Available online at www.sciencedirect.com
SCIENCE DIRECT®
Review of Economic Dynamics 7 (2004) 243–264
Review of Economic Dynamics
www.elsevier.com/locate/red

Self-control, revealed preference and consumption choice

Faruk Gul and Wolfgang Pesendorfer *

Princeton University, Princeton, NJ, USA

Received 9 December 2002

Abstract

We provide a time-consistent model that addresses the preference reversals that motivate the time-inconsistency literature. The model subsumes the behavior generated by the time-inconsistency approach in finite settings but, unlike the time-inconsistent models, allows for self-control. This paper provides a brief summary of theoretical results shown elsewhere [Gul and Pesendorfer, Econometrica 69 (2001) 1403; Econometrica, (2002), in press; Rev. Econ. Stud. (2002), in press] and contrasts the predictions and welfare implications of our model and the time-inconsistent $\beta-\delta$ model.

1. Introduction

Experiments find evidence that individuals resolve the same intertemporal trade-off differently depending on when the decision is made. (See Loewenstein, 1996 for a recent survey.) In the typical experiment, subjects choose between a smaller, date 1 reward and a larger, date 2 reward. If the choice is made at date 1 then the smaller–earlier reward is chosen. If the choice is made earlier (at date 0) then the larger–later reward is chosen.

We interpret the behavior documented in experiments on time preference as part of a broader phenomenon of struggling with *temptations*. In the decision problem above, subjects find *immediate* rewards tempting. When the decision is made in period 1, the smaller–earlier reward can be consumed immediately and hence constitutes a temptation. As a result, the agent is more inclined to choose the smaller–earlier reward. When the

* Corresponding author.
E-mail address: pesendor@princeton.edu (W. Pesendorfer).

1094-2025/$ – see front matter
doi:10.1016/j.red.2003.11.002

decision is made in period 0, neither of the two rewards can be consumed immediately and hence the decision is unaffected by temptations.

How can we decide which alternatives are temptations? Consider again the above decision problem and suppose that the agent must choose between the following two alternatives in period 0. The first alternative offers commitment to the larger–later reward. The second alternative offers the option of choosing either reward in period 1. If the agent expresses a strict preference for the first alternative over the second, we say that the smaller–earlier reward is a temptation. In that case, the agent strictly prefers to eliminate the smaller–earlier reward from his period-1 options. Hence, an alternative is identified as a temptation if its availability makes the agent worse-off.

When the agent cannot exclude a temptation from the period-1 choice set he will either succumb or exercise costly self-control. Self-control describes a situation where the agent does not choose the most tempting alternative. In the example above, the agent would exercise self-control if he strictly prefers a situation where the smaller–earlier reward is *not* available but still chooses the larger–later reward when the smaller–earlier reward *is* available. We interpret this combination of period-0 and period-1 choices as situation where the individual exercises costly self-control in period 1. If the smaller–earlier reward were not available the individual would be better off because he would not incur the cost of self-control.

We discuss the conceptual ideas in more detail in Section 2. There we also provide a representation of preferences for a two-period decision problem. Section 3 extends the model to finite horizon decision problems.

Our work is related to the time-inconsistency literature. Starting with Strotz (1956), authors have analyzed models of changing (time-)preference. That approach assumes that the agent is an independent decision-maker in every period and resolves conflicts between the various "selves" in a game theoretic manner.[1] In Section 4, we show that for finite decision problems we can re-interpret the time-inconsistency model as a temptation model. Hence, for finite decision problems our model subsumes the behavior of time-inconsistent agents. The converse is not true because time-inconsistent models do not allow for self-control. In Section 5, we discuss applications and illustrate how predictions differ for the time-inconsistent model and the model with self-control.

Section 5.2 considers a competitive economy with liquid and illiquid assets. As Kocherlakota (1996) points out, the time-inconsistency approach predicts that agents will specialize in their asset holdings and either hold only liquid or only illiquid assets. In contrast, our model allows for mixed asset holding and hence comes closer to matching observed behavior. Section 5.3 discusses the conditions under which agents benefit from the introduction of an intergenerational transfer. In a recent paper, Imrohoroglu et al. (2002) point out that time-inconsistent agents will typically not benefit from such a system. We illustrate how a different conclusion may emerge in a model with self-control.

[1] Strotz (1956) assumes that every period the decision maker maximizes utility among all plans that are consistent with incentives of future selves. He refers to this behavior as consistent planning. Peleg and Yaari (1973) note that consistent plans are not well defined unless the decision problem is finite and use Nash equilibrium to resolve the conflict among the various selves.

Section 5.4 considers a model in which gambles that offer immediate returns constitute temptations. Such preferences may help explain why some types of gambles are prohibited or regulated while other risky investments such as stocks are not. The difference between gambling and stock investing lies in the timing of returns. Stock investing offers risky returns with a delay whereas casino gambling offers risky returns with immediate rewards. For agents described in Section 5.4 only the latter is a temptation and hence only the latter offers room for welfare improving regulation.

Section 6 analyzes infinite horizon choice problems. In particular, we focus on a standard consumption–savings model. As is the case for standard time-separable preferences, optimal behavior in our model is described as the solution of a standard dynamic programming problem. This is in contrast to the time-inconsistency approach which must confront a folk-theorem like multiplicity of outcomes in this case.[2] Our model does not exhibit this kind of multiplicity and, in particular, assigns to each decision problem a unique value.

