

The Social Discount Rate

Andrew Caplin and John Leahy*

January 2004

Abstract

What discount rate should be applied to social investments? It is standard to use the market interest rate, thereby respecting private preferences. We show that application of this “revealed preference” criterion rests on faith, not logic. It is justified only if preferences over all choices, including *past choices*, are time invariant. The conditions for this to be true are stringent. Under more reasonable conditions, policy makers should be more patient than private citizens, whose choices define the most short-sighted Pareto optimum.

1 Introduction

In the early days of the field, many economists argued that policy makers should be more patient than private citizens. Yet most of their arguments were paternalistic, such as Ramsey’s claim that it was “ethically indefensible” to discount the future (Ramsey [1928]). Assertions of this sort entirely neglect personal preferences, and proved vulnerable to the “dynamic revealed preference” critique that effectively ended the debate:

“It is hard to see why the revealed preference of individuals should be disregarded in the realm of time, where it is accepted, broadly

*New York University. We thank Andy Abel, Fernando Alvarez, Marco Bassetto, Susanto Basu, Gary Becker, Kfir Eliaz, Edward Glaeser, Jeffrey Campbell, Jean-Jacques Laffont, David Laibson, Robert Lucas Jr., Casey Mulligan, Barry Nalebuff, Debraj Ray, Mario Rizzo, Nancy Stokey, Robert Solow, Lars Svensson, and Richard Zeckhauser for helpful comments and discussions. Caplin thanks the C.V. Starr Center at NYU and Leahy thanks the National Science Foundation for financial support.

speaking, in evaluating current commodity flows.” (Arrow and Kurz [1970], p. 12)

In this paper, we demonstrate that the dynamic principle of revealed preference is far less compelling than its static counterpart. Furthermore, we show that under reasonable conditions, policy makers should be more patient than private citizens. Unlike earlier arguments to this effect, ours is non-paternalistic and is therefore immune to the revealed preference critique.

The basic model of evolving preferences introduced in section 2 is due to Strotz [1956]. He allowed for utility flows from past consumption to cover “the possibility that a person is not indifferent to his consumption history but enjoys his memories of it” (Strotz [1956], p.167). His model enables us to illustrate our critique of the dynamic principle of revealed preference, which derives from an asymmetry in intertemporal control rights. Decision makers are powerless to alter earlier decisions. As a result, nothing can be gleaned from present behavior concerning preferences over the past. Preferences revealed in the market do not adequately represent tastes.

In section 3 we study the issue of impatience. We characterize the extent of disagreement over time in a simple variant of the Strotz model, and provide conditions guaranteeing that the revealed preference approach to policy selects the most impatient Pareto optimum. In the standard case in which future consumption is discounted, our impatience result holds unless current consumption grows more important as it passes into the past. We argue that such growth is unrealistic.

While the argument of section 2 is completely general, our argument in section 3 concerning the impatience of revealed preference is model-specific. Section 4 opens up the question of how general is the result. We discuss also whether or not it is possible to observe preferences over past consumption. Section 5 concludes.

While our rejection of the dynamic principle of revealed preference represents a departure from current practice, it has historical roots. In the work that laid the foundations of modern welfare economics, Pigou [1952], preceded us not only in our criticism of using choice behavior as a guide to policy, but also in our hypothesis concerning the implied directional bias:

“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)

2 Revealed Preference and Time

We consider a standard finite horizon model of consumption and investment. The discrete time intervals are indexed by t , $1 \leq t \leq T$. We define $C = (c_1, \dots, c_T) \in R_+^T$ to be a generic consumption program chosen from a rectangular consumption set, $A = \prod_{t=1}^T A_t$, with each set $A_t \subset R_+$ comprising a non-empty interval (possibly unbounded). The question of interest concerns choice among consumption programs in a convex production possibility set, $Y \subset R_+^T$.

Following Strotz [1956], we allow a different utility function for each point in time, $U^t : A \rightarrow R$, for $1 \leq t \leq T$, representable in the following simple form,

$$U^t(c_1, \dots, c_T) = \sum_{s=1}^{t-1} \lambda(s-t)u_s(c_s) + u_t(c_t) + \sum_{s=t+1}^T \lambda(s-t)u_s(c_s). \quad (1)$$

In this equation, all functions $u_s : A_s \rightarrow R$ are strictly monotone, concave, and continuous.

