maximizing the utility of a representative agent represents a form of dictatorship in which the perspective of the current self is given full weight in the social welfare function. We label this approach the *dictatorship of the present*. Our view of this dictatorship is that it has all the same ethical pros and cons as other forms of dictatorship. It certainly helps resolve questions in a quick and easy manner, but it may do so at the cost of doing tremendous and predictable harm to someone else in society, in this case the future selves. We see no normative reason to favor the present at the expense of the future. To paraphrase Arrow and Kurz: it is hard to see why dictatorship should be embraced in the realm of time, where it is rejected, broadly speaking, in evaluating current allocations.

Once we reject dictatorship, it follows immediately that the socially optimal plan places greater weight on future utility than does the representative

agent. Any social welfare function that places weight on future perspectives, also places weight on perspectives that rank concerns of the future above those of the present. The appropriate social discount rate is therefore lower than that of the representative agent. By extension, competitive equilibrium is myopic.

In section 2, we discuss the arguments for discounting the past and the future, and make the case for retrospective time inconsistency. Section 3 presents our criterion for social optimality. We characterize the social discount rate in terms of the agent's prospective and retrospective discount rates and the properties of the social welfare function. In general, a social planner should place greater weight on future consumption than does the agent, and for many parameterizations, the social discount rate is actually negative. This bias is decreasing in the weight that the agent places on the past, so that it is in fact greatest if the past does not matter at all.

We conclude section 3 with a discussion of possible criteria for choosing a social welfare function. We first consider the merits of a neutral view that weights equally the perspective of each period of an agent's life. We show that this approach does not generally lead to equal weights on the utility from consumption in each period. Instead these weights depend on the relationship between the prospective and retrospective discount rates and the position of the period in the agents life. We next consider time consistency as a criterion for selection. We show that a social planner's optimal policy is time consistent only if the planner places quasi-hyperbolic weights on the perspectives of future periods. We show that the planner does not have to place very much weight on future perspectives to drive the social discount rate to zero.

In section 4, we discuss a variety of issues raised by our approach. These include the policy implications and the feasibility of applying the theory in practice. We discuss how discounting the past sheds light on rational addiction and procrastination. We close this section by discussing the relationship to the literature. While we develop our approach in the context of time consistency, our methodology is directly applicable to more general models of time changing tastes.

2 Disagreement in the Exponential Model

2.1 Retrospective Time Inconsistency

Consider an agent choosing how to allocate consumption and resources across two periods. The agent is endowed with a certain amount of a raw material and a technology for transforming this material into period one and period two consumption. We assume that the production possibility set is convex. To keep things simple we suppose that the agent chooses consumption in period one to maximize a time additive utility function:

$$U_1(C_1, C_2) = u(C_1) + \beta u(C_2),$$

where C_t represents consumption in period $t = \{1, 2\}$, $u(\cdot)$ is the felicity from consumption within a period, and $\beta \in [0, 1]$ is the discount factor. To avoid later confusion, we will reserve the term "utility" for the welfare of the agent in period t, U_t, and the term "felicity" for the contribution to utility of consumption in period t in the absence of discounting, $u(C_t)$.

The solution to this problem is straightforward and can be found in almost any undergraduate textbook. The optimal policy is to choose the point on the production possibility frontier associated with the highest indifference curve of the utility function U_1. This solution is represented by point $\mathcal{A}$ in the Fisher diagram (Figure 1).

What is not generally appreciated is that this solution represents the optimal choice only from the perspective of the first period, and does not generally maximize second period utility. To see this, let's proceed as is common in dynamic models and treat the past as sunk both for choice and for welfare. In this case second period utility is simply

$$U_2(C_1, C_2) = u(C_2). \tag{1}$$

Given this utility function, the second period indifference curves are horizontal and, from the second period perspective, the optimal choice of consumption is represented by point $\mathcal{B}$ in Figure 1.

The reason that tastes change in this example is that the agent places no weight on period 1 consumption in period 2. One might object that (1) is really only a partial representation of the individual's welfare. It represents that component of the agent's decision that is still "active" at the time of

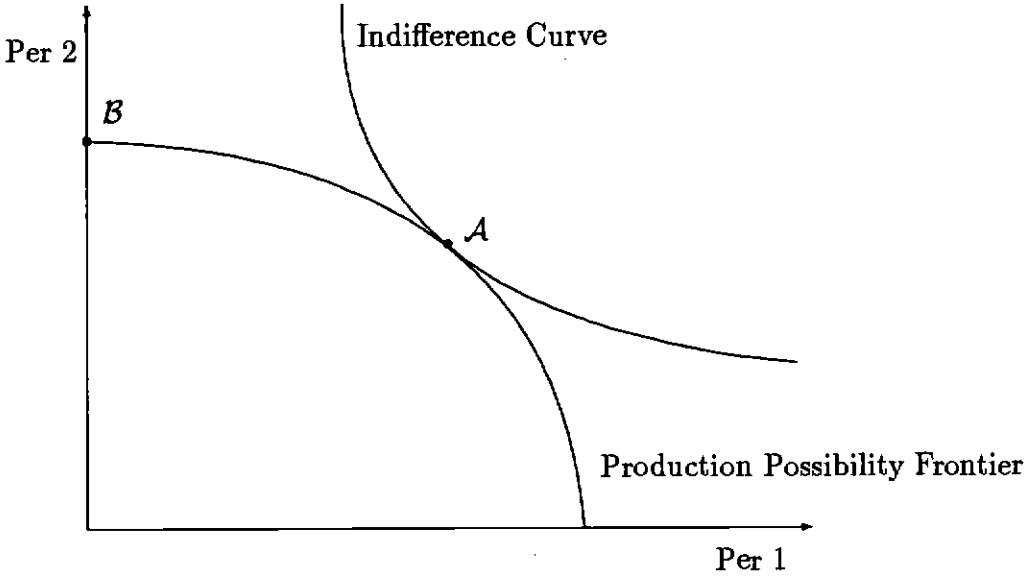

Figure 1: The Fisher Diagram

the period t decision, and ignores the past only because the past is sunk. As Deaton [1992] has emphasized, a complete specification of preferences must always rank the entire consumption stream.