Section 7 contrasts welfare analysis in our model and the time-inconsistent model. In standard economic models the welfare of an agent is synonymous with his choices. That is, if the agent chooses alternative a over alternative b then this means that a leads to higher welfare than b. Our model retains this feature of standard theory. In particular, if our model predicts that agents "should" prefer a particular policy (for example, smokers should prefer an increase in cigarette taxes) then behavior contradicting that prediction (for example, smokers voting against an increase in cigarette taxes) is evidence that the model is incorrect. However, our model differs from standard economic models in that the agent values commitment. Hence, to evaluate welfare, we have to keep track of both the individuals consumption choices and his commitment choices.

In contrast, the time-inconsistency literature assumes that each decision-maker consists of a sequence of distinct agents—called the (multi)selves. Each self has a different preference over consumption streams. Hence the period-t self's choice of alternative a over b reflects only the fact that given the predicted behavior of the subsequent selves, a leads to a consumption stream that is better for the period-t self than the one induced by b. Other selves may be and often are made worse-off by this choice. In the time-inconsistency literature the selves do not value commitment per se; commitment has value only as a vehicle for one of the selves to impose his preferences on subsequent selves.

With multiple selves, finding an adequate welfare criterion is difficult. This is why researchers have used a variety of alternative criteria. For example, O'Donoghue and Rabin (1999) maximize the utility of the period-0 self. There are at least two problems with this welfare criterion. First, since the preferences of the period-0 self can no longer be observed, welfare trade-offs are made based not on observable choice but on the modeler's conjectures regarding what these choices would have been if they could been put to the period-0 self. Hence, an assertion of the form "policy a is better than policy b" can never be refuted. Second, even if the preferences of the period-0 self could be verified, it does not seem reasonable for a social planner to align himself with an agent who is no longer present. Applied to an economy with multiple agents (as opposed to multiple selves)

[2] As Krusell and Smith (2003) show, refinements such as Markov perfection do not alleviate this multiplicity.

such as an overlapping generations model or a dynastic model, this criterion would yield absurd policy recommendations. In the concluding section of our paper, we point out the shortcomings of the other welfare criterion for the multiselves approach.

2. A model of self-control

We begin with a simple two-period model. Consumption takes place in period 1. In period 0, the agent takes an action that affects the set of alternatives available in period 1. We can think of the period-0 problem as a choice among sets of alternatives. The interpretation is that in period 1 the agent must pick an alternative from the set chosen in period 0.

The model takes as given a preference relation (denoted $\succcurlyeq$) over sets of consumption lotteries. This preference captures the agent's period-0 behavior. Our objective is to identify assumptions on period-0 behavior that can be interpreted as coming from an individual that expects to struggle with temptation in the period 1 and may have the ability to exercise self-control. A complete analysis of this model can be found in Gul and Pesendorfer (2001). Here we illustrate the main ideas.

Let $c \in C$ denote consumption in period 1 and let p denote a consumption lottery. We use c to denote also the degenerate lottery that yields c with probability 1. Let B denote a set of consumption lotteries.

Suppose the agent is a standard expected utility maximizer. Then, he has a utility function u such that the period-1 choice solves

$$\max_{p\in B} \int u(c)\,\mathrm{d}p.$$

Therefore, in period 0 this agent prefers the set B to the set B' if

$$\max_{p\in B} \int u(c)\,\mathrm{d}p \geqslant \max_{p\in B'} \int u(c)\,\mathrm{d}p. \tag{1}$$

This standard agent satisfies the following key property (Kreps, 1979):

$$B \succcurlyeq B' \quad \Rightarrow \quad B \sim B \cup B'. \tag{S}$$

That is, if the set B is weakly preferred to the set B' then adding the alternatives in B' to B has no effect on the agent's welfare. Kreps shows that this property characterizes the standard agent. That is, if the preference over sets satisfies (S) then we can find a utility function u such that $B \succcurlyeq B'$ if and only if inequality (1) holds for that u. Assumption (S) rules out the possibility that the agent may benefit from the alternatives in the inferior set B'. [3] Moreover, (S) rules out the possibility that the agent is harmed by the addition of the alternatives in B'. Our model of temptation and self-control relaxes the latter aspect of (S).

The preference $\{c\} \succ \{c, c'\}$ expresses a desire to commit. In our model, the availability of c' is undesirable because it represents a temptation in period 1. For example, c' may

[3] Kreps (1979), Dekel et al. (2001) explore preferences where this aspect of (S) may be violated.

be some unhealthy food or cigarettes or some other consumption good that is commonly associated with a preference for commitment. When the availability of c' makes the agent worse-off, we require that $\{c, c'\} \succcurlyeq \{c'\}$. We interpret $\{c, c'\} \sim \{c'\}$ as a situation where the agent succumbs to the temptation presented by c'. In contrast, $\{c, c'\} \succ \{c'\}$ is situation where temptation lowers the agent's welfare but the availability of c remains beneficial. We interpret this as an instance of self-control. The agent chooses c in period 1 but must bear a disutility of self-control triggered by the presence of c'.