This formulation borrows from Strotz not only the assumption that utility depends on past as well as present and future consumption, but also his assumption that the weight on any given period depends only on its temporal relationship to the present (Strotz [1956], equation (2), p.167). For integers $d \in [1-T, T-1]$, the discount factors $\lambda(d) \geq 0$ define the weight given in any given period to consumption in a period that differs from it by fixed amount d . As does Strotz, we normalize to $\lambda(0) = 1$, implying that the weight on own period consumption is strictly positive. Generalizations are considered in section 4.

In the two period case, it is standard to solve consumption-investment models in the Fisher diagram, as in Figure 1. The convex set of productively feasible consumption choices is represented in the diagram, as are indifference curves representing period 1 preferences, $U^1(c_1, c_2) = u_1(c_1) + \lambda(1)u_2(c_2)$. The individual in period 1 chooses the point P on the production possibility frontier associated with the highest indifference curve of the utility function.

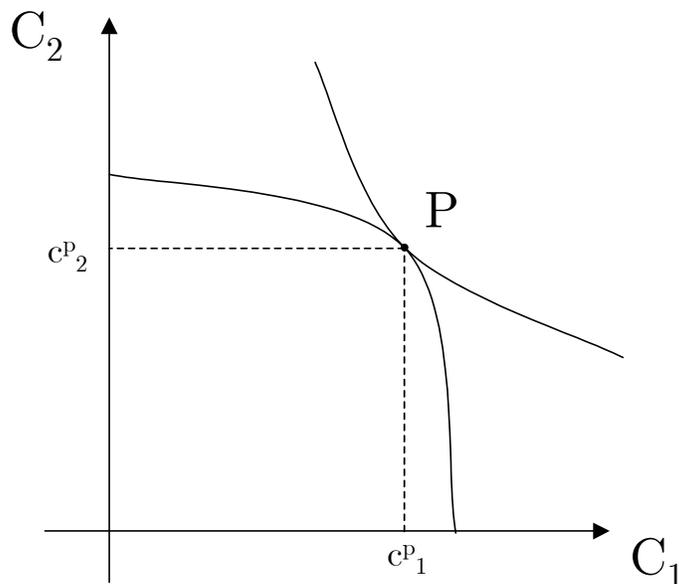


Figure 1: Fisher Diagram

To read a policy moral into the Fisher diagram, one stresses the analogy with static choice. In a static setting one would automatically identify the optimal social choice as that yielding highest private welfare. In precisely this spirit, it seems natural to assert that point $P = (c_1^P, c_2^P)$ is not only privately but also socially optimal. It appears that dynamic questions concerning both private and social choice can be settled entirely by reliance on individual preferences.

Compelling as this argument may appear, it tells only half the story. The indifference curve in figure 1 represents period 1 preferences over consumption programs, and it is from this viewpoint that point P is optimal. Yet the diagram contains no reference whatsoever to period 2 preferences over these programs, $U^2(c_1, c_2) = \lambda(-1)u_1(c_1) + u_2(c_2)$. The standard Fisher diagram is therefore an incomplete summary of the Strotz model.

Why does the Fisher diagram exclude any reference to period 2 preferences? It is because it respects the principle of revealed preference. Information irrelevant to decision making is ignored, and period 2 preferences are of just this form. After the decision maker in period 1 selects c_1^P , the decision maker in period 2 is left only with the trivial choice of whether or not to

consume all remaining resources. With standard assumptions such as those of Strotz, the answer is in the affirmative. The arrow of time deprives the individual in period 2 of choice as a means of expressing preference.

While neglecting period 2 preferences makes sense from the viewpoint of private decision making, social policy is quite a different matter. The fact that only period 1 and not period 2 preferences over consumption streams can be put into effect should not be confused with a policy doctrine. The asymmetry in intertemporal control rights should no more determine policy in this case than it would were we dealing with two distinct individuals. It is in this respect that the dynamic principle of revealed preference differs from its static counterpart, since there is a single unified perspective on the static problem.

The fact that revealed preference does not allow one to gather complete information on period 2 preferences opens a gaping hole in the current rationale for subordinating public to private choice. A priori, we have no reason to believe that these unobserved preferences coincide with the period 1 preferences. If they do not coincide, there will be disagreement over time concerning the optimal consumption program. The individual in period 2 would be better off had a different choice been made in period 1. Unfortunately, this decision is sunk from the period 2 viewpoint, and acts as a constraint on choice. We turn now to the general model to analyze how preferences change over time, and the implications for social choice.