Even if the past enters period 2 welfare, however, the conditions under which tastes remain unchanged are very stringent. Consider the following specification of preferences in which the agent in period 2 cares about both period one and period two consumption

$$U_2(C_1, C_2) = \delta u(C_1) + u(C_2), \tag{2}$$

Here δ represents the "discount" factor applied to past felicity. We place the word discount in quotation marks because we want to allow for the possibility that δ is greater than one. Note (2) reduces to (1) when we restrict the domain of U_2 to current consumption.

From (2) it is easy to see that a necessary and sufficient condition for tastes to be unchanging over time is that $\delta = \beta^{-1}$. In this case, U_2 is simply a monotonic transformation of U_1.[1]

If, on the other hand, $\delta \neq \beta^{-1}$ preferences over consumption sequences change over time. A natural concomitant of $\delta \neq \beta^{-1}$ will be regret concerning past consumption choices. In the example above, the consumer regrets in period 2 the choice of $\mathcal{A}$ in period 1. In period 2, the consumer wishes that the choice had been closer to point $\mathcal{B}$. We see this as a form of time inconsistency that operates in a backward-looking rather than in a forward-looking manner and refer to the phenomenon as retrospective time inconsistency in order to distinguish it from the more familiar prospective time inconsistency.

The argument easily extends to a multi-period setting. Consider an infinite horizon, $t \in \{0, 1, 2, \ldots\}$. We can represent the agent's utility over consumption sequences by a function that discounts both the past and the future:

$$U_t(C_0, C_1, \ldots) = \sum_{m=1}^{t} \delta^m u(c_{t-m}) + u(c_t) + \sum_{n=1}^{\infty} \beta^n u(c_{t+n}), \tag{3}$$

Figure 2 illustrates both the discount factor of an individual in period zero and an individual in period ten under the assumption that $\delta = \beta = .9$.

[1]Deaton [1992] and Mas-Collel, Whinston and Green [1995] implicitly assume $\delta = \beta^{-1}$ when they assume that preferences over consumption sequences are time invariant.

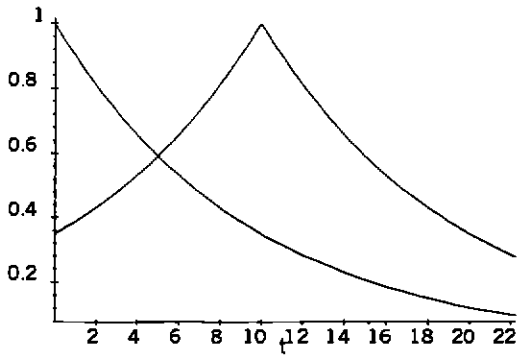

Figure 2:

In period zero the agent places exponentially declining weights on future felicity. In particular, period zero felicity receives greater weight than period ten felicity. In period ten, the agent place exponentially declining weights on both past and future felicity, and the relative weights on period zero and period ten are reversed.

We can see from (3) that retrospective time inconsistency is perfectly consistent with exponential discounting and time consistency in the standard sense.[2] Given any dates $s' > s > t$, the marginal rate of substitution between consumption at s and s' is independent of the period of evaluation t:

$$\frac{\beta^{s'-s}u'(C_{s'})}{u'(C_s)}.$$

This property implies that optimal plans made to maximize (3) are time consistent: a plan chosen to maximize utility in period t will remain optimal in all subsequent periods.

We also see from (3) that time consistency does not imply a preference ordering over sequences of consumption that is time invariant. Time consistency only requires that preferences over future consumption are unchanging. It does not rule out disagreements over the past. Retrospective time consistency requires exponential discounting and $\delta = \beta^{-1}$.

[2]Strotz [1956] showed that exponential discounting is necessary and sufficient for time consistency.

The question whether choices today maximize utility from the perspective of all future time periods therefore boils down to how much agents care about the past. Is it reasonable that $\delta = \beta^{-1}$ or do preferences change over time? In the next section, we review the case for discounting the future and argue that the case for discounting the past is at least as persuasive.

2.2 The Argument for Discounting: Future and Past

Ever since Böhm-Bawerk [1959] first proposed discounting as one of the prime reasons for positive interest rates, there has been near unanimity among economists that individuals discount future felicity relative to present felicity.[3] The main argument in favor of discounting is empirical. It appears to be the revealed preference of individuals that current pleasure is preferred to future pleasure, and future pain to current pain. Introspection supports this view, as do market interest rates in excess of the growth rate of population and productivity, and so does the available experimental evidence from economics and psychology.[4]

Given that individuals discount the future, invariance requires that agents weigh the past more heavily than they do the present. In other words, it requires that a thirty year old cares more about an apple that was consumed at age three than about an apple to be consumed today. Moreover, the importance of that apple at age three relative to an apple today must grow at an exponential rate as the consumer ages.

This implication of invariance is difficult to accept. Rather than viewing past consumption as progressively more important as the consumer ages, it is more plausible that agents discount past consumption. When asked most people would prefer having completed an onerous task last week over having to perform one this evening; they would prefer a pleasurable meal this evening to one ten years ago; and they would prefer to be leaving on vacation to returning to work when their vacation is over.[5,6] Whenever one

[3]For a rare statement of dissent see Becker and Stigler [1977].

[4]Although supportive of discounting, the experimental evidence does not necessarily support exponential discounting. See Ainslie [1992].

[5]This last statement assumes that the vacation is a pleasurable experience. For some families this may not be true. In this case the inequality is reversed, but the point is the same.

[6]There will, of course, be exceptions. There may be that particularly pleasurable

thinks, "I'm glad it is finally three o'clock and that lecture is over," it reflects discounting of the past. The disagreeable experience is less painful when it is in the past than when it is being experienced in the present.[7]

Beyond revealed preference, the strongest theoretical argument in favor of discounting, involves what Böhm-Bawerk referred to as the "brevity and uncertainty of human life." According to this view, people discount future felicity because they may not be around to enjoy it. An individual weighs pleasures at date t by the conditional probability of living until date t.