More generally, the key assumption of our model is *set betweenness*:

$$B \succcurlyeq B' \quad \Rightarrow \quad B \succcurlyeq B \cup B' \succcurlyeq B'. \tag{SB}$$

Set betweenness captures the idea that the source of the preference for commitment are temptations. Moreover, there is a ranking of alternatives according to how tempting they are. The agent's well-being is affected only by the most tempting alternative. To see the connection between SB and our temptation interpretation consider the choice problems B, B' with $B \succcurlyeq B'$. Temptations can only lower the agent's utility. Hence, $B \succcurlyeq B \cup B'$. If the most tempting alternative from $B \cup B'$ is in B then the addition of B' to B does not affect the agent's welfare and we have $B \cup B' \succcurlyeq B$ (which implies $B \cup B' \sim B$ since $B \succcurlyeq B \cup B'$). If the most tempting alternative from $B \cup B'$ is in B' then the addition of B to B' cannot lower the agent's utility and we have $B \cup B' \succcurlyeq B'$. In either case, SB is satisfied.

In Gul and Pesendorfer (2001) we show that SB together with the standard axioms that yield expected utility imply (and are implied by) the following representation of the preference $\succcurlyeq$: there are two von Neumann–Morgenstern utility functions u and v such that $B \succcurlyeq B'$ if and only if

$$\begin{aligned} &\max_{p\in B} \int \big(u(c) + v(c)\big)\,\mathrm{d}p - \max_{p\in B} \int v(c)\,\mathrm{d}p \\ &\quad \geqslant \max_{p\in B'} \int \big(u(c) + v(c)\big)\,\mathrm{d}p - \max_{p\in B'} \int v(c)\,\mathrm{d}p. \end{aligned} \tag{2}$$

The function u represents the agent's ranking over alternatives when he is committed to a single choice. To see this, note that when B consists of only one element the v-terms in the above formula drop out.

When the agent is not committed to a single choice then his welfare is affected by the temptation utility represented by v. Consider the two element choice set $B = \{c, c'\}$ and assume that $u(c) > u(c')$. We can distinguish three cases.

(1) When $v(c) \geqslant v(c')$ then the commitment and temptation utilities agree. In this case there is no preference for commitment since $\{c, c'\} \sim \{c\}$. The agent chooses c in period 1.
(2) When $v(c') > v(c)$ and $u(c') + v(c') > u(c) + v(c)$, then c' is a temptation, that is, $\{c\} \succ \{c, c'\}$. Moreover, $\{c, c'\} \sim \{c'\}$ since

$$\max_{\tilde{c}\in\{c,c'\}} \big(u(\tilde{c}) + v(\tilde{c})\big) - \max_{\tilde{c}\in\{c,c'\}} v(\tilde{c}) = u(c').$$

In this case, the agent succumbs to the temptation and chooses c' in period 1.

(3) When $v(c') > v(c)$ and $u(c) + v(c) > u(c') + v(c')$, then as in case (2) the alternative c' is a temptation. However, in this case $\{c\} \succ \{c, c'\} \succ \{c'\}$. This is a case where the agent exercises self-control: he chooses c in period 1 but incurs a utility penalty of $v(c') - v(c) > 0$ which we interpret as the cost of self-control.

Period 1 choices maximize $u + v$. Hence, period-1 behavior maximizes a utility function that is a "compromise" between the temptation and the commitment utilities. Recall that our model makes assumptions only on period-0 behavior. In Gul and Pesendorfer (2001) we also provide an extended model that assumes we observe the agent's behavior in period 0 (the preference $\succcurlyeq$) and his behavior in period 1 (the choice from B). In that model, we give conditions on behavior in both periods such that period-0 preferences are represented by a formula given in (1) and period-1 choices maximize $u + v$.

3. Dynamic models of self-control

Applications of self-control often require a more elaborate dynamic setting. For example, we may want to consider the behavior of a household who faces a consumption–savings problem. Each period, this individual makes a decision that yields a consumption for that period and wealth for the next period.

More generally, a finite horizon decision problem should be thought of as a decision tree. Every period $t = 1, \ldots, T$, the agent chooses from a set of alternatives. In the final period (period T), the decision problem specifies a (compact) set of consumption choices. We allow consumption to be stochastic hence period-T choices yield a consumption lottery. In period $T - 1$, a decision problem is a set of alternatives, each yielding a lottery over consumption for period $T - 1$ and decision problems for period T. We continue in this fashion to define decision problems for periods $T - 2, T - 3, \ldots$, etc. The notation B_t is used to denote a period-t decision problem and $\mathcal{B}_t$ denotes the (collection of) period-t decision problems. Hence, B_t is a (compact) set of lotteries which yield a period-t consumption $c_t \in C$ and a decision problem for period $t + 1$, B_{t+1}. We use p_t to denote an element of B_t. Note that p_t is a probability measure defined on $C \times \mathcal{B}_{t+1}$.

To describe the behavior of the agent in this multi-period setting we analyze preferences over decision problems starting in period 1, $\mathcal{B}_1$. These preferences capture the period-0 behavior of this agent. Note that we assume that in period 0 there is no consumption.

Standard assumptions needed for a separable expected utility representation together with a multi-period version of (SB) yield the following recursive representation of self-control preferences:

$$W_{t-1}(B_t) = \max_{p_t \in B_t} \int \big(u_t(c) + W_t(c, B_{t+1}) + V_t(c, B_{t+1})\big)\, \mathrm{d}p_t - \max_{p_t \in B_t} \int V(c, B_{t+1})\, \mathrm{d}p_t. \tag{3}$$

The function W_{t-1} represents the agent's preference over choice problems that "start" in period t, that is, choice problems where prior to period t the agent is committed to some consumption path. The continuous function $u + W_t$ is the commitment utility in period t

(analogous to u in the 2-period problem). The continuous function V_t is the temptation utility in period t (analogous to v in the two-period problem). In the terminal period (period T), there is no continuation problem and hence

$$W_{T-1}(B_T) = \max_{p_T \in B_T} \int \big(u_T(c) + V_T(c)\big)\,\mathrm{d}p_T - \max_{p_T \in B_T} \int V_T(c)\,\mathrm{d}p_T.$$

Note that once v is substituted for V_T and u is substituted for u_T the formula for terminal date choice problems is the same as the formula in the previous section.

