Time Inconsistency, Self-Control, and Remembrance

STEVEN L. GREEN*

Steven L. Green is Professor of Economics and Statistics, Baylor University (TX).

For I know that nothing good dwells in me—my unspiritual self, I mean—for though the will to do good is there, the ability to affect it is not. The good which I want to do, I fail to do; but what I do is the wrong which is against my will; and if what I do is against my will, clearly it is no longer I who am the agent, but sin that has its dwelling in me (Romans 7:18–20, Revised English Bible).

Economists use the term "time inconsistency" to refer to a situation in which a plan for some periods *t+n* and beyond that is preferred when made at some earlier time *t* (and is in some sense globally optimal) is not preferred when time *t+n* arrives. With respect to an individual agent, attempted solutions to the time inconsistency problem involve self-control issues—ways in which the agent can bind him/herself to adhere to the globally optimal plan.

This paper provides an overview of the time inconsistency problem and the related issue of self-control. It provides a discussion of some self-control strategies identified in the psychology literature and argues that the Christian understanding of human nature—specifically, the idea of remembrance—may be able to provide useful insights to economists working in this area. The paper also specifies and solves a simple model of self-control based on these insights and explores its implications.

The paper proceeds as follows. Section I provides an overview of the basic model of intertemporal choice used by economists. Section II solves a very simple intertemporal choice model to illustrate the time inconsistency problem. Section III reviews the recent time inconsistency literature, while Section IV considers strategies for self-control identified (primarily) by laboratory experiments reported in the psychology literature. The relation of religious faith to these issues is discussed in Sections V and VI, while Section VII specifies and solves a simple model of self-control that is informed by the discussion in previous sections. Section VIII asks whether the conventional approach to teaching introductory economics may create problems for students by leading them to dismiss legitimate self-control strategies. Section IX summarizes the paper and offers some concluding remarks.

__Author's Note:__ The author thanks Chuck North for insightful discussions and two anonymous referees for helpful comments.

I. The Discounted Utility Model

Many interesting issues in economics—macroeconomics in particular—involve decision-makers considering trade-offs over time. A consumer increasing savings today incurs the cost of lower consumption today in exchange for the benefit of increased consumption capability in the future. Firms contemplating investment choices compare current costs to expected future cash flows. When members of Congress debate patent policy, they consider the benefits of competitive markets for existing products (an implication of weak patent protection) against the prospect of future innovation (a product of strong patent protection).

The standard framework used by neoclassical economists in their analysis of intertemporal choice is the discounted utility (DU) model, which was first proposed by Samuelson (1937). Consider a consumer formulating a plan for consumption spending between the present (date 0) and some future date T. Let c_t denote consumption in period t, for t = 0, 1, ..., T. Then the consumer in the DU model solves the following optimization problem (subject to relevant constraints):

$$(1) \qquad \max_{\{c_0, c_1, \ldots, c_T\}} \; U = \sum_{t=0}^{T} \rho_t \cdot u(c_t),$$

where $u(\cdot)$ is the (time invariant) single-period utility function in all periods and ρ_t is the period *t* "discount factor." The standard DU model incorporates the assumption of *exponential discounting* as follows:

$$(2) \qquad \rho_t = \left(\frac{1}{1+\delta}\right)^t,$$

where $\delta > 0$ is the"rate of time preference." Note that equation (2) implies $\rho_0 = 1$. A larger value of δ means a lower value of ρ_t for any given t, which means the consumer is more impatient—i.e., has a stronger preference for present consumption vis-à-vis future consumption. In this case, one can also write $\rho_t = \lambda^t$, where $\lambda = \frac{1}{1+\delta}$. Thus with exponential discounting we have the following discount factors: $\rho_0 = 1$, $\rho_1 = \lambda$, $\rho_2 = \lambda^2$, $\rho_3 = \lambda^3$, . . . , and $\rho_T = \lambda^T$.

Exponential discounting implies that $\frac{\rho_i}{\rho_j}$ depends only on the value of i-j. For example, $\frac{\rho_i}{\rho_{i+1}} = \lambda$ for all i . That is, the relative desirability of utility at different points in time is independent of the proximity of the time periods to the present. For example, the value of utility in period 9 relative to utility in period 10 is no different from the value of utility in period 0 relative to utility in period 1: $\frac{\rho_0}{\rho_1} = \frac{\rho_9}{\rho_{10}} = \lambda$.

An important alternative to exponential discounting is hyperbolic discounting. In this case, the discount factors are given by:

$$(3) \qquad \rho_t = \frac{1}{(1+\alpha \cdot t)^{\gamma/\alpha}}$$

With hyperbolic discounting, discount factors fall more rapidly than exponential discounting for t close to zero and less rapidly as t gets large. That is, $\frac{\rho_0}{\rho_1} > \frac{\rho_9}{\rho_{10}}$, so people are more inclined to prefer a smaller/sooner reward to a larger/later reward the closer to the present the two rewards are offered. A more analytically tractable formulation with this same basic property is quasi-hyperbolic discounting, in which case the set of discount factors $\{\rho_1, \rho_2, . . . , \rho_T\}$ is given by $\{1, \beta\delta, \beta\delta^2, . . . , \beta\delta^T\}$, with $\beta < 1$. In this case, $\frac{\rho_0}{\rho_1} = \frac{1}{\beta\delta}$, while $\frac{\rho_j}{\rho_{j+1}} = \frac{1}{\delta}$ for all $j > 0$.