There is a parallel between mortality and forgetfulness. The past dies when we forget.[8] Do we value today meals eaten ten years ago, if we cannot even remember what we ate? How is our current utility reduced by past pains that we cannot even remember? Imperfect memory justifies discounting the past in much the same way that mortality justifies discounting the future, and just as mortality suggests that we discount the far off future more than the near future, imperfect memory suggests that we discount the far off past more than the recent past.

Finally, heterogeneity in discount rates across individuals undermines the case for retrospective consistency. If, as is commonly believed, the rich and educated are more forward looking, do they place relatively *less* weight on the past as invariance would require? To the extent that individual discount rates are endogenous, as argued by Becker and Mulligan [1997], does the weight on the past adjust to maintain $\delta = \beta^{-1}$? Or is the weight on the past also endogenous? After all, we spend considerable resources keeping the past alive, through stories, photographs, souvenirs, and diaries. If the weight on the past is endogenous, wouldn't it be a remarkable quirk of fate of this

---

occurrence which yields memories for years to come, memories so precious that one is glad that the experience was in the past so that the memories could be enjoyed over an extended period of time. Similar exceptions, however, apply to discounting the future. One may, for example, want to delay a particularly pleasurable event in order to savor the feelings of anticipation that it engenders (Loewenstein [1987]). The existence of a few exceptions, however, does not alter the general rule.

[7]Some suggestive evidence comes from DeGenova [1992] who found that when people looked back on their lives they most regretted not spending more time on education. Education, being an activity with current costs and future rewards, is a prime candidate for retrospective time inconsistency.

[8]The experimental evidence in psychology indicates that exponential decay of recall probabilities fits the data rather well, although a power function may perform better (See Crovitz and Schiffman [1974]).

weight were exactly equal to the inverse of the discount factor?

We conclude that in almost any reasonable formulation, future selves weight current consumption less heavily than does the current self. The implication is that tastes change over time. We now consider the welfare implications of this observation.

3 Dynamic Welfare Theory

As with most economic theory, the roots of the dynamic case are to be found in a close examination of the static case. In static welfare theory we begin with a set of individuals with preferences defined on an appropriate commodity space. We add a social planner who has available a set of policy tools, each of which is identified with a particular consumption bundle for each individual in the economy. The social planner is then assumed to pick a policy that satisfies some set of ethical desiderata. The most universally accepted of these is Pareto optimality. Beyond Pareto optimality one might look to load on other properties directly, as in cooperative bargaining theory, or implicitly, as when maximizing some social welfare function.

The case of a dynamic representative agent is no different. All that is needed is to identify the set of policy choices available to the social planner, the relevant set of individuals, their payoffs, and the nature of the social welfare criterion.

Identifying the set of policy choices is relatively straightforward. This set is usually given as part of the data of the problem or as the set of equilibrium outcomes of some model.

The set of individuals is not a trivial matter. We have seen that even in representative agent models with exponential discounting, tastes may change over time. Which preferences should the planner maximize? Instead of constructing a theory around a fixed answer to this question, we adopt a framework that is flexible enough that it allows for the possibility that the perspective of each point in time may affect welfare. We take the set of individuals to be the set of temporally distinct selves indexed by time. Given this collection of individuals, the payoff function for the date t self is given by the U_t in equation (2).

One way to think about the set of individuals is to consider a Rawlsian thought experiment: which self shows up to represent the interests of the

individual in the initial position? When thinking about retirement, for example, does the individual take the perspective of a retired person receiving a pension or a young person saving? Our view is that each of these perspectives has some merit and that some form of compromise is natural.[9]

This leaves only the determination of the welfare criterion. We will begin by analyzing the set of Pareto optimal allocations, and then turn to the issues involved in choosing among Pareto optima.

3.1 Properties of Pareto Optima

For the purposes of discussion we shall assume a finite horizon $T > 0$ and we shall assume that the set of policy choices are limited by the budget constraint

$$\sum_{t=0}^{T} \frac{C_t}{(1+r)^t} \leq W$$

where r represents the rate of return on saving and W represents the initial resource endowment.

Let $\mathcal{P}$ denote the set of Pareto optimal allocations $C^T = \{C_0, C_1, \ldots C_T\} \in \mathbb{R}^T$. We state the following standard result without proof.

Lemma Consider the problem of a social planner maximizing the weighted average of the individual's utilities:

$$S(U_0, U_1, \ldots, U_T) = \sum_{t=0}^{T} \alpha_t U_t. \tag{4}$$

Let $\mathcal{P}_{SP}$ denote the set of consumption streams C^T that achieve the maximum of S for some weights $\alpha_t \geq 0$, for all t. Then $\mathcal{P}_{SP} = \mathcal{P}$.

[9] Rawls [1999] reaches essentially the same conclusion:

> The question is settled by reference to the original position;...There is no reason for the parties to give any weight to mere position in time....If they make a distinction between earlier and more remote periods because, say, the future states of affairs seem less important now, the present state of affairs will seem less important in the future.

Rawls concludes that pure time preference is unjust. In section 3.2, however, we show that an equal weighting of each period in life does not imply an absence of time preference, only more forward looking behavior.

Hence to study the set of Pareto optima, we may study the properties of the class of linear Bergsonian social welfare functions of the form (4). Note that standard approach of maximizing the present value of the agent's utility is a member of this class of welfare functions. In this case, the period zero self receives all of the weight in the social welfare function: $\alpha_0 > 0$ and $\alpha_t = 0$ for all $t > 0$. This is a form of dictatorship, which we term the dictatorship of the present.

All social welfare functions of the form (4) are non-paternalistic in the sense that welfare economics uses the term (see Mas-Collel, Whinston, and Green [1995, p. 895]). They are functions only of the agent's own utility and the planner does not have any direct preference over the allocation. Note, however, that the social choice will not in general be the same as the private choice. Whether or not this welfare function is paternalistic in the broader sense of the term depends on how the society decides to implement the social optimum: through dictate or by altering the incentives that mould the agent's own choices.