3 Choice and Impatience

We provide conditions such that all weights $\lambda(d)$ in equation (1) are of a simple form. This yields a correspondingly simple characterization of the conditions under which private choice yields the most impatient Pareto optimum. We close the section by discussing the plausibility of these conditions.

3.1 Time Consistency

Our simple characterization of preferences rests on two assumptions concerning the extent of agreement over time. The first assumption is time consistency of the sort analyzed by Strotz. In behavioral terms, the importance of this assumption is that it allows the decision maker to stick with plans for the future made at some earlier point of time. If it did not hold, there would

be disagreement over time concerning optimal future choices, which would render the dynamic principle of revealed preference ill-defined.

The formal treatment of time consistency in definition 1 invokes the qualifier *prospective*, since we will in definition 2 introduce a distinct form of time consistency that looks at the past rather than the future.

Definition 1 A family of utility functions $U^t : A \rightarrow R$, $1 \leq t \leq T$, is *prospectively time consistent* if, given any date $\tau \in [1, T]$, consumption programs $C, C' \in A$ with $c_s = c'_s$ for all $s < \tau$, and $t \leq \tau$,

$$U^t(C) \geq U^t(C') \text{ if and only if } U^1(C) \geq U^1(C').$$

The assumption of prospective time consistency brings order to preferences over future consumption. It corresponds to a case in which one future is chosen over another because it is always viewed as better than the alternative. Definition 2 formalizes the natural retrospective counterpart to definition 1. The family of utility functions will be referred to as *retrospectively* time consistent if, whenever two consumption programs differ only before date τ , they are ranked identically from the viewpoint of all ensuing periods (and in particular the final period). As with its prospective counterpart, retrospective time consistency limits the degree to which the current period can play a role in determining preferences. Last year's consumption may be more important than that of the previous year, but only to the degree that it will be viewed as more important from even later dates.

Definition 2 A family of utility functions $U^t : A \rightarrow R$, $1 \leq t \leq T$, is *retrospectively time consistent* if, given any date $\tau \in [1, T]$, consumption programs $C, C' \in A$ with $c_s = c'_s$ for all $s \geq \tau$, and $t \geq \tau$,

$$U^t(C) \geq U^t(C') \text{ if and only if } U^T(C) \geq U^T(C').$$

In contrast with definition 1, note the use in definition 2 of the weak inequality, $c_s = c'_s$ all $s \geq \tau$. Only if the difference between two programs lies in the *strict* past from the viewpoint of period τ must the period τ and period T views agree. For this reason, definition 2 is consistent with the past being irrelevant to current utility.

3.2 A Characterization

Strotz [1956] shows that geometric discounting is necessary and sufficient for prospective time consistency for a family of utility functions satisfying (1). Our result is the natural generalization of this result for utility functions that are both retrospectively and prospectively time consistent.

Proposition 1: A family of utility functions satisfying equation (1) is retrospectively and prospectively time consistent if and only if there exist constants $\beta > 0, \delta > 0$ and $K \geq 0$ such that $\lambda(d) = \beta^d$ for all integers $d \in [1, T - 1]$ and $\lambda(d) = K\delta^{d+1}$ for all integers $d \in [-T + 1, -1]$.

Proof: *Sufficiency.* Given that $c_s = c'_s$ for all $s < \tau$ and $t \leq \tau$, the only difference between the condition for $U^t(C) \geq U^t(C')$ and $U^1(C) \geq U^1(C')$ is the additional factor β^{t-1} that multiplies both sides. Such multiplications leaves the order unchanged, establishing prospective time consistency. Retrospective time inconsistency follows from an equivalent argument based on powers of δ rather than β .