Strotz (1955–1956) raises the interesting question of whether consumption plans are consistent over time. Let $c_t^{(j)}(W)$ denote the optimal level of period t consumption determined from a plan formulated in period j when the consumer has wealth W at the beginning of period j. Thus in period $j = 0$, the optimal plan is $\{c_0^{(0)}(W), c_1^{(0)}(W), c_2^{(0)}(W), . . . , c_T^{(0)}(W)\}$ where W denotes wealth at the beginning of period 0. The plan formulated in period $j = 2$ yields $\{c_2^{(2)}(W), c_3^{(2)}(W), c_4^{(2)}(W), . . . , c_T^{(2)}(W)\}$, where W denotes wealth at the beginning of period 2. (Although the definition of W is potentially ambiguous, it will always be clear in the discussion below.) Strotz raised the interesting question of whether $c_t^{(j)}(W)$ for a given t is the same for all j. For example, is the optimal level of period 2 consumption in the plan formulated in period 2, $c_2^{(2)}(W)$, equal to the optimal level of period 2 consumption foreseen at period 0, $c_2^{(0)}(W)$? If the answer to Strotz's question is "yes", the period 0 plan is said to be "time consistent". If not, it is "time inconsistent."

Strotz demonstrates that optimal plans in the DU model are time consistent only with exponential discounting—that is, if and only if $\rho_t = \lambda^t$ for all t. Any other pattern of discount factors implies time inconsistency. I will now proceed to illustrate this result in a simple model.

II. Time Inconsistency: A Simple Example[1]

Suppose in period 0 a consumer has a certain level of wealth, W, to allocate to consumption over three periods *(0, 1, and 2)*. One might think of W as including the present value of current and future income along with other assets. For simplicity, assume a zero rate of interest and logarithmic utility in each period. Thus the consumer's problem in period 0 is given by:

$$(4) \qquad \max_{\{c_0,c_1,c_2\}} U = \ln(c_0) + \rho_1 \cdot \ln(c_1) + \rho_2 \cdot \ln(c_2), \quad \text{s.t. } c_0 + c_1 + c_2 = W$$

In period 0, the solution to the above problem is denoted $\{c_0^{(0)}(W), c_1^{(0)}(W), c_2^{(0)}(W)\}$. In period 0 actual consumption will be $c_0^{(0)}(W)$, so the consumer enters period 1 with total wealth of $W - c_0^{(0)}(W)$. In period 1 the consumer again formulates a consumption plan by solving

$$(5) \qquad \max_{\{c_1,c_2\}} U = \ln(c_1) + \rho_1 \cdot \ln(c_2) \quad \text{s.t. } c_1 + c_2 = W - c_0^{(0)}(W)$$

The solution to this problem is denoted $\{c_1^{(1)}(W), c_2^{(1)}(W)\}$. Table 1 shows the optimal levels of consumption in each period according to the period 0 plan and the period 1 plan when initial wealth in period 0 is W_0 . Wealth at the beginning of the second period is $[W_0 - c_0^{(0)}(W_0)]$ in both plans, so that $c_1 + c_2 = [W_0 - c_0^{(0)}(W_0)]$ in either case. We see in the table that $c_1^{(0)}(W) = c_1^{(1)}(W)$ if and only if [2]

$$(6) \qquad \frac{\rho_1}{\rho_1 + \rho_2} \cdot \left[W - c_0^{(0)}(W)\right] = \frac{1}{1+\rho_1} \cdot \left[W - c_0^{(1)}(W)\right].$$

Condition (5) holds if and only if $\rho_2 = \rho_1^2$, which is the case of exponential discounting. If $\rho_2 \neq \rho_1^2$, then $c_1^{(0)}(W) \neq c_1^{(1)}(W)$ and the plans are time inconsistent.

At first glance one might think that this "time inconsistency problem" is really no problem at all for economists. In the above example, for instance, there would appear to be a clear prediction that observed consumption will be $\{c_0^{(0)}(W), c_1^{(1)}(W), c_2^{(1)}(W)\}$. That is, the consumer will follow the $t=0$ plan in $t=0$ and the $t=1$ plan in $t=1$ and $t=2$. (At $t=2$ there is no need to re-compute the plan because the budget constraint dictates that the only option is to consume remaining wealth, $c_2 = W - c_0 - c_1$). This conclusion, however,

Table 1. Optimal consumption plans in 3-period log-utility model
Plans beginning at t=0 and t=1

	Plan formed at time t = 0	*Plan formed at time t = 1*
c_0	$c_0^{(0)}(W) = \frac{1}{1+\rho_1+\rho_2} \cdot W_0$	--
c_1	$c_1^{(0)}(W) = \frac{\rho_1}{\rho_1+\rho_2} \cdot \left[W_0 - c_0^{(0)}(W_0)\right]$	$c_1^{(1)}(W) = \frac{1}{1+\rho_1} \cdot \left[W_0 - c_0^{(0)}(W_0)\right]$
c_2	$c_2^{(0)}(W) = \frac{\rho_2}{\rho_1+\rho_2} \cdot \left[W_0 - c_0^{(0)}(W_0)\right]$	$c_2^{(1)}(W) = \frac{\rho_1}{1+\rho_1} \cdot \left[W_0 - c_0^{(0)}(W_0)\right]$

assumes that the person in period *0* does nothing at all in period *0* to constrain her behavior in period *1*. If she is somehow able to do so in period 0, however, she may choose to pre-commit to the period 0 plan and thereby be somehow "locked in" to $\{c_0^{(0)}(W), c_1^{(0)}(W), c_2^{(0)}(W)\}$. Thus whether the model predicts observed period *1* consumption of $c_1^{(0)}(W)$ or $c_1^{(1)}(W)$ depends on the ability of the individual in period *0* to place effective constraints on her future choices. In fact, period *1* consumption may turn out to be equal to *neither* $c_1^{(0)}(W)$ or $c_1^{(1)}(W)$ if the pre-commitment technology in period *0* is effective but imperfect. The simple DU model no longer has unambiguous observable implications.