After substituting for the U_t, we have that maximizing S is equivalent to maximizing

$$\sum_{t=0}^{T} q_t u(C_t)$$

where

$$q_t = \sum_{s=0}^{t-1} \alpha_s \beta^{t-s} + \alpha_t + \sum_{s=t+1}^{T} \alpha_s \delta^{s-t}.$$

The next result follows immediately:

Proposition 1: If $\delta \leq \beta^{-1}$, then for all $t \in \{0, 1, \ldots T-1\}$, $q_{t+1}/q_t \geq \beta$. The inequality is strict if and only if $\delta < \beta^{-1}$ and $\alpha_s > 0$ for some $s > t$.

q_{t+1}/q_t is the social discount factor relating period t and period $t+1$ felicity. The proposition states the sense in which the dictatorship of the present is the most present biased of the Pareto optima. If $\delta < \beta^{-1}$, then any social welfare function that places weight on any future perspective discounts future felicity by less than the dictatorship of the present.

There is a sense in which increasing the weight on the perspective of period t in the social planner's problem shifts consumption towards period

t. Let $\sigma_t = u'(C_{t+1})/u'(C_t)$. σ_t denotes the ratio of the marginal utility of consumption in period $t+1$ to the marginal utility in period t.

Proposition 2: Given $\delta < \beta^{-1}$, $d\sigma_s/d\alpha_t \leq 0$ for $s < t$ and $d\sigma_s/d\alpha_t \geq 0$ for $s \geq t$.

Proof: The statement follows from manipulating the social planer's Euler equation.□

The proposition states that an increase in the weight given the period t perspective causes marginal felicity to fall more rapidly prior to t and rises less rapidly after t. If δ is small enough and the weight on future perspectives is high enough it is possible for σ_t to become negative, so that consumption is rising for some period of time.

To see how these propositions work, consider the two period example with $\delta = 0$, and consider the following social welfare function:

$$S = aU_1 + (1-a)U_2.$$

In this case, varying the choice of a from zero to one traces the production possibility frontier from $\mathcal{B}$ to $\mathcal{A}$. All of these points are Pareto optimal. Any movement from any one of them either violates the resource constraint or reduces welfare from the perspective of one of the periods. The standard practice of maximizing the utility of representative agent is equivalent to setting $a = 1$. Note that standard practice does not yield a random selection from the set of Pareto optima. Rather point $\mathcal{A}$ is the most present oriented element of $\mathcal{P}$. Any $a < 1$ produces an outcome that is more favorable towards the future than $a = 1$.

If we raise δ the period two optimum $\mathcal{B}$ moves along the production possibility frontier towards $\mathcal{A}$. The range of Pareto optima shrinks, until at $\delta = \beta^{-1}$ preferences coincide. Intuitively, the less backward looking are second period preferences, the more forward looking is the set of Pareto optima. This result also holds with many periods. The less backward looking agents are the more forward looking the set of Pareto optima.

Proposition 3: Let $\mathcal{P}_\delta$ represent the set of Pareto optimal allocations given a retrospective discount rate of δ. If $\delta_1 > \delta_2$, then $\mathcal{P}_{\delta_1} \subseteq \mathcal{P}_{\delta_2}$.

Proof: See Appendix.□

It is interesting to consider how the social discount rate behaves at long horizons. Suppose $\alpha_t = \alpha^t$ so that the welfare function discounts future utility at rate α, and let the time horizon T go to infinity. It is easy to calculate the following result.

Proposition 4: Let $\alpha_t = \alpha^t$ and $T = \infty$, then $\lim_{t\to\infty} q_{t+1}/q_t = \max\{\alpha, \beta\}$.

As the horizon becomes long, the relative weights that the welfare function places on felicity in adjacent periods approaches the maximum of the prospective discount factors of the planner and the agent. To understand this limit, note that there are two channels by which the felicity in the far future impacts welfare. The first is through individuals in the near future who care about felicity in the far future. These individuals discount future felicity by β. The second is through individuals in the far future who care about felicity in the near future. The welfare function discounts the utility of these individuals by α. Which channel is more important in the limit depends on which discount factor is larger.

3.2 Choosing a Social Welfare function

The set of possible social welfare functions is potentially large. While we find dictatorship inherently troubling, eliminating dictatorial welfare functions eliminates only a few scattered elements in this set. How is one to choose a social welfare function?

One potential criterion for narrowing the set is symmetry. Placing equal weights on each perspective, $\alpha_t = 1$, generates the following weights on felicity:

$$q_t = \sum_{s=0}^{t-1} \beta^{t-s} + 1 + \sum_{s=t+1}^{T} \delta^{s-t}$$

Figure 3 illustrates these weights for several values of δ. For $\delta = \beta^{-1}$, all perspectives agree to discount later periods and q_t declines exponentially. For $\delta = \beta$, q_t is hump-shaped, rising for $t < T/2$ and falling thereafter. Felicity in the middle periods receives greater weight because it is closer in time to all

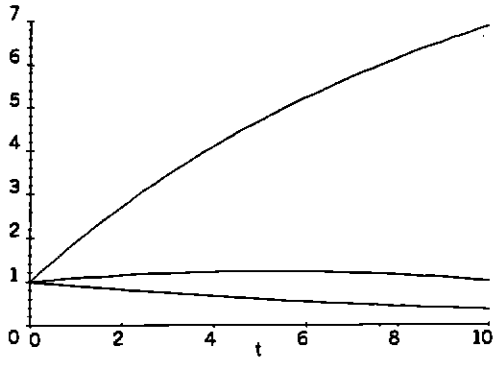

Figure 3: Social weights on felicity for various values of delta.[10]

periods of evaluation and is therefore not discounted by as much as felicity in early or late periods. For $\delta = 0$, q_t is monotonically rising. q_t rises because felicity in the future matters more than felicity in the past.

We draw two conclusions from figure 3. First, we see that equal weights on utility in all periods does not imply equal weights on felicity in all periods. Second, we see that the less the agent cares about the past, the more the social planner cares about the future.

While symmetry may have some appeal when the horizon is finite, it has undesirable consequences when the horizon is infinite. As noted by Koopmans [1960], a social welfare function which gives equal weight to utility at all dates is insensitive to changes in utility at any given point in time. One potential solution is to discount future perspectives, choosing $\alpha_t = \alpha^t$. This, however, leaves the choice of α to be determined.