As before, the representation suggests that the agent chooses a lottery $p_t \in B_t$ to maximize

$$\int \big(u_t(c) + W_t(c, B_{t+1})\big)\,\mathrm{d}p_t + \int V_t(c, B_{t+1})\,\mathrm{d}p_t \tag{4}$$

in subsequent periods. This behavior represents the optimal compromise between commitment and temptation utilities.

To illustrate dynamic self-control preferences consider the following three-period example. There are two consumption periods, 1 and 2, and one consumption good $c \in [0, \bar{c}\,]$. The temptation utility depends only on current consumption. In particular, assume that

$$V_1(c, B_{t+1}) = V_2(c) = v(c), \qquad u_t(c) = u(c)$$

with $u(c)$ and $v(c)$ increasing in c. Then,

$$W_1(B_2) = \max_{p_2 \in B_2} \int \big(u(c) + v(c)\big)\,\mathrm{d}p_2 - \max_{p_2 \in B_2} \int v(c)\,\mathrm{d}p_2$$

and

$$W_0(B_1) = \max \int \big(u(c) + v(c) + W_1(B_2)\big)\,\mathrm{d}p_2 - \max_{p_2 \in B_2} \int v(c)\,\mathrm{d}p_2.$$

For the special case, where all the elements B_2 are deterministic, we have

$$\arg\max_{B_2} u(c) = \arg\max_{B_2} v(c),$$

since both sides are simply the maximal consumption in B_2. In that case, temptation plays no role in period 2 and

$$W_1(B_2) = \max_{B_2} u(c).$$

In period 1, the agent is tempted to choose the maximally feasible consumption for that period whereas the commitment utility wishes to maximize $u(c_1) + u(c_2)$. We will use this setting in the following sections to provide simple examples of economies with self-control preferences and to contrast our model with the dynamically inconsistent models of hyperbolic discounting.

4. Self-control and time-inconsistency

Strotz (1956) proposes a model of changing preferences. Each period, the decision-maker is thought to have a distinct utility function. Consistent planning requires that the

decision-maker take into account future changes in the utility function and "*reject any plan that he will not follow through. His problem is then to find the best plan among those he will actually follow*" (Strotz, 1956).

Our approach does not postulate a change in preference. Nevertheless, the behavior of consistent planners as defined by Strotz emerges as a special case of temptation preferences if we restrict to finite deterministic choice problems. Consider a finite subset of consumption choices $\widehat{C} \subset C$. Let $\widehat{\mathcal{B}}_t$ be the collection of deterministic t-period choice problems corresponding to $\widehat{C}$. That means that if $B_t \in \mathcal{B}_t$ then all elements of B_t yield a particular consumption $c \in \widehat{C}$ and a choice problem $B_{t+1} \in \widehat{\mathcal{B}}_{t+1}$. We denote with (c, B_{t+1}) the degenerate lottery that yields c and B_{t+1} with probability 1.

Note that every $B_t \in \widehat{\mathcal{B}}_t$ is a finite set and, in addition, there are finitely many decision problems in $\widehat{\mathcal{B}}_t$. Therefore, we can choose $\lambda > 0$ large enough so that

$$\arg\max_{B_t}(u_t + W_t + \lambda V_t) \subset \arg\max_{B_t} V_t \tag{5}$$

for all $B_t \in \widehat{\mathcal{B}}_t$. But this means that the agent maximizes V_t in every choice problem $B_t \in \widehat{\mathcal{B}}_t$. Plugging Eq. (5) into the representation yields

$$\begin{aligned} &W_{t-1}(B_t) = u_t(c) + W_t(B_{t+1}) \\ &\text{subject to} \quad (c, B_{t+1}) \in \arg\max_{B_t} V_t(c, B_{t+1}). \end{aligned} \tag{6}$$

Equation (6) can be interpreted as the Strotz model of consistent planning: In period t, the agent maximizes V_t. In period 0, the agent evaluates period-t choices with a different utility function $u + W_t$. Consistent planning means that the individual treats period-t choices as a constraint when evaluating B_t. To further illustrate the connection to the time-inconsistency literature we draw on the Krusell et al.'s (2002, 2003) work:

$$u_t(c) = \delta^{t-1}u(c), \qquad V_t(c, B_{t+1}) = u(c) + \beta W_t(B_{t+1}), \qquad V_T(c) = \delta^{T-1}u(c).$$

Then, setting $\widetilde{W}_t = \delta^{t-1}W_t$ we can rewrite (6) as

$$\begin{aligned} &\widetilde{W}_{t-1}(B_t) = u(c) + \delta\widetilde{W}_t(B_{t+1}) \\ &\text{subject to} \quad (c, B_{t+1}) \in \arg\max_{B_t} u(c) + \beta\delta\widetilde{W}_t(B_{t+1}) \end{aligned} \tag{7}$$

with $\widetilde{W}_{T-1}(B_T) = \max_{c\in B_T} u(c)$. The behavior of this agent corresponds to the $\beta-\delta$ model first introduced by Phelps and Pollak (1968) and later used by Laibson (1997). The commitment utility $u + W_t$ and the temptation utility $u + \beta W_t$ differ in how they discount the immediate future. Behavior maximizes $u + \beta W_t$ in every period but in period 0 this behavior is evaluated with the more patient utility function $u + W_t$.