In recent years a number of economists have developed models with time inconsistent preferences. Most often these models assume quasi-hyperbolic discounting, which is analytically the most tractable approach. Two notable examples in this literature are Angeletos *et al.* (2001) and Laibson (1997). Both of these papers view the existence of illiquid assets such (e.g., real estate) as imperfect but nevertheless effective self-control vehicles. If an individual places a substantial fraction of her wealth in illiquid form, she will find it costly to increase her consumption in future periods by drawing down wealth. Models with hyperbolic discounting provide explanations for well-established puzzles such as the tendency of consumption to move with income more closely than the relationship implied by the permanent income hypothesis.

More generally, economists and psychologists have in recent years explored in detail the ability of agents to pre-commit to future actions. I now proceed to a selective discussion of a few important papers in this literature.

III. The Recent Time Inconsistency Literature

The important contribution by Strotz (1955–56) did not have a strong immediate impact on the economics literature. The notion of time inconsistency began to receive much more attention with the publication of the paper by Kydland and Prescott (1977), who examined time inconsistency issues in economic policymaking. In their model, time inconsistency occurs in games in which one player is the policymaker and the other is a group of agents affected by policy. The policymaker forms an optimal plan at time *0* that may not be optimal in some future date. For example, a central bank may announce a low-inflation policy only to find later that, if agents have made decisions (involving wages or interest rates on loans to the Treasury) based on expectations of low inflation, the subsequent benefits to the central bank of surprise inflation may exceed the costs. The game between the policymaker and the private agents is essentially a Prisoner's Dilemma, and the equilibrium of the game is suboptimal relative to the outcome that can occur if the policymaker can pre-commit its future actions. Kydland and Prescott viewed this result as a rationale for rules constraining the behavior of policymakers.

The Kydland and Prescott paper spawned a large literature on policy games, and the original focus of Strotz on an individual decision-maker was largely ignored until the 1990s. In the last 10–15 years, economists have returned to the self-control issues addressed by Strotz. The focus of this more recent literature is on issues such as the empirical relevance of exponential discounting and the ability of agents to constrain their own future actions. A very useful survey of this literature is provided by Tirole (2002) in his 2001 Presidential Address to the European Economic Association. Perhaps the most important feature of this literature is a consistent attempt to incorporate relevant insights from the field of psychology. Psychologists, unsurprisingly, have been interested in self-control issues for quite some time, and Rachlin (2000) provides a remarkably interesting and helpful overview of this literature.

There is a large body of experimental evidence on intertemporal choice behavior of humans and animals (rats and pigeons, mainly). One consistent finding in this literature is that discount factors do not decay over time in ways fully consistent with exponential discounting (Rachlin 2000, p. 43). Specifically, observed discount factors are in fact consistent with hyperbolic discounting, which means that as j increases from $j = 0$ in equation (1) ρ_j falls *more* rapidly than exponential discounting for low values of j (in particular, as j goes from 0 to 1) and *less* rapidly than exponential discounting for higher values of j.

Psychologists and economists view this evidence as consistent with the everyday impression that people have self-control problems. A proposed course of action today that seemed like it would be optimal when planning yesterday often no longer seems optimal today. My optimal consumption plan might involve eating pizza for lunch today and salad every day in the future, but when tomorrow comes the salad is sometimes no longer the optimal choice. More generally, unhealthy eating involves a "smaller/sooner" reward (great taste today), while healthy eating generates a larger/later reward (better health and longer life).

Suppose the consumer must choose between a smaller/sooner alterative (SS) that may be consumed at date n and a larger/later alternative (LL) that may be consumed at date $n+1$. Exponential discounting implies that the consumer's preference between SS and LL is independent of n. For example, consider the following two scenarios:

Scenario A: SS = one apple in period 0
LL = two apples in period 1

Scenario B: SS = one apple in period 9
LL = two apples in period 10

$n=0$ in scenario A, while in $n=9$ in scenario B. Exponential discounting implies that if the consumer prefers LL in scenario B, she will also prefer LL in scenario A. Rachlin (2000) notes, however, that a large body of evidence indicates this consistency does not hold. Instead, people (as well as rats and pigeons) are more likely to prefer SS to LL the closer the proximity of the choice to the present.

Because an individual at time n knows that she may not make the "right" choice when period $n+k$ rolls around, she may look for ways to control her subsequent behavior in period $n+k$. Economists view this as the prototypical self-control problem. The following section overviews some strategies of self-control identified in the literature.

IV. Strategies of Self-Control

Individuals attempt to control future choices in a number of ways. One obvious way involves making commitments today so as to penalize short-sighted choices in the future. One can check into a "fat farm" to make it more costly to obtain unhealthy food, for example, or one may sell one's television to avoid its pernicious influence.

It is true that in modern society one cannot voluntarily give up one's rights as a free individual, so it is very difficult to arrange an unbreakable commitment. One can always leave the fat farm or buy another television. Even so, to the extent that the commitments associate a cost with making the undesirable future choice, they improve the chances that behavior will be time consistent.

Individuals in the real world do make use of opportunities to pre-commit. Perhaps the most widely used of these opportunities are illiquid assets such as 401-K retirement accounts. Withdrawals from these accounts before retirement involve significant costs, so placing funds in them is a form of pre-commitment against excessive spending prior to retirement. Laibson, Repetto, and Tobacman (1998) develop a simulation model of the savings behavior of consumers with time inconsistent preferences and show that, under certain circumstances, the availability of pre-commitment mechanisms such as 401-K retirement plans can result in significantly higher levels of aggregate saving.

Apart from explicit strategies such as fat farms, there are a number of other more subtle strategies people use in attempts at self-control. I will now review very briefly a few of these strategies that may be relevant for the intersection of religious faith and economic analysis.