Another criterion that significantly reduces the set of Pareto optima is time consistency. If we think of the social optimum as being implemented over time by a social planner, then the planner may be tempted to alter the policy as time passes. Why, for example, should the social planner care about the feelings of past selves, when these feelings are past and cannot be altered? At date t the planner may decide to reset α_s to zero for all $s < t$ and $\alpha_s = \alpha_{s-t}$ for all $s \geq t$. The next proposition presents a class of social

[10]For all curves beta=.9 and the horizon is 10.

welfare functions that are time consistent and immune to such behavior.

Proposition 5: Suppose $\delta < \beta^{-1}$. Let $\alpha_0 = 1$ and, for $t > 0$, let $\alpha_t = (\gamma - \beta)\gamma^{t-1}$ where $\gamma \in [\beta, \infty)$. Then the social planner discounts future felicity at the constant rate γ.

The proposition provides the utility weights that generate exponential discounting of felicity. It states that if the planner has a certain quasi-hyperbolic discount rate, then the planner is time consistent. In an infinite horizon, this approach is robust in to Koopmans' critique if $\gamma \in [\beta, 1)$. It still leaves γ undetermined.

One recurrent theme in welfare economics is a dislike for discounting. A preference for equal weights on felicity is implicit in Ramsey's [1928] contention that discounting the future was "ethically indefensible," and is also implied in the quotation from Pigou [1929] cited in the introduction. To generate equal weights on felicity the planner needs to choose $\alpha_0 = 1$ and $\alpha_t = 1 - \beta$ for all $t > 0$. Given conventional calibrations of β, the weight on the perspective of each future year is then $1/20$ of the weight on the current perspective. What is surprising here is that it does not take much weight on the future perspectives to generate a very forward looking social welfare function.

A final possibility is to deduce the social welfare function empirically through the revealed preference of policy makers. This could be done by analyzing situations in which agents make choices that they believe would be in the best interests of other agents and comparing these choices to the choices that agents make for themselves.

4 Discussion

We begin with an number of observations concerning the practical implications of our approach and conclude with a discussion of the related literature.

4.1 Policy Implications

Our analysis has immediate implications for public policy: agents discount the future too much and therefore governments should promote future oriented policies. Theoretically, our analysis provides a foundation for what

Malinvaud [1985] calls an Allais equilibrium (See also Allais [1947]). An Allais equilibrium is a situation in which individual preferences are retained for the choice between consumption goods at the same date, but not necessarily between consumption goods at different dates.

Relative to the market evaluation, our analysis tips the scales in favor of policies with short run costs and long run benefits. In the area of fiscal policy, it favors subsidies to capital accumulation. In the area of monetary policy, it favors low inflation. In the area of natural resource extraction, it favors conservation. In general, it favors investment and saving at the expense of consumption.

4.2 Procrastination

O'Donoghue and Rabin [1999a] model procrastination using time inconsistent preferences that exhibit a bias towards current felicity. They argue that procrastination occurs when a time inconsistent individual acts later than a time consistent individual would act, and that preponation occurs when a time inconsistent individual acts relatively sooner. They argue that procrastination is more likely when an agent incurs the costs of an action prior to the reward and that preponation is more likely when rewards precede costs.

Discounting past felicity provides an alternative characterization of procrastination and preponation. In this characterization, agents are time consistent, but they regret their past actions. Procrastination is associated with a wish that an action had been taken at an earlier date. Preponation is associated with a wish that an action had not been taken at an earlier date.

As in O'Donoghue and Rabin, procrastination will tend to occur when costs precede rewards. At the time that an act is postponed, these costs appear salient and the future rewards are discounted. In the future, the agent comes to regret this delay, because, if the action had been taken, the rewards which would now be received would be salient, and the costs which would have been incurred in the past would be discounted. In a similar manner, preponation is associated with rewards that precede costs.

4.3 Addiction

Becker and Stigler [1977] and Becker and Murphy [1988] model addiction as rational. Agents are time consistent and have stable felicity functions. They

weigh the present value of costs and benefits of consuming the addictive substance and consciously choose to become addicts.

Because addiction is the result of a rational choice in this framework, there is the impression that addicts are "happy" or at least better off than they would have been if the option to become addicted had never existed. This vision of the "happy addict" has caused some economists to question the usefulness of the framework (O'Donoghue and Rabin [1999b]).

This welfare calculation, however, assumes that agents are retrospectively time consistent. Harmful addiction is a case in which rewards precede costs. The addict decides that the benefits to current consumption of the addictive substance outweigh the future costs of increased addiction. As in the case of procrastination, agents who discount the past may come to regret their past behavior as these future costs become salient. Rational choice may maximize utility from the current perspective, but this does not imply that future utility is also maximized. Rational addicts need not be happy addicts. There may therefore be a rationale for prohibition as a means of protecting the agent's future self from his current decisions.

4.4 Political Feasibility

One might argue that the case for the dictatorship of the present is positive rather than normative. If social decisions result from a democratic process, then the current self would indeed appear to be a dictator, since only the current self can vote. The dictatorship of the present, being the only criterion that is practically implementable, would then arise as a type of constrained optimum.

It is not obvious, however, that political constraints favor the dictatorship of the present. Far sighted policy may arise in a number of ways. Sen [1961] and Marglin [1963] have shown how altruism may lead agents to vote for policies that are more forward looking than the choices they would make for themselves. Their argument is framed in an intergenerational context, but is easily cast in an intrapersonal setting.

Example: Consider two altruistic agents. Each cares about their own utility and the welfare of the other agent. They maximize

$$U_0(C_0, C_1) + \phi S(\tilde{C}_0, \tilde{C}_1)$$

where U_0 is a utility function of the form (2) and C_t is their own consumption, and S is a social welfare function of the form (4) and $\tilde{C}_t$ is the other agent's consumption. ϕ is a weight that characterizes the degree of altruism. Absent any possibility of coordination, each will choose a savings rate that maximizes U_0. This savings rate will not generally maximize S. If in this position they were then offered a policy that would jointly increase both their savings rates, they would gladly vote to accept it. The loss in utility, U_0, would more than be replaced by the gain in S.