When there are only two decision periods, the restriction to finite choice problems is not necessary. As we show in Gul and Pesendorfer (2001), the Strotz model is a special case of the self-control model if we do not require preferences to be continuous. When the model has more than 2 periods and choice problems are not finite, it is no longer possible to describe the Strotz model as the solution to a maximization problem (see Peleg and Yaari, 1973; Gul and Pesendorfer, 2002b). Instead, starting with Peleg and Yaari, authors have used game theory to solve dynamically inconsistent decision problems. One

implication of the game theoretic approach is that period-0 behavior cannot be described by a preference relation. Hence, the revealed preference approach runs into difficulty at a very basic level. In contrast, the self-control preferences described above are well-defined for general compact choice problems.

5. Predictions and evidence

Among psychologists and medical professionals it is held as self evident that "people often act against their self-interest in full knowledge that they are doing so; they experience a feeling of being 'out of control' " (Loewenstein, 1996). Evidence of out-of-control behavior is sought in the actions of drug addicts, or people who are subject to extreme emotional or physical stress.

Revealed preference theory *defines* the interest of people to be what they do. Since there is no objective standard of behavior it is unclear what it would mean for an agent to act against his self-interest. To incorporate visceral influences into a revealed preference theory, we must identify the presence of visceral influences from the agent's behavior alone without reference to an external standard. In other words, we must find a subjective notion of 'acting against one's self-interest.'

Our model does this by analyzing behavior at a stage where the agent is not (yet) subject to temptation (or other visceral influences) but rather chooses among situations with differing temptations. These choices reveal how the agent evaluates the impact of those temptations and hence establishes a subjective notion of self-interest free of visceral influences. Hence, to distinguish between standard utility maximizing behavior and behavior that is subject to visceral influences we must consider behavior where subjects choose between "choice situations," that is, decision problems.

Standard behavior is identified by Axiom (S). In particular, Axiom (S) rules out any preference for commitment, that is, a strict preference for fewer options. In contrast, temptation preferences but also models of dynamically inconsistent behavior allow for a preference for commitment. Hence, violations of (S) that exhibit a *preference for commitment* would be a natural starting point to look for evidence.

There are few studies that examine whether subjects have a preference for commitment. Wertenbroch (1998) finds evidence that people buy smaller quantities of tempting goods even when those goods are sold with quantity discounts. Ariely and Wertenbroch (2002) let students choose whether to impose deadlines for class assignments and find that many students choose to impose deadlines that constrain their future choices. Green and Rachlin (1996) conduct experiments with pigeons that finds evidence in favor of a preference for commitment.

A much larger body of evidence deals with a related phenomenon, preference reversals. Subjects are asked to choose between smaller–earlier and larger–later rewards. Experiments document the following preference reversal. When the delay to both rewards is increased, subjects tend to switch their preference from the earlier to the later reward (see Frederick et al., 2002 for a survey of this and related experimental literature).

To see how the temptation model can account for preference reversals, we consider the 3-period example introduced at the end of Section 3. All choices are assumed to be

deterministic, there is one consumption good and $c_t \in [0, \bar{c}\,]$. The period-0 preference over decision problems is represented by the following utility function W_0:

$$W_0(B_1) = \max\bigl(u(c_1) + \delta W_1(B_2) + v(c_1)\bigr) - \max v(c_1), \qquad W_1(B_1) = \max u(c_2).$$

The above utility function represents an individual who is tempted by *immediate consumption*. In the terminal period (period 2), he consumes all the remaining endowment and hence temptation plays no role ($W_1 = \max u(c_2)$). In period 1, the agent is tempted by current consumption since the temptation utility v depends on c_1 only.

In period 1, the agent maximizes $u + v + \delta W_1$. Plugging in for W_1 we can simplify the agent's decision problem in period 1 to

$$\max\bigl(u(c_1) + v(c_1) + \delta u(c_2)\bigr).$$

Suppose the agent (at date 1) can choose to transfer resources to date 2 (or vice versa) at a rate $1 + r$. Assuming u and v are differentiable the agent will reduce consumption in period 1 if

$$1 + r > \frac{u'(c_1) + v'(c_1)}{r\delta u'(c_2)}.$$

Now suppose that the agent is asked to choose consumption in periods 1 and 2 at date 0. By making the choice in date 0 we "increase the delay to both rewards." To be consistent with the experimental findings the agent should now make more patient choices.

Note that if the agent commits in period 0 to a consumption choice (c_1, c_2) then there is no temptation in period 1. That is, the maximal period-1 consumption in c_1 and hence the v-terms drop out. Therefore commitment to (c_1, c_2) in period 0 yields the utility

$$u(c_1) + \delta u(c_2).$$

In that case, the agent will reduce consumption in period 1 if

$$1 + r > \frac{u'(c_1)}{\delta u'(c_2)}.$$

Hence, period-0 commitment choices imply a lower rate of time preference. The reason is that commitment avoids the utility cost of temptation associated with transferring resources from period 1 to 2.

Note that our interpretation of the experimental evidence assumes that the agent is committed to his period-0 consumption choices. In particular, this means that at date 1 he cannot "undo" his choices by borrowing or simply consuming a greater share of his period-1 income. If the agent is not committed to a consumption path but rather faces a lifetime budget constraint then neither self-control preferences nor dynamically inconsistent models would predict a preference reversal. By a simple arbitrage argument the agent would pick the alternative with the higher net present value at the market interest rate independent of the timing of the choice. Therefore, the discounting evidence is at best indirect evidence for the temptation or the dynamic inconsistency model. Unlike a direct test based on choice problems, the discounting evidence relies on the assumption that the delayed rewards offer commitment.