Rachlin (2000) documents extensively that effective solutions to self-control problems often involve a focus by the individual on *patterns* of behavior (habits) rather than specific acts. Hoch and Loewenstein (1991) provide an example of this strategy, which they call "bundling":

> [r]ather than myopically view the eating of an ice cream cone as an isolated act, a dieter may attempt to reframe it as the first in an endless stream of self-control violations. By bundling eating of the current cone with eating of future cones, the consumer may view the extra cost of the cone not as 250 calories, but as an extra 250 calories a day for the foreseeable future, with obesity as the inevitable outcome (pp. 502–503).[3]

Successful self-control strategies seem to involve an individual keeping her- or himself *focused* on desirable patterns. One effective way of doing so in many contexts is for the individual simply to *monitor* his or her behavior.

People who monitor closely their spending tend to save more and spend less. Smokers—even those with no professed desire to quit—smoke less when they keep track carefully how many cigarettes they smoke. "Counting calories" is a time-honored approach to weight control. When confronted with an impulse to deviate from some plan that is optimal in the long-run, the individual is more likely to be aware of the long-run consequences of impatience if he or she sees the choice as breaking a desirable habit.

Rachlin (2000, pp. 125–127) also describes a related but more subtle self-control strategy that has been proven successful in experimental studies. He calls this strategy "reduction of variability." The idea is that the individual makes a choice with the internal commitment to abide by the choice consistently for a period of time. A smoker can choose to smoke any number of cigarettes per day, for example, but he or she must smoke *exactly the same number* each day for a week. Experimental evidence shows that with this commitment, the smoker (even with no professed desire to quit) will smoke less. Rachlin notes that, with this approach, "[t]he pleasures of consumption are no longer restricted to the moment; they are extended in time and therefore more easily compared with future disadvantages" (p. 127).

Tice and Ciarocco (1998) find that "exerting self-control in one area (by inhibiting the response to engage in impulsive behavior) can undermine efforts at a subsequent self-control task even in an unrelated area" (p. 228). That is, self-control operates like an internal resource that becomes depleted with use. Moreover, Muraven, Baumeister, and Tice (1999) find that the capacity for self-control can improve with regular exercise. Thus self-control capacity is rather like muscle strength—it depletes with use (irrespective of the specific nature of the use), recovers with non-use, and improves with exercise.

This is only a very small sample of the enormous literature on self-control. It does provide a sense of some of the remarkably robust characterizations of human behavior identified in the literature. The integration of these psychological insights into economic models has become a "hot" topic in economic research in recent years. Evidence of the importance of this line of research is the fact that two scholars doing major research in the area have recently received the most prestigious recognitions offered by the profession: Daniel Kahneman won the 2002 Nobel Prize in Economic Science, while Matthew Rabin won the 2001 John Bates Clark Medal.

V. Self-Control and Religious Faith

Self-control researchers acknowledge the usefulness of religion in the establishment and maintenance of self-control. Rachlin (2000, p. 13) notes:

> The great advantage of the religious point of view is that it offers a way to achieve self-control. Buddhism, for instance, suggests certain mental and physical exercises; Judaism and Christianity suggest study of sacred texts. All suggest prayer. The end result of self-control from the religious point of view is a body under the control of the best part of the soul. What that is exactly, and how it may be achieved, depends on the religion. While the particular advice that religions prescribe may not be accepted by every person in every modern culture, religions at least offer practical access to self-control. Oddly enough, modern cognitive and physiological psychology, with all its scientific regalia, scarcely attempts to find practical methods of self-control.

When considering the relationship between religious faith and the economics of self-control, two perspectives are possible.

The first perspective asks how economic analysis can be used to understand religious behavior. This is the "economics of religion" approach surveyed by Iannaccone (1998). From this viewpoint one might consider different religions as (among other things) different self-control technologies. One might then expect individuals to sort into different religion groups according to how well the groups meet their self-control needs. One might also use the presence or absence of religiosity as an indicator of the extent of self-control. Barro (1999, p. 1137) makes the following comments on religion in his model of economic growth with hyperbolic discounting:

> Commitment can also be provided by personal discipline and self-control. The extent of this discipline may differ for cultural reasons across societies, but such differences are hard to quantify. Religious principles are dedicated, in part, toward curbing lavish expenditures and excessive debt. Thus, one potentially observable influence on commitment is religiosity, measured by church attendance or outlays on organized religion. Differences in affiliations also matter because attitudes toward material pleasures differ across religions.

The second perspective on the relationship between religious faith and the economics of self-control is to ask whether religion in general (and Christianity in particular) has insights about behavior that might be usefully incorporated into economic models of self-control. In contrast to the first perspective, which is generally straightforward and which is gaining acceptance in the mainstream economics literature, the second perspective is unknown territory.

If religion is ever going to make a useful contribution to economic methodology, this would certainly seem to be a likely place. Self-control might, after all, be reasonably viewed as the ability to resist temptation—something about which religion surely has as much or more to say than anything else economists study.[4] In the following section I propose one possible Christian perspective on the self-control strategies discussed above, and Section 8 presents a simple model that illustrates those ideas.

VI. Remembrance

In his letter to the Galatians, the Apostle Paul notes:

> But the fruit of the Spirit is love, joy, peace, patience, kindness, goodness, faithfulness, gentleness, self-control; against such there is no law. And those who belong to Christ Jesus have crucified the flesh with its passions and desires (Gal 5:22–24, RSV).