Even in the absence of altruism forward looking policies might survive. Once in office, candidates would do well to note that tastes change over time. If they wish to remain in office, they may not want to maximize constituents' current utility, but their utility at the time of the next election.[11] Other policy makers, especially Presidents, seem to be very concerned about how they will be judged by history. The perspective of history, however, is the perspective of the future. These policy makers may also prove to be more far sighted than the average voter.

Finally, there are many social decisions that lie outside of the democratic process. Governors of the Federal Reserve System and Federal judges are appointed for long terms partly to allow these decision makers to take a longer term perspective. Beyond government, many individuals, such as parents and teachers, often find themselves placed in situations where their decisions affect the welfare of others at various points in time. These individuals too may behave like our social planner.

Given the complexities of the political constraints cited above, our approach is to first characterize the unconstrained optimum.

4.5 Consumer Sovereignty

Much of welfare economics as it now stands is based on the idea of consumer sovereignty: no one is in a better position to know what an agent wants than the agent himself. This is the belief that underlies the the quote of Arrow and Kurz presented in the introduction. Some may find our analysis trou-

[11] It may be possible to estimate the extent to which agents discount the past by observing how politicians voting patterns vary with their remaining term in office.

bling because it appears to dispense with consumer sovereignty, and seems to replace the dictatorship of present with the dictatorship of the social planner.

Consumer sovereignty, however, is not an axiom upon which to base a theory but a conclusion that one draws after one has understood the problem. There are many cases in which society chooses to discard consumer sovereignty. Parents are given power over children. Suicide is discouraged.

The reason that consumer sovereignty is so appealing is that there are well known problems with real world social planners. In practice, planners rarely have the information necessary to maximize the utility of those under their care or sufficient incentive to do the job well. All too often they have agendas of their own which interfer with their trusteeship. Society must therefore take caution in how it implements the dictates of any social welfare function. It is often best to let individuals make their own mistakes, then to impose society's mistakes upon them.

This is more a question of how best to implement the social optimum than a question of what form of the social welfare function should take. The real lesson of consumer sovereignty is that it is often best to alter incentives and allow agents to make their own decisions than to dictate actions. In the consumption saving example of the last section, society does not need to make all of the consumption and savings decisions for the individual in order to maximize a social welfare function of the form (4). The government only needs to determine the appropriate tax rates. Consider the condition for optimal consumption when an agent maximizes a utility fuction of the form (2):

$$u'(C_1) = \beta(1+r)u'(C_2).$$

If the government wishes that the agent behave as if he discounted future felicity by $\beta^* > \beta$, it needs to subsidize saving so that the agent perceives a gross rate of return of $(1+r)\beta^*/\beta$ rather than $1+r$. The government already practices this type of intervention in the price system. In the case of income taxes and capital gains taxes, however, the government generally pushes returns in the opposite direction.

To the extent that the social optimum can be implemented through the price system, it would not appear to do terrible damage to the principle of consumer sovereignty.

4.6 On the Interpretation of the Utility Function

Until this point we have treated the utility function U_t as the welfare of the agent in period t, but we have been deliberately vague as to what we meant by this. For most of the argument all we need is that U_t represents the agent's ranking of consumption sequences in period t. This is enough to motivate retrospective time inconsistency and the need to aggregate the perspectives of different periods. Whereas the precise weights that the planner places on the different perspectives will require some notion of cardinal utility, the presumption that the planner should tilt the scale towards the future requires only an aversion to dictatorship.

In order to better understand the welfare problem, however, it is useful to consider how these choice rankings relate to the level of satisfaction experienced by the agent over time. This is essentially a question of why the future and the past matter to an agent in the present. As Loewenstein [1992] has pointed out, there are essentially two answers to this question, one due to Jevons and the other to Samuelson.[12] In order to avoid confusion we will use the words pleasure and happiness when referring to this experience of satisfaction over time, and contrast these with utility U_t and felicity u_t.

According to the Jevonian perspective the future and the past matter because they directly influence the agent's present happiness. Agents feelings today are influenced by memories of the past and anticipations of the future. In this view, U_t represents the pleasure that the agent experiences in period t. Future and past felicity are discounted in U_t because these experiences are less immediate than current ones.

Given the Jevonian perspective the social welfare function (4) counts felicity multiple times, but this is because felicity is experienced multiple times. Felicity at date t is first experienced as anticipation at dates $s < t$, then as part of current date t enjoyment, and finally as memory at dates $s > t$. The weights on felicity in (4) reflect these multiple perspectives, as well as the weights that the social planner assigns to these perspectives. According to the Jevonian perspective, the dictatorship of the present is extremely myopic. It focuses solely on current pleasure and completely ignores all future pleasure.

At the other extreme, the Samuelsonian perspective states that the pleasure that an agent experiences in period t is associated with the period t

[12]See Caplin and Leahy [2000] for a more detailed discussion of these views and their relationship to welfare theory.

felicity $u(c_t)$. Only current consumption makes the agent happy. In order to come to grips with the fact that agents care about future consumption but receive utility only from current consumption, the Samuelsonian perspective divides the period t self into two components: a decision maker and a consumer. The consuming self experiences period t felicity. The decision making self essentially acts as an internal social planner who aggregates the preferences of the consuming selves of various periods. In this view, U_t represents the preferences of the internal social planner at date t.[13]

According to the Samuelsonian perspective, when $\delta \neq \beta^{-1}$ the internal social planner's rankings of consumption sequences changes over time, and the social welfare function (4) represents an aggregation of these rankings. Consumption at date t provides pleasure only at date t, but the weight placed on date t felicity reflects an average of the weights placed on date t felicity by internal social planners at all dates.

4.7 Related Literature

In the rare cases in which the past is explicitly considered or in which the perspectives of different temporal selves are considered, economists commonly assume that agents discount the past. Pigou [1929], Wolf [1970], and Page [1977] consider utility functions in which the past enters with diminishing weight. Pigou and Page note that this implies that tastes change over time.