As we argued above, our model includes the behavior of agents with β–δ preferences as a special case. However, unlike the latter model, we allow for self-control. As with

visceral influences, it is not clear how we could determine whether an agent exercises self-control when he makes a particular choice. Our method of identifying self-control is to study behavior prior to moment when the need for self-control may arise. When the agent chooses among decision problems he may reveal whether or not he *expects* to exercise self-control. Consider a situation where the agent prefers commitment to c over the choice between the lotteries c and c'. As we argued in Section 2, the agent expects to use self-control when facing $B = \{c, c'\}$ if

$$\{c\} \succ \{c, c'\} \succ \{c'\}.$$

We interpret the above preference to represent a situation where the agent chooses c but is tempted by c' and therefore incurs a positive cost of self-control. The dynamic inconsistency literature assumes that self-control does not occur. That is,

$$\{c, c'\} \sim \{c'\}.$$

In other words, the agent does not benefit when c is added to a choice problem that contains c'. We are not aware of any experimental evidence about self-control.

In the following two subsections we contrast the behavior of agents with and without self-control when they have commitment opportunities such as illiquid assets. In these settings clear differences emerge between agents who exercise self-control and dynamically inconsistent agents who do not exercise self-control.

5.1. *Illiquid assets*

A preference for commitment suggests that agents have a demand for commitment devices. One example of an institution that may provide some commitment are *illiquid assets*, such as housing. In this section, we provide a simple example of a two-period representative household economy with liquid and illiquid assets.

We first analyze a model with self-control preferences. In the competitive equilibrium the representative household holds both liquid and illiquid assets. The illiquid asset carries a commitment premium since it allows the agent to reduce the cost of self-control.

Kocherlakota (2001) analyzes a version of this example for time-inconsistent preferences. He demonstrates (and we illustrate below) that the model with one household does not admit a competitive equilibrium in this case. The reason is that the agent can never simultaneously hold the liquid and the illiquid asset. To find a competitive equilibrium we must assume a continuum of identical households with some fraction of households only holding the liquid assets and the remainder holding only illiquid assets.

Consider the following three-period economy. There is one consumption good $c \in [0, \bar{c}\,]$. In period 0, there is no consumption but agents trade assets. The period-0 endowment of assets is the only endowment in the economy.

There are three assets, indexed $j \in \{1, 2, 3\}$. Asset 1 is traded in period 0 and returns two units of consumption in period 1. Assets 2 is traded in periods 0 and 1 and returns one unit of consumption in period 2. Assets 3 is traded only in period 0 and returns one unit of consumption in period 2. Hence assets 2 and 3 have the same return but asset 3 is not traded in period 1. Furthermore, short selling of assets is prohibited. Therefore, asset 3 offers commitment because it cannot be sold (or borrowed against) in period 1. For

example, if the agent holds one unit of asset 3 and no other assets, then in period 1 he is committed to 1 unit of consumption in period 2. Let z_{tj} denote the agent's holdings of asset j in period t and let q_{tj} denote the price of asset j in period t.

We assume a representative agent who is endowed with 1 unit of each asset.

5.1.1. Self-control preferences

Consider the following self-control preference

$$W_0(B_1) = \max_{B_1}\big(u(c_1) + \lambda u(c_1) + W_1(B_2)\big) - \lambda \max_{B_1} u(c_1),$$
$$W_1(B_2) = \max_{B_2} u(c_2),$$

where u is concave and $u(c) - \lambda u(x + c)$ is concave for all $x \in [0, \bar{c}\,]$.

In this economy, the choice problems B_1 and B_2 depend on the agent's asset holdings. We write $B_1(z_0)$ to denote the choice problem generated by the asset holding $z_0 = (z_{01}, z_{02}, z_{03})$. The choice problem B_2 is trivial: the agent simply consumes all the remaining wealth in period 2. Therefore, we can simplify $W_0(B_1(z_0))$ as follows:

$$W_0\big(B_1(z_0)\big) = \max_{(c_1,z_{12})\in b_1(z_0)} \big(u(c_1)(1+\lambda) + u(z_{03} + z_{12})\big) - u(2z_{01} + q_{12}z_{02}) \tag{8}$$

with

$$b_1(z_0) = \big\{(c_1, z_{12}) \mid c_1 + q_{12}z_{12} = 2z_{01} + q_{12}z_{02}\big\}.$$

We normalize $q_{03} = 1$. The agent's budget constraint in period 0 is

$$b_0 := \big\{(z_{01}, z_{02}, z_{03}) \mid 1 - z_{03} = q_{01}z_{01} + q_{02}z_{02}\big\}.$$

In period 0, the agent chooses z_0 to solve

$$\max_{z_0\in b_0} W_0\big(B_1(z_0)\big).$$

Since this is a representative agent economy, in a competitive equilibrium the household must choose $z_{0j} = 1$ and $z_{12} = 1$. Moreover, consumption must equal 2 in both periods. A straightforward calculation (using the first order condition of the maximization problem on the right-hand side of (8)) implies that

$$u'(2)(1+\lambda) = u'(2)q_{12}$$

and hence

$$q_{12} = 1/(1+\lambda).$$

Plugging q_{12} into $W_0(B_1(z_0))$ and substituting for the equilibrium quantities then allows us to solve for the equilibrium prices. Normalizing $q_{03} = 1$, we get

$$q_{01} = 2\frac{u'(2)(1+\lambda) - \lambda u'(2 + 1/(1+\lambda))}{u'(2)},$$
$$q_{02} = \frac{u'(2) - \frac{\lambda}{1+\lambda}u'(2 + 1/(1+\lambda))}{u'(2)},$$
$$q_{03} = 1.$$

For $\lambda = 0$, this yields the expected result that $q_{02} = q_{03}$. For $\lambda > 0$, asset 3 commands a commitment premium. The reason is that holding asset 3 reduces the cost of self-control in period 1.