Necessity. We show first that if the stated conditions on $\lambda(d)$ for $d \in [1, T - 1]$ are not met, prospective time consistency fails. We note first that $\lambda(d) > 0$ for all $d \geq 1$. If $\lambda(d) = 0$ for some $d \geq 1$, we pick C interior to A and $x > 0$ to derive the following contradiction to prospective time consistency,

$$U^1(C) = U^1(C + xe_{d+1}) \text{ but } U^{d+1}(C) < U^{d+1}(C + xe_{d+1}),$$

where e_k is the k^{th} unit vector. The characterization $\lambda(d) = \beta^d > 0$ for all integers $d \in [1, T - 1]$ follows provided $\frac{\lambda(d+1)}{\lambda(d)} = \frac{\lambda(d)}{\lambda(d-1)}$ all $d \geq 1$, where $\lambda(0) \equiv 1$. If $T = 2$, the statement is vacuous. If $T \geq 3$, suppose to the contrary that $\frac{\lambda(d+1)}{\lambda(d)} \neq \frac{\lambda(d)}{\lambda(d-1)}$ for some $d \geq 1$. Now pick C interior to A , and consider any $(x, y) \neq 0 \in R^2$ such that

$$U^1(C) = U^1(C + xe_{d+1} + ye_{d+2}) \equiv U^1(C') :$$

that $(x, y) \in R^2$ with this property can be found follows from the continuity and strict monotonicity of all functions $u_t(c_t)$. Substitution in equation (1) yields,

$$\frac{\lambda(d+1)}{\lambda(d)} [u_{d+2}(c_{d+2} + y) - u_{d+2}(c_{d+2})] = u_{d+1}(c_{d+1} + x) - u_{d+1}(c_{d+1}). \quad (2)$$

Now note that $U^2(C) \neq U^2(C')$, since the corresponding indifference condition involves substituting the distinct number $\frac{\lambda(d)}{\lambda(d-1)}[u(c_{d+2}) - u(c'_{d+2})]$ on the left hand side of equation (2), contradicting prospective time consistency.

Finally, we show that if the stated conditions on $\lambda(d)$ for $d \in [-1, -T+1]$ are not met, retrospective time consistency fails. Define $\lambda(-1) \equiv K$. Note first that $K > 0$ implies $\lambda(d) > 0$ for all integers $d \in [-T+1, -1]$. If $T = 2$, the statement is vacuous. If $T \geq 3$, let $\bar{d} \in [-2, -T+1]$ be the largest integer such that $\lambda(\bar{d}) = 0$. Once again take C interior to A and $x > 0$ and note that,

$$U^{\bar{d}}(C) < U^{\bar{d}}(C + xe_1) \text{ but } U^{\bar{d}+1}(C) = U^{\bar{d}+1}(C + xe_1), \quad (3)$$

contradicting retrospective time consistency. Given that $\lambda(d) > 0$ for all integers $d \in [-T+1, -1]$, constancy of $\frac{\lambda(d)}{\lambda(d-1)} = \delta > 0$ follows as in the prospective case. Finally, suppose $K = 0$. In this case, retrospective time consistency requires $\lambda(d) = 0$ all $d < 0$ by defining $\bar{d} \in [-2, -T+1]$ be the largest integer such that $\lambda(\bar{d}) > 0$, and flipping inequality (3). ■

One interesting point to note is that discounting into the future must be exponential, while that on the past can be hyperbolic (Laibson [1997]). There can be a once-off change in the importance of an event as it slips into the past. After this, its weight must change at a constant rate.¹

3.3 The Pareto Set

Proposition 1 implies that there is intertemporal disagreement unless $K = 1$ and $\delta = \beta^{-1}$. In the face of such disagreement, the natural procedure is to study the Pareto set. It is clear that the revealed preference solution is in this set, since it is optimal from the viewpoint of period 1 preferences. The first welfare theorem is therefore valid in our economy. However, the second welfare theorem is not valid, since the choices made in period 1 impose external effects that constrain agents in later periods. The question of interest is how the Pareto optimum selected by private choice differs from those that are ruled out by constraints of this form.

Proposition 2 establishes conditions under which the actual consumption program selected by the agent is the most impatient of all Pareto optima.

¹This additional degree of freedom follows from the use in definition 2 of the weak inequality, $c_s = c'_s$ all $s \geq \tau$. With the strict inequality, we would conclude that $K = 1$. The reason for not making this stronger assumption is that it rules out the case in which past consumption has no impact on current utility.