Thus Scripture plainly states that self-control (*enkrateia*) is a fruit of the Spirit. I propose that the idea of *remembrance* may be a useful way to think about the self-control benefits of religion. By "remembrance" I mean the subset of an individual's knowledge that is within his consciousness when he makes a decision. Assume the following:

(i) If an individual consciously "remembers" her long-term priorities, goals, and commitments at the exact instant she faces a decision, she will more likely to exercise self-control and avoid temptation; and

(ii) Remembrance is difficult to maintain—in economics language, it depreciates rapidly.

These two assumptions imply that ongoing strategies to preserve remembrance will lead to improved self-control.[5] I believe this is a reasonable way to think about the effectiveness of monitoring and bundling as self-control mechanisms. An important function of religious ritual is to help believers remember what God has done and their commitment to God. At communion, the words of institution quote Christ: "This is my body which is given for you. Do this in remembrance of me" (Luke 22:19). The idea of remembering, or not forgetting, shows up many places in Scripture. Here are two examples:

> These things I have spoken to you, while I am still with you. But the Counselor, the Holy Spirit, whom the Father will send in my name, he will teach you all things, and *bring to your remembrance* all that I have said to you (John 14:25–26, RSV, emphasis added).

> Only take heed, and keep your soul diligently, *lest you forget* the things which your eyes have seen, and lest they depart from your heart all the days of your life; make them known to your children and your children's children (Deuteronomy 4:9, RSV, emphasis added)

A decision-maker generally has a large amount of knowledge, including knowledge of past experiences, commitments, and perspectives. At any given time, however, most of this knowledge is not within her frame of remembrance —that is, she is not explicitly conscious of it.

This idea is certainly consistent with many aspects of experience. People tend to lose focus. Political consultants are valuable to the extent they can keep a candidate "on message." Individuals experiencing tragedy such as the death of a family member or a struggle with illness often find that they have lost sight of what is really important. Bringing those ideas into consciousness will, change behavior. When I really remember how much I love my wife and son, I take better care of myself by exercising more and making healthier food choices. Christians have several practices to maintain or restore focus. These practices include worship, prayer, fasting, and fellowship with like-minded believers.

How might these ideas be built into an economic model? The following section specifies and solves a simple model that incorporates some of these ideas.

VII. A Simple Model of Self-Control

Table 1 outlines the consumption plans formed at times *t=0* and *t=1* in the very simple three period model developed in Section II. The formulas show that the period 0 plan is time consistent [that is, $c_1^{(0)}(W) = c_1^{(1)}(W)$ and $c_2^{(1)}(W) = c_2^{(1)}(W)$] only in the case of exponential discounting. If discounting is not exponential, a self-control problem arises in period *1* in the sense that the optimal period *1* and *2* consumption levels, viewed from the perspective of period *1*, differ from the optimal period 1 consumption levels planned in period *0* (that is, $c_1^{(0)}(W) \neq c_1^{(1)}(W)$ and $c_2^{(0)}(W) \neq c_2^{(1)}(W)$). This section describes an extended version of the Section II model that considers the self-control problem explicitly.

Assume that discounting is quasi-hyperbolic rather than exponential, so that $\rho_1 = \lambda \cdot \rho$ and $\rho_2 = \lambda \cdot \rho^2$, where $0 < \lambda < 1$ and $0 < \rho < 1$. Suppose at time *t=0* the consumer understands that her preferences are potentially time inconsistent. This idea may be represented by allowing the consumer to contemplate two possible outcomes. The first is the *time consistent* (TC) outcome, which is simply the optimal plan formed at time *t=0*: $\{c_0^{(0)}(W), c_1^{(0)}(W), c_2^{(0)}(W)\}$. The second is the *time inconsistent* (TI) outcome, which involves following the time *t=0* plan in period *0* and the time *t=1* plan in periods *t=1* and *t=2*: $\{c_0^{(0)}(W), c_1^{(1)}(W), c_2^{(1)}(W)\}$. The time consistent plan involves delaying gratification from period *1* to period *2*, so that $c_1^{(0)}(W) < c_1^{(1)}(W)$ and $c_2^{(0)}(W) > c_2^{(1)}(W)$. Period *1* savings is therefore higher with the time consistent plan than with the time inconsistent plan.

I allow for only these two outcomes (TC and TI) to facilitate analytical tractability. In effect, I allow only for the extreme cases of perfect self-control (TC) and no self control at all (TI). Intermediate outcomes are also possible in which period *1* consumption is less than the TC outcome and more than the TI outcome, but those are beyond the scope of the present discussion.

The individual at time $t=0$ considers her consumption plan for periods 0, 1, and 2. The TC plan is clearly optimal, but in the absence of perfect self-control there is some chance she will deviate from it and follow the TI plan instead. The discussion in Sections IV–VI above makes clear that there are courses of action the individual can take to increase the chance that the TC plan will be followed. In general, however, it is not possible to *insure* that the TC plan will be followed.

Let p denote the probability in period 0 that the TC plan will be followed, so (1-p) is the TI probability. Suppose p is determined as follows:

$$(7) \quad p = \alpha \cdot A + \beta,$$

where α = a parameter representing the relative importance of internal factors for the ability of the individual to exercise self-control ($0 \leq \alpha \leq 1$),
A = a variable representing the scale of actions taken by the individual to enhance self-control (which might include efforts at remembrance),
β = a variable representing the influence of factors external to the individual on the individual's ability to exercise self-control, with ($0 \leq \beta \leq 1 - \alpha$).

For simplicity, assume A is a binary choice variable taking on the value $A=1$ if the individual chooses to "invest" in self-control in period 0 and $A=0$ otherwise. Equation (7) implies that $p=(\alpha+\beta)$ when $A=1$ and $p=\beta$ when $A=0$. A more general model would allow A to take on any value between 0 and 1.