Notice that the self-control model allows us to infer asset prices just like in a standard Lucas tree economy. The commitment premium is derived by assessing the marginal effect of asset holdings on the cost of self-control.

5.1.2. Time-inconsistent preferences

Next assume that the agent has $\beta-\delta$ preferences. In particular, assume that the period-0 utility function is

$$u(c_1) + u(c_2),$$

whereas the period-1 utility function is

$$u(c_1) + \beta u(c_2)$$

with $0 < \beta < 1$. In period 2, the agent simply maximizes $u(c_2)$.

In this case, there does not exist an equilibrium for strictly concave u. Intuitively, the failure of existence can be explained as follows: In the dynamically inconsistent model, illiquid assets can only offer valuable commitment if the agent does not hold *any* liquid asset. As long as the agent holds some liquid assets he is not committed *at the margin* and therefore is not willing to pay a commitment premium. But if the commitment premium is zero then there is a corner solution where the agent achieves full commitment.

To verify the nonexistence of equilibrium formally, first recall that equilibrium requires $c_t = 2,\ t = 1, 2,\ z_{0j} = 1,\ j = 1, 2, 3$ and $z_{12} = 1$. Furthermore, at date 1 the equilibrium price of asset 1 must satisfy

$$q_{12} = \beta$$

since in period 1 the agent must be indifferent between transferring wealth between periods 2 and 3. Normalize $q_{01} = 2$. Suppose the agent increases the holding of asset 3 by ϵ and finances this increase by an ϵ reduction in the holding of assets 2. At the allocation $z_{0j} = 1, \forall j$, this change does not affect the consumption in any period. Since u is strictly concave and $q_{12} = \beta$, the agent will not change his asset holding in period 1. Therefore, the change in the portfolio leaves the agent indifferent. But this implies that at an equilibrium we must have

$$q_{02} = q_{03},$$

and since no arbitrage requires,

$$q_{02} = q_{01}q_{12}/2,$$

it follows that

$$q_{02} = q_{03} = \beta.$$

However, at these prices the agent is not willing to hold his endowment of assets. Holding the endowment yields a utility of

$$u(2) + u(2).$$

By holding only assets 1 and 3 the agent can achieve the utility

$$\max u(c_1) + u(c_2)$$
$$\text{subject to} \quad c_1 + \beta c_2 = 2(1+\beta)$$

which is greater than the utility of holding the endowment.

To get an equilibrium for the time-inconsistent model, we must assume that there is a continuum of identical households. Some fraction α will hold *only* the illiquid asset whereas the fraction $1-\alpha$ of households holds the liquid asset. As Kocherlakota (2001) emphasizes, the prediction of *exclusive holding* is contradicted by the data. People who hold illiquid assets such as houses or IRAs, also hold liquid assets such as bank accounts. This suggests that *if* illiquid assets have a role as commitment devices then the self-control model is the more appropriate model to incorporate this role.

5.2. Social security

In this section, we illustrate the welfare effects of a simple "inter-generational" transfer. Consider the following 3-period model. In period 0, the agent must choose a transfer $\tau \in [0, 1]$. In period 1, the agent is endowed with 1 unit of wealth. The government taxes τ units of the agent's endowment. The agent must decide how much to consume (c_1) and how much to save ($1-\tau-c_1$). Saving is constrained to be non-negative. In period 2, the agent receives a subsidy of τ and hence his period-2 wealth (and period-2 consumption) is $\tau + (1+r)(1-\tau-c_1)$.

5.2.1. Self-control preferences

The agent has the utility function

$$W_0(B_1) = \max\big(u(c_1) + W_1(B_2) + \lambda u(c_1)\big) - \max_{B_1} \lambda u(c_1), \qquad W_1(B_2) = \max_{B_2} u.$$

The choice problem B_1 depends on the transfer τ. We write $B_1(\tau)$ to denote the choice problem when the transfer is τ.

Assume that $\lim_{c\to 0} u'(c) = \infty$. Then for τ sufficiently small, optimal behavior satisfies $c_1 < 1-\tau$. That is, the agent chooses to save some of his period-1 wealth. Then, we can plug in the budget constraint to get the following expression for $W_0(B_1(\tau))$:

$$W_0\big(B_1(\tau)\big) = \max_{c_1}\big(u(c_1)(1+\lambda) + u\big((1+r)(1-c_1) - r\tau\big)\big) - \lambda u(1-\tau).$$

An increase in τ has two effects; it reduces the cost of self-control in period 1 ($\lambda u(1-\tau)$ is reduced) and therefore increases welfare. An increase in τ also affects the consumptions c_1, c_2. When $r > 0$, a straightforward revealed preference argument establishes that this effect is negative because an increase in τ reduces the agent's wealth. For r sufficiently close to zero the overall welfare effect of an increase of τ is positive.