The proof involves only the simplest of algebra. Given the convexity of the production set and the concavity of all utility functions, the utility possibility set for the economy is convex. Hence the Pareto optimal allocations coincide with the maximizers of simple (non-paternalistic) Bergsonian social welfare functions which define the supporting hyperplanes to this set,

$$S \equiv \sum_{t=1}^T \alpha_t U^t, \quad (4)$$

with $\alpha_t \geq 0$, for all t , and $\sum_{t=1}^T \alpha_t = 1$. Maximizing S is equivalent to maximizing $\sum_{t=1}^T q_t u_t(c_t)$, where,

$$q_t = \sum_{s=1}^{t-1} \alpha_s \beta^{t-s} + \alpha_t + K \left[\sum_{s=t+1}^T \alpha_s \delta^{s-t} \right]. \quad (5)$$

Straightforward manipulation of equation (5) yields,

$$q_{t+1} - \beta q_t = \alpha_{t+1}(1 - K\beta\delta) + K \left[\sum_{s=t+2}^T \alpha_s \delta^{s-t} \left(\frac{1}{\delta} - \beta \right) \right],$$

for $t \leq T - 1$. This equation makes it straightforward to identify conditions under which the under which the revealed preference solution is the most short-sighted Pareto optimum.

Proposition 2: Consider any family of utility functions satisfying the conditions of proposition 1. If $\max[\delta, K\delta] \leq \beta^{-1}$, then,

$$\frac{q_{t+1}}{q_t} \geq \beta, \quad (6)$$

for all integers $t \in [1, T - 1]$. If $\max[\delta, K\delta] < \beta^{-1}$, inequality (6) is strict provided $\alpha_s > 0$ for some $s \geq t + 1$.

To interpret this result, note that the ratio q_{t+1}/q_t specifies the relative weight placed on $u_{t+1}(c_{t+1})$ as opposed to $u_t(c_t)$ by the social welfare function S in equation (4). Proposition 2 implies that if $\delta < \beta^{-1}$ and $K \leq 1$, all Pareto optima other than the revealed preference optimum place at least as high a relative weight on $u_{t+1}(c_{t+1})$ as opposed to $u_t(c_t)$ for all t , and a strictly higher weight for some t . Moreover, if any weight is placed on the final period perspective, the weight on $u_{t+1}(c_{t+1})$ as opposed to $u_t(c_t)$ is strictly higher for all t .

3.4 Discounting the Past

The critical question concerns the plausibility of the condition on discount rates assumed in proposition 2, $\max[\delta, K\delta] \leq \beta^{-1}$. There is one specific context in which this subject has attracted the attention of the economics profession. If one interprets each period's consumption as corresponding to one generation in a dynastic family, the standard assumption is that parents care about the utility of their children, but children do not care about the utility of their parents (Barro [1974]). This gives rise to a utility function with geometric discounting of the future, and no weight on the past. Abel [1987] and Kimball [1987] consider two-sided altruism, with the explicit assumption that $K = 1$ and $\delta < \beta^{-1}$. Proposition 2 is satisfied in all of these cases.

When one interprets the model as corresponding to the lifetime of a single agent, economists have had little to say about the importance of the past. In those rare cases in which past consumption has been mentioned, it has been assumed that $K = 1$ and $\delta < 1$ (Wolf [1970], Page [1977], and Ray and Wang [2001]). No one has ever included a growing weight as the event passes from the present ever further into the past.

What psychological arguments are relevant to the relationships between β on the one hand and δ and K on the other? In terms of the Strotz model, the argument for $K \leq 1$ seems clear-cut: if this were not true, then one would strictly prefer to have had a wonderful meal last month rather than having one to look forward to this month. Given the universal assumption that $\beta^{-1} < 1$, the case for $\delta < \beta^{-1}$ seems equally powerful. This is the case in particular if memory is the chief method of generating utility from past consumption, as suggested by Strotz [1956]. Memory is not only imperfect, but also tends to fade over time. Just as memories tend to fade, so may their impact on current utility.

Overall, we hypothesize that the conditions of proposition 2 apply whenever an individual with Strotzian preferences would prefer to have consumed less in the past, and to have more wealth left over for present and future consumption. We conjecture that this is the normal state of affairs. Intuition suggests that many would prefer in the past to have invested more in the past, so that in the future they could consume more. In support of this hypothesis, DeGenova [1992] found that when people looked back on their lives, they most regretted not spending more time on education.