For the problem to be non-trivial, the investment in self-control by the individual must involve some cost. In general, this cost can be explicit or implicit. An example of an explicit cost is the monetary cost of participation in a weight-control program. An example of an implicit cost is the opportunity cost of time spent in efforts to retain remembrance of one's basic values and priorities (as described above in Section VI). This model makes the simple assumption that a choice of $A=1$ imposes a monetary cost $S>0$ on the individual.

Thus at time $t=0$, the individual faces not only a choice about how much to consume in that period, but also a choice of whether to invest in self-control for the future. Let W_0 denote the baseline wealth of the individual in period 0. If she chooses $A=0$ and therefore does not invest in self-control, she will have total wealth of W_0 to allocate across the three periods. If she decides to invest in self-control and chooses $A=1$, the cost of the self-control investment (S) means her total wealth available for consumption will fall to (W_0 - S).

To facilitate tractability, assume the self-control investment decision in period 0 is separate from and prior to the consumption decision. In general, a choice of $A=1$ only increases the chance of the time consistent outcome and does not guarantee it. In terms of equation (7), $A=1$ only increases p relative to the case when $A=0$. As long as $\alpha+\beta<1$, $A=1$ does not imply that $p=1$.

Thus the individual faces four possible consumption paths: (i) If she chooses $A=1$ and self-control turns out to be possible, she follows the TC path while allocating wealth of (W_0 - S) to consumption; (ii) If $A=1$ and self-control is not possible, she follows the TI path while allocating wealth of (W_0 - S) to consumption; (iii) If she chooses not to invest in self-control ($A=0$) and self-control is nevertheless possible, she follows the TC path while allocating wealth of W_0 to consumption; (iv) If $A=0$ and self-control is not possible, she follows the TI path and allocates wealth of W_0 to consumption.

Thus wealth can take on two values: $W=W_0$ when $A=0$, and $W=W_0 - S$ when $A=1$. Recall the utility maximization problem shown in equation (4) above. Let $V_{TC}(W)$ denote the indirect utility in period 0 associated with initial wealth W and the *time consistent* consumption path:

$$(8) \quad V_{TC}(W) = \ln[c_0^{(0)}(W)] + \lambda \cdot \rho \cdot \ln[c_1^{(0)}(W)] + \lambda \cdot \rho^2 \cdot \ln[c_2^{(0)}(W)].$$

Note that the level of wealth is now shown explicitly in the optimal consumption values. Similarly, let $V_{TI}(W)$ denote the indirect utility in period 0 associated with wealth W and the *time inconsistent* consumption path:

$$(9) \quad V_{TI}(W) = \ln[c_0^{(0)}(W)] + \lambda \cdot \rho \cdot \ln\{c_1^{(1)}[W-c_0^{(0)}(W)]\} + \lambda \cdot \rho^2 \cdot \ln\{c_2^{(1)}[W-c_0^{(0)}(W)]\}.$$

Define the term "expected meta-utility," denoted E(M), as the following weighted average:

$$(10) \quad E(M) = p \cdot V_{TC}(W) + (1-p) \cdot V_{TI}(W),$$

where E(·) is the expected value operator. Assume the individual chooses the value of A (either $A=0$ or $A=1$) that maximizes $E(M)$. This decision determines whether the individual chooses to invest in self-control.

When $A=1$, initial wealth is (W_0-S) and equation (7) implies that $p=\alpha+\beta$ and expected meta-utility is given by:

$$(11) \quad E(M \mid A=1) = (\alpha+\beta) \cdot V_{TC}(W_0 - S) + (1-\alpha-\beta) \cdot V_{TI}(W_0 - S)$$

When $A=0$, initial wealth is W_0 and equation (7) implies that $p=\beta$ and the expected value of meta-utility is given by:

(12) $E(M \mid A=0) = \beta \cdot V_{TC}(W_0) + (1-\beta) \cdot V_{TI}(W_0)$.

The individual chooses to invest in self-control (*A=1*) when $E(M \mid A=1) > E(M \mid A=0)$. Conversely, the individual chooses *A=0* when $E(M \mid A=0) > E(M \mid A=1)$.

To facilitate comparative static analysis, define Δ as follows:

(13) $\Delta = E(M \mid A=1) - E(M \mid A=0)$.

A change in any parameter or exogenous variable that increases Δ makes it more likely the individual will invest in self-control. Straightforward but rather tedious algebra implies (i) $\frac{\partial \Delta}{\partial \alpha} > 0$, (ii) $\frac{\partial \Delta}{\partial \beta} = 0$, (iii) $\frac{\partial \Delta}{\partial W} > 0$, and (iv) $\frac{\partial \Delta}{\partial S} < 0$. (i) says that as self-control investments become more effective, the individual is more likely to invest in self-control. (ii) implies that the tendency of the individual to invest in self-control is *independent* of the extent to which self-control efforts are supported by the external world. (iii) means that wealthier individuals are more likely to invest in self-control – i.e., self-control investment is a normal good. (iv) indicates that a rise in cost reduces the tendency of the individual to invest in self-control

The partial derivatives $\frac{\partial \Delta}{\partial \lambda}$ and $\frac{\partial \Delta}{\partial \rho}$ are quite complex, and as a result sign determination is difficult. Simulations, however, suggest strongly that (v) $\frac{\partial \Delta}{\partial \lambda} < 0$ and (vi) $\frac{\partial \Delta}{\partial \rho} > 0$. (v) means that as the individual becomes more impatient in the near term (i.e., as λ falls), investment in self-control is more likely to be undertaken. (vi) implies that a higher value of the discount parameter ρ makes self-control investment less likely. This is because a higher value of ρ decreases relative impatience in the present as compared to the future.[6]

If self-control actions take the form of explicit purchases (e.g., weight-loss programs), the above comparative-static results have implications for the variables one might include in the associated demand function. The discussion in Section 7 above, however, focused on a kind of self-control investment (remembrance of core values) that in all likelihood will be difficult to observe directly. This fact makes it necessary to ask whether the model has any unusual implications for *observable* data.