In the literature on time inconsistency, Phelps and Pollak [1968] and Laibson [1996] discuss the welfare effects of changes in the savings rate in the presence of hyperbolic discounting. Each considers the effect of these changes on the utility of the various temporal selves. In each case, future selves are assumed to place no weight on past consumption. Neither note that this assumption implies that tastes change even if preferences are time consistent.

In models with overlapping generations it is common to consider how parents and children care about each others utility. In Barro's [1974] model of imperfect altruism, for example, parents care about the utility of their children, but children do not care about the utility of their parents. This

[13]Kahneman, Wakker and Sarin [1997] argue that the external social planner should ignore the perspective of the internal social planner and focus solely on the various consuming selves. The problem with this approach is that it entirely divorces intertemporal social choice from intertemporal private choice.

gives rise to a utility function for each generation in which future felicity is discounted exponentially, but past felicity is given no weight. Others have considered two-sided altruism (Abel [1987], Kimball [1987]). Here the common assumption is that $\delta < \beta^{-1}$ which gives rise to a discount factor kinked at the present as in Figure 2. The focus of these papers, however, is on the validity of Ricardian equivalence, not social welfare.[14]

The literature contains several other arguments in support of a lower social discount rate. These arguments tend to focus on intergenerational conflicts. Many authors have argued that it is troublesome to rank the needs of the present generation above those of future generations (e.g. Ramsey [1928] and Solow [1974]). Others have argued for a reduced social discount rate based on the fact that some private decisions have external effects on future generations. One example is capital left after death (Pigou [1929]).

These arguments are really dynamic extensions of static welfare arguments. Shell [1971] has shown that the overlapping generations model can be reinterpreted as a static Arrow-Debreu model in which all agents trade simultaneously. The first argument is therefore no different than the static argument for income redistribution. The second argument rests on a missing market. In the example, it is the market for annuities. In contrast to these arguments, our story has no static counterpart. Time is essential for tastes to change.

These intergenerational arguments also focus attention on long horizons. These arguments make a big difference over a span of 50 to 100 years, but they tend to be ignored over shorter horizons such as the business cycle. There is a tendency in the literature to only discuss the issue of social discounting when considering long term issues such as global warming or resource extraction. The argument presented in this paper, in contrast, is operative in the short run as well as the long run. It applies as well to monetary policy as social security, as well to unemployment insurance as to the environment.

[14]There is also a literature on welfare in overlapping generations models. This literature focus on the question of whether expected lifetime utility should be evaluated unconditionally (Muench [1977]) or conditional on the information available in the first period of life (Peleg [1982]). In our view, both approaches suffer from a dictatorship of the young and ignore the fact that agents may view consumption profiles quite differently when old.

5 Conclusion

We formulate the problem of a social planner attempting to maximize welfare over time. We argue against the currently practice of equating the social discount rate and the private discount rate. Our analysis has immediate implications for public policy: agents discount the future too much and governments should promote future oriented policies.

6 Appendix: Proof of Proposition 3

Let $\delta_1 > \delta_2$. Suppose $C^T \in \mathcal{P}_{\delta_1}$. We show that $C^T \in \mathcal{P}_{\delta_2}$. Since $C^T \in \mathcal{P}_{\delta_1}$, there exist weights $\{\alpha_t\}$ such that C^T maximizes S subject to the constraint on consumption. The $\{\alpha_t\}$ are associated with weights on felicity $\{q_t\}$. We show that for $\delta_2 < \delta_1$, there exist non-negative weights $\{\hat{\alpha}_t\}$ which yield weights on felicity $\{q_t\}$. Therefore with δ_2 C^T solves the social planner's problem with weights $\{\hat{\alpha}_t\}$, and $C^T \in \mathcal{P}_{\delta_2}$.

The proof is by induction. We begin with $T = 2$. $S(U_1, U_2) = \alpha_1 U_1(C^2) + \alpha_2 U_2(C^2) = (\alpha_1 + \delta_1\alpha_2)u(C_1) + (\beta\alpha_1 + \alpha_2)u(C_2)$. Hence $q_1 = \alpha_1 + \delta_1\alpha_2$ and $q_2 = \beta\alpha_1 + \alpha_2$. To maintain these weights on felicity, the weights on utility must be: $\alpha_1 = \frac{q_1 - \delta_1 q_2}{1-\delta\beta}$ and $\alpha_2 = \frac{-q_2 + \beta q_1}{1-\delta\beta}$. If the $\alpha_1$and α_2 that yield q_1 and q_2 are non-negative for δ_1, then the $\hat{\alpha}_1$ and $\hat{\alpha}_2$ that yield q_1 and q_2 are non-negative for $\delta_2 < \delta_1$.

Suppose now that the Proposition is true for $T = N$. Consider $T = N + 1$. In this case, $q_t = \sum_{s=0}^{t-1} \alpha_s \beta^{t-s} + \alpha_t + \sum_{s=t+1}^{N+1} \alpha_s \delta_1^{s-t}$. Consider the first N of these equations, the equations for $\{q_1, \ldots, q_N\}$, and replace $\alpha_N + \alpha_{N+1}\delta_1$ in these equations with A_N. The resulting system gives the weights on utility $\{\alpha_0, \ldots \alpha_{N-1}, A_N\}$ necessary to generate weights on felicity $\{q_1, \ldots, q_N\}$. Since the proposition holds for $T = N$, given $\delta_2 < \delta_1$, there exist non-negative $\{\hat{\alpha}_0, \ldots \hat{\alpha}_{N-1}, \hat{A}_N\}$ that generate $\{q_1, \ldots, q_N\}$. It remains to show that $\hat{A}_N$ can be decomposed into non-negative α_N and α_{N+1}. We can use the equations for q_{N+1}, q_N, and q_{N-1} to solve for α_N and α_{N+1} in terms of the $q's$. It is easily seen that the resulting α_N and α_{N+1} are positive. This completes the proof.□

References

[1] Abel, Andrew, "Operative Gift and Bequest Motives," *American Economic Review*, vol. 77, 1037-1047.