4 Open Questions

4.1 How General is Proposition 2?

To what extent would our results on directional bias survive with richer models of changing tastes? The most obvious generalization is to allow for general time separable preferences,

$$U^t(c_1, \dots, c_T) = \sum_{s=1}^T u_s^t(c_s), 1 \leq t \leq T, \quad (7)$$

where all functions $u_s^t : A_s \rightarrow R$ are monotone, concave, and continuous, with u_t^t strictly monotone. A somewhat modified Strotzian representation is available for such preferences, provided they are both prospectively and retrospectively time consistent. However this representation requires the use of two distinct utility functions for period s , one relevant to period s and before, and the other relevant to later periods. Of course, this renders more intricate the appropriate analog of the impatience result in proposition 2. However we conjecture that some version of the result will hold provided utility functions are such that we would prefer to have a great meal next week rather than having already had one last week.

Ultimately, it will be important to go beyond the separable case. Separability is particularly unappealing in the context of preferences over the past, particularly if these are determined by the properties of memory. It is plausible that one remembers better extreme levels of consumption than more accustomed levels. If this is the case, preferences over past consumption will not be separable. One way to capture forces of this sort is to look for a generalization of the recursive model of Koopmans [1960]. The novelty would be that rather than having these preferences separate according only to current and future utility, one would need a third term that captures the past. Formulating such a model and analyzing the implied directional bias are tasks that lie ahead of us.²

²We are grateful to Fernando Alvarez for pointing out the salience of the recursive formulation.

4.2 Can Preferences over the Past be Measured?

If we are to develop a social policy that reflects future preferences, the key question is how to measure preferences over past consumption. To make progress on this matter, it will be important to understand the extent to which the Strotzian utility function *itself* represents an effort at social planning by the individual (the “Samuelsonian perspective” in the language of Loewenstein [1992]) as opposed to directly reflecting enjoyable anticipations and memories (the “Jevonian” perspective). Yet for economists, choice behavior would provide far more evidence than would discussions of the sources of utility. In this light, the most interesting question is whether or not choices are made that would allow preferences over past consumption to be measured. Can preferences over the past be revealed through choice behavior?

Empathic choices offer intriguing opportunities for observing preferences over past consumption. Parents in particular often play the role of social planner for their children. In this role, they work hard to ensure that each child look’s back fondly on his or her upbringing. Parental choices may therefore provide information on the role of memory in utility. Whether or not there are reasonable conditions under which this information can be combined with other choice-based information to understand utility from past consumption is an open question. Of course, even if this could be done, the question remains as to whether the utility function would characterize the child, the parent, or some hybrid defined by parental beliefs.³

5 Conclusion

The principle of revealed preference has been oversold, in particular when used to justify use of the market interest rate in evaluating social investments. We highlight the limitations of the principle in dynamic settings, and provide conditions under which policy makers should be more patient than private citizens.

³Ray and Wang [2001] powerfully make the case that preferences over past consumption *do* have an important influence on current behavior.

References

- [1] Abel, A., 1987, "Operative Gift and Bequest Motives," *American Economic Review*, vol. 77, 1037-1047.
- [2] Arrow, K., and M. Kurz, 1970, *Public Investment, the Rate of Return, and Optimal Fiscal Policy*, Baltimore: Johns Hopkins University Press.
- [3] Barro, R., 1974, "Are Government Bonds Net Wealth?" *Journal of Political Economy*, vol. 81, 1095-1117.
- [4] 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.
- [5] Kimball, M., 1987, "Making Sense of Two-Sided Altruism," *Journal of Monetary Economics*, vol. 20, 301-326.
- [6] Koopmans, T., 1960, "Stationary Ordinal Utility and Impatience," *Econometrica*, vol. 28, 287-309.
- [7] Laibson, David, 1997, "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics* CXII, 443-477.
- [8] 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.
- [9] Page, T., 1977, "Conservation and Economic Efficiency," Baltimore: Johns Hopkins University Press.
- [10] Pigou, A., 1952, *The Economics of Welfare*, London: MacMillan and Co.
- [11] Ramsey, F., 1928, "A Mathematical Theory of Saving," *Economic Journal*, vol. 38, 543-559.
- [12] Ray, D., and R. Wang, 2001, "On Some Implications of Backward Discounting," Working Paper, New York University.

- [13] Strotz, R., 1956, "Myopia and Inconsistency in Dynamic Utility Maximization," *Review of Economic Studies*, vol. 23, 165-180.
- [14] Wolf, C., 1970, "The Present Value of the Past," *Journal of Political Economy*, vol. 78, 783-792.