One possible observable implication of the model involves a relationship between the cross-section mean and variance of consumption. Consider c_1, consumption in period *1*. We know from above that $c_1 = c_1^{(0)}(W)$ with probability *p* and that $c_1 = c_1^{(1)}(W)$ with probability *(1-p)*. For a given level of initial wealth, the implied mean and variance of c_1 are:

(14) $E(c_1)=[x + p\cdot(1-x)]\cdot c_1^{(0)}(W)$

(15) $Var(c_1) = p\cdot(1-p)\cdot(1-x)^2\cdot[c_1^0(W)]^2$,

where $x = \frac{1+\rho}{1+\lambda\cdot\rho}$. Equation (14) implies that, for a given value of wealth, the mean of period 1 consumption is strictly increasing in *p*. Equation (15), however, implies a non-linear relationship between *p* and the variance of period *1* consumption.[7] This variance is proportional to the product *p·(1-p)*, which is the familiar binomial variance expression. $Var(c_1) = 0$ if either *p=1* or *p=0*. When *p=1*, the time consistent value of c_1 is always chosen, and the time inconsistent value of c_1 is always chosen when *p=0*. As *p* begins to rise from 0 and move toward 1, $Var(c_1)$ at first rises, reaches a maximum at *p=0.5*, and then falls.

Recall that $p = \alpha+\beta$ when *A=1* and $p=\beta$ when *A=0*. We know that individuals with higher values of α are, other things equal, more likely to choose *A=1*. In a cross-section of individuals, each member of the cross-section would belong in one of three categories: (I) individuals with $\alpha <$ 0.4; (II) individuals with values $0.4 \leq \alpha \leq 0.6$; or (III) individuals with $\alpha > 0.6$. Equation (14) implies that, controlling for income and other determinants of consumption not represented in this simple model, the mean values of consumption will be highest in group III, second highest in group II, and lowest in group I. Equation (15) implies that, again controlling for income and other factors, the cross-section variance of consumption will be higher in group II than in either group I or group III, with the relative variance of group III vis-à-vis group I is indeterminate.

Suppose, plausibly, that wealth and income are highly correlated (positively) across individuals. In that case, higher-than-average consumption implies a lower-than-average savings rate. The key observable implication of the present model is this: *Controlling for income and other determinants of consumption, there will be higher variability in savings rates among "mid-range savers" than among* either *"high savers" or "low savers."*

Section VI describes possible self-control strategies involving remembrance that do not necessary involve specific observable market actions (such as the purchases of weight loss programs). Even so, however, religious participation may be positively correlated with α, the efficiency of self-control investment. This correlation could arise because religious activities such as prayer and worship do in fact bring remembrance of an individual's core values. Suppose this is the case. Then those individuals who have almost no religious participation would have relatively low values of α, and those who are at church every time the door opens would have relatively high values of α. Those who attend only occasionally would have values of α between those of the two other groups. In that case (and as long as "high" values of α are above 0.5 and "low" values of α are below 0.5), the above model predicts that variability in consumption across those who attend only occasionally may be higher than for either those who do not attend or for those who attend heavily.

VIII. Does the Study of Economics Inhibit Self-Control?

In a widely-discussed paper, Frank, Gilovich, and Regan (1993) present evidence based on experiments using the ultimatum game that the study of economics inhibits cooperation. Although their arguments have been challenged[8], the impact of their paper on the profession suggests that we ask the same question about self-control.

Introductory microeconomics courses generally rely heavily on the rational choice paradigm, focusing on the idea that consumer choices reflect utility maximization. The conventional discussion of choice assumes a consumer is choosing over two or more goods—often abstract goods such as x_1 and x_2. Not much thought is given to the nature of x_1 and x_2 (though sometimes "cute" examples are used such as x_1 = pizza and x_2 = ice cream). The self-control literature discussed above makes it clear, however, that *defining* x_1 *and* x_2 *properly is a key to self-control*. In particular, the goods over which the consumer is maximizing should be "bundled" so that a seemingly small indiscretion today can be seen for what it is—an action that could be very costly because it involves breaking a healthy pattern.

Consider a student who has just earned a grade of "A" in introductory microeconomics. The student, let us suppose, has an emerging alcohol abuse problem. His friends express some concerns to him one day, and in response to their concerns he resolves to change his behavior. That very night, however, he goes with some other friends to party at which alcohol is readily available. As a good economics student, he may determine that, at the margin, the benefits of "one more evening" of alcohol consumption outweigh the costs. He will be drawn to that conclusion if he has been encouraged, implicitly if not explicitly, to think of this choices as involving *single* actions rather than *patterns* of actions. One could make similar arguments regarding any number of other self-destructive behaviors.

As another example, consider an individual who wants to lose weight and has decided join a health club. She is confronted with two choices: (i) pay $50 per month, or (ii) pay $600 in advance for 12 months. Conventional economic analysis would point to option (i) for at least two reasons. First, if interest rates are positive, the present value of (i) is less than the present value of (ii). Second, there is at least some chance that the club will go out of business before 12 months have expired. Conventional discussions may tend to ignore, or at least not emphasize sufficiently, an important potential benefit of option (ii): it can be a signal that reminds her future selves of the values and priorities of her current self.

IX. Conclusion

This paper has considered the time inconsistency question and related issues of self-control. Religion facilitates self-control and should therefore in principle provide insights that may be of use in the development of economic models of self-control phenomena. If religion is ever going to have anything useful to say to economics (in terms of methodology), surely it can do so when the issue involves resistance to temptation.