[2] Ainslie, G., 1992, *Picoeconomics*, Cambridge, U.K.: Cambridge University Press.

[3] Allais, M., 1947, *Economie et Intérêt*, Paris: Imprimerie Nationale.

[4] Arrow, K., and M. Kurz, 1970, *Public Investment, the Rate of Return, and Optimal Fiscal Policy*, Baltimore: Johns Hopkins University Press.

[5] Barro, R., 1974, "Are Government Bonds Net Wealth?" *Journal of Political Economy*, vol. 81, 1095-1117.

[6] Barro, R., 1997, "Myopia and Inconsistency in the Neo-Classical Growth Model," NBER Working Paper No. 6317.

[7] Barro, R., and Xavier Sala-i-Martin, 1995, *Economic Growth*, New York: McGraw Hill.

[8] Becker, Gary, and George Stigler, 1977, "De Gustibus Non Et Disputandum," *American Economic Review*, 76-90.

[9] Becker, Gary, and Kevin Murphy, 1988, "A Theory of Rational Addiction," *Journal of Political Economy*, vol. 96, 675-7000.

[10] Becker, Gary, and Casey Mulligan, 1997, "The Endogenous Determination of Time Preference," *Quarterly Journal of Economics*, vol. 112, 729-758.

[11] Böhm-Bawerk, E., 1959, *The Positive Theory of Capital*, South Holland, Il: Libertarian Press.

[12] Caplin, Andrew, and John Leahy, 2000, "Welfare Theory and Experienced Utility," New York University working paper.

[13] Chari, V.V. and P. Kehoe, 1999, "Optimal Fiscal and Monetary Policy," NBER Working Paper No. 6891.

[14] Crovitz, H., and H. Schiffman, 1974, "Frequency of Episodic Memories as a Function of their Age," *Bulletin of the Psychonomic Society*, vol. 4, 517-518.

[15] Deaton, A., 1992, *Understanding Consumption*, Oxford: Clarendon Press.

[16] DeGenova, M., 1992, "If You Had Your Life to Live Over Again: What Would You Do Differently," *International Journal of Aging and Human Development*, vol. 34, 135-143.

[17] Fisher, I., 1930, "The Theory of Interest," London: MacMillan.

[18] Kahneman, D., P. Wakker, and R. Sarin, 1997, "Back to Bentham: Explorations of Experienced Utility," *Quarterly Journal of Economics*, vol. 112, 375-406.

[19] Kimball, Miles, 1987, "Making Sense of Two-Sided Altruism," *Journal of Monetary Economics*, vol. 20, 301-326.

[20] Koopmans, T., 1960, "Stationary Ordinal Utility and Impatience," *Econometrica*, vol. 28, 287-309.

[21] Laibson, D., 1996, "Hyperbolic Discount Functions, Undersaving, and Savings Policy," NBER Working Paper No. 5635.

[22] Loewenstein, G., 1987, "Anticipation and the Valuation of Delayed Consumption", *Economic Journal*, vol.97, 666-684.

[23] Loewenstein, G. 1992, "The Fall and Rise of Psychological Explanations in the Economics of Intertemporal Choice," in G. Loewenstein and J. Elster, eds., *Choice over Time*, New York: Russel Sage Foundation, 3-34.

[24] Lucas, R. E. Jr., and N. Stokey, 1983, "Optimal fiscal and Monetary Policy in an Economy without Capital," *Journal of Monetary Economics*, vol. 12, 55-92.

[25] Malinvaud, E., 1985, "Lectures on Microeconomic Theory," New York: North Holland.

[26] Marglin, S., 1963, "The Social Rate of Discount and the Optimal Rate of Investment," *Quarterly Journal of Economics*, vol. 77, 100-109.

[27] Mas-Collel, Andreu, Michael Whinston, and Jerry Green, 1995, *Microeconomic Theory*, Oxford: Oxford University Press.

[28] O'Donoghue, Ted, and Matthew Rabin, 1999a, "Doing it Now or Later," *American Economic Review*, vol. 89, 103-124.

[29] O'Donoghue, Ted, and Matthew Rabin, 1999b, "Addiction and Self Control," U. C. Berkeley, mimeo..

[30] Muench, T., 1977, "Optimality, the Interaction of Spot and Futures Markets, and the Nonneutrality of Money in the Lucas Model," *Journal of Economic Theory*, vol. 15, 325-244.

[31] Page, T., 1977, "Conservation and Economic Efficiency," Baltimore: Johns Hopkins University Press.

[32] Peleg, D., 1982, "Informational Diversity over Time and the Optimality of Monetary Equilibria," *Journal of Economic Theory*, vol. 28, 255-274.

[33] Phelps, E., and R. Pollak, 1968, "On Second-Best National Saving and Game-Equilibrium Growth", *Review of Economic Studies,* vol. 35. 573-598.

[34] Pigou, A., 1929, *The Economics of Welfare,* third edition, London: MacMillan and Co..

[35] Ramsey, F., 1928, "A Mathematical Theory of Saving," *Economic Journal*, vol. 38, 543-559.

[36] Rawls, J., 1999. *A Theory of Justice: Revised Edition,* Cambridge: Harvard University Press.

[37] Romer, D. 1996, *Advanced Macroeconomics*, New York: McGraw Hill.

[38] Sen, A., 1961, "On Optimizing the Rate of Saving," *Economic Journal*, vol. 71, 479-496.

[39] Shell, K., 1971, "Notes on the Economics of Infinity," *Journal of Political Economy*, vol. 78, 1002-1011.

[40] Shelling, T., 1984, "Self-Command in Practice, in Policy, and in a Theory of Rational Choice," *American Economic Review Papers and Proceedings*, vol. 74, 1-11.

[41] Solow, R., 1974, "The Economics of Resources or the Resources of Economics," *American Economic Review Papers and Proceedings*, vol. 64, 1-14.

[42] Strotz, R., 1956, "Myopia and Inconsistency in Dynamic Utility Maximization," *Review of Economic Studies*, vol. 23, 165-180.

[43] Wolf, Charles, 1970, "The Present Value of the Past," *Journal of Political Economy,* vol. 78, 783-792.