After a review of the self-control literature, the paper developed a Christian perspective on self-control based on the Biblical idea of remembrance. That perspective, when embedded in a simple model, yields a somewhat novel implication of a non-linear relationship between the cross-section mean and cross-section variance of consumption. It also has another implication about the possible effect of religious participation on the variability of consumption. The paper has therefore provided one specific example of how the Christian understanding of human behavior can inform and guide the specification of an economic model.

Endnotes

1 The model in this section is a slight generalization of the simple example presented by Hall (1998).

2 If $c_1^{(0)}(W) = c_1^{(1)}(W)$, the budget constraint implies that $c_2^{(0)}(W) = c_2^{(1)}(W)$. Any plan that is consistent for period *1* will necessarily also be consistent for period *2*.

3 Hoch and Loewenstein note that advertisers therefore try to "unbundle" costs with tactics such as "only pennies a day."

4 For example, the title of a recent paper published in *Econometrica* by Gul and Pesendorfer (2001) is "Temptation and Self-Control."

5 These assumptions relate to the debate over *akrasia* — a debate which goes back (at least) to Aristotle. Akrasia involves ". . . free (and therefore uncompelled) intentional actions that are contrary to the agent's consciously held better judgment at the time of action" (Mele 2002, p. 153). Socrates believed akratic actions were impossible: why would someone who knew an action was not best nevertheless take it? Aristotle, in contrast, argued that akratic actions do occur as the result of the passions. (*Oxford Companion to Philosophy*, 2002) I do not presume to make a contribution to this debate. My purpose instead is to show how a particular Bible-based perspective on self-control might be incorporated into an economic model.

6 The relative discount applied to period *1* consumption compared to period *0* consumption (what we might call "near-term impatience") is $\lambda\cdot\rho/1$. The relative discount applied to period *2* consumption compared to period *1* consumption ("future impatience") is $\lambda\cdot\rho^2/\lambda\cdot\rho = \rho/1$. A rise in ρ increases $\rho/1$ more than it increases $\lambda\cdot\rho/1$, so near-term impatience falls relative to future impatience, which means less of a self-control problem.

7 Similar expressions hold for period *2* consumption, c_2.

8 See Yezer, Goldfarb, and Poppen (1996) for a discussion of both sides of this issue.

References

Angeletos, George-Marios, David Laibson, Andrea Repett, Jeremy Tobacman, and Stephen Weinberg. 2001. "The Hyperbolic Consumption Model: Calibration, Simulation, and Empirical Evaluation." *Journal of Economic Perspectives.* 15:3 (Summer), pp. 47–68.

Barro, Robert J. 1999. "Ramsey Meets Laibson in the Neoclassical Growth Model." *Quarterly Journal of Economics.* CXIV:4 (November), pp. 1125–1152.

Frank, Robert H., Thomas Gilovich, and Dennis T. Regan. 1993. "Does Studying Economics Inhibit Cooperation?" *Journal of Economic Perspectives.* 7:2 (Spring), pp. 159–171.

Gul, Faruk, and Wolfgang Pesendorfer. 2001. "Temptation and Self-Control." *Econometrica.* 69:6 (November), pp. 1403–1435.

Hall, Robert E. 1998. "Comments and Discussion." *Brookings Papers on Economic Activity.* 98:1, pp. 173–177.

Hoch, Stephen J., and George F. Loewenstein. 1991. "Time-Inconsistent Preferences and Consumer Self-Control." *Journal of Consumer Research.* 17 (March), pp. 492–507.

Iannaccone, Laurence R. 1998. "Introduction to the Economics of Religion." *Journal of Economic Literature.* 36:3 (September), pp. 1465–1495.

Laibson, David. 1997. "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics.* CXII:2 (May), pp. 443–477.

Laibson, David I., Andrea Repetto, and Jeremy Tobacman. 1998. "Self-Control and Saving for Retirement." *Brookings Papers on Economic Activity.* 1, pp. 91–172.

Kydland, Finn E., and Edward C. Prescott. 1977. "Rules Rather than Discretion: The Inconsistency of Optimal Plans." *Journal of Political Economy.* 85:3 (June), pp. 473–492.

Mele, Alfred R. 2002. "Akratics and Addicts." *American Philosophical Quarterly.* 39:2 (April), pp. 153–67.

Muraven, Mark, Roy F. Baumeister, and Dianne M. Tice. 1999. "Longitudinal Improvement of Self-Regulation through Practice: Building Self-Control Through Repeated Exercise." *Journal of Social Psychology.* 139 (August), pp. 446–457.

Oxford Companion to Philosophy. 2002. "Akrasia." http:/www.xrefer.comentry.jsp?xrefid=551239&secid=.-&hh=1 (accessed 11/7/02).

Rachlin, Howard. 2000. *The Science of Self-Control.* Cambridge, MA: Harvard University Press.

Samuelson, Paul. 1937. "A Note on Measurement of Utility." *Review of Economic Studies.* 4, pp. 155–161.

Strotz, R. H. 1955–1956. "Myopia and Inconsistency in Dynamic Utility Maximization." *Review of Economic Studies.* 23:3, pp. 165–180.

Tice, Dianne M. and Nichole J. Ciarocco. 1998. "Inhibition and Self-Control." *Psychological Inquiry.* 9:3, pp. 228–230.

Tirole, Jean. 2002. "Rational Irrationality: Some Economics of Self-Management." *European Economic Review.* 46, pp. 633–655.

Yezer, Anthony M., Robert S. Goldfarb, and Paul J. Poppen. 1996. "Does Studying Economics Inhibit Cooperation? Watch What We Do, Not What We Say or How We Play?" *Journal of Economic Perspectives.* 10:1 (Winter), pp. 177–186. ■