5.2.2. Time-inconsistent agents

Imrohoroglu et al. (2002) analyze equilibria of an overlapping generations model with agents who are time-inconsistent. Their calibration (Table 7) shows that unfunded social

security is unlikely to find support in such an economy. At most 6% of the population would vote for introduction of social security in their model.

Below, we illustrate why agents with time-inconsistent preferences may not benefit from the introduction of a social security transfer. In period 1, the agent has the utility function

$$u(c_1) + \beta u(c_2).$$

In period 0, the agent evaluates the subsequent choices with the utility function

$$u(c_1) + u(c_2).$$

Note that there is no consumption in period 0 and hence period-0 utility function does not exhibit a presence bias. If we set $\beta = 1/(1+\lambda)$ then the β–δ model and the self-control model make the same predictions about behavior in periods 1 and 2. That is, both models predict the same consumption pattern c_1, c_2. In addition, both models predict the same consumption choices if the agent can commit to c_1, c_2 in period 0.

However, the two models differ in how the agent evaluates the effect of the transfer τ in period 0. Again, consider the case where τ is sufficiently small so that the optimal choice for c_1 is strictly less than $1-\tau$. In this case, the optimal consumption choice solves the following unconstrained maximization problem:

$$\begin{aligned} &\max u(c_1) + \beta u(c_2) \\ &\text{subject to} \quad c_1 + c_2/(1+r) \leqslant 1 - r\tau. \end{aligned}$$

An increase in τ therefore amounts to a reduction in overall wealth in this case. Such a reduction can never increase the agent's utility in period 0.

While the self-control model predicts that a small intergenerational transfer may benefit the individual such a transfer can never benefit the time-inconsistent agent in period 0. The reason is that even after the transfer is introduced, young agents save ($c_1 < 1-\tau$) and hence are unconstrained in their period-1 consumption. But this implies that there is no commitment benefit to the transfer. It is this feature of the time-inconsistent model that leads to the Imrohoroglu et al. (2002) conclusion that at most 6% of the US population is likely to support social security.

In contrast, self-control preferences allow for the possibility that agents benefit from non-binding constraints. The intergenerational transfer can be interpreted as such a non-binding constraint: it reduces wealth available for consumption when the agent is young but does not constrain the agent because even after its introduction saving is positive. As our example illustrates, when interest rates are low young agents will benefit from such a transfer.

A complete analysis of the effects of a transfer must also consider the price effects of such a policy. Imrohoroglu et al. (2002) do this for the case of time-inconsistent preferences. A similar analysis for self-control preferences has not been conducted.

5.3. Gambling and risk aversion

This section considers an agent who is tempted to gamble. This can be captured by a temptation utility that is a convex function of current consumption. This form of temptation

has two effects: agents are tempted by risky instantaneous lotteries but are risk averse when evaluating risky assets that pay-off in future periods.

As an illustration, we again consider a simple three-period model. A related analysis for the infinite horizon case can be found in Gul and Pesendorfer (2002a). There are two consumption periods, 1 and 2, and one consumption good $c \in [0, \bar{c}\,]$. The temptation utility depends only on current consumption. Let $u(c)$ and $v(c)$ be increasing functions with v convex and u concave. The preference over decision problems in period 2 is represented by

$$W_1(B_2) = \max_{p_2 \in B_2} \int \big(u(c) + v(c)\big)\,\mathrm{d}p_2 - \max_{p_2 \in B_2} \int v(c)\,\mathrm{d}p_2$$

and the preference over decision problems in period 1 is represented by

$$W_0(B_1) = \max_{p_2 \in B_2} \int \big(u(c) + v(c) + W_1(B_2)\big)\,\mathrm{d}p_2 - \max_{p_2 \in B_2} \int v(c)\,\mathrm{d}p_2.$$

Suppose, the agent has a constant endowment $\bar{c}$ and cannot borrow or save. Assume that a lottery with the (stochastic) period-1 payoff $R \in (-I, I)$ is offered to the agent at no cost. The agent will accept that lottery if

$$E\big(u(\bar{c} + R) + v(\bar{c} + R)\big) \geqslant u(\bar{c}) + v(\bar{c}).$$

Note that the continuation problem remains unaffected by the decision because we assumed the agent cannot transfer wealth between periods. Next, assume that the same lottery is offered but now the payment of the return is delayed by one period. That is, the agent decides in period 1 whether to accept the stochastic reward R for period 2. In this case, the agent compares two choice problems. If he rejects the lottery the period-2 choice problem is $B_2^r = \{\bar{c}\}$, i.e., the agent is committed to the consumption $\bar{c}$ in period 2. If he accepts the lottery the choice problem is $B_2^a = \{\bar{c} + R\}$, i.e., the agent is committed to the stochastic consumption $\bar{c} + R$. The agent accepts the lottery if

$$\begin{aligned} E\big(W_1\big(B_2^a\big)\big) &= E\big(u(\bar{c} + R) + v(\bar{c} + R) - v(\bar{c} + R)\big) = Eu(\bar{c} + R) \\ &\geqslant W_1\big(B_2^r\big) = u(\bar{c}\,). \end{aligned}$$

If v is convex and u is concave then $u + v$ is less risk averse than u and hence the agent is more willing to take risk with respect to lotteries that pay off *immediately*.

A risk loving temptation utility may help explain why the purchase of certain lotteries are prohibited or regulated while other risky investments such as stocks are not. In our interpretation, the difference between gambling and stock investing lies in the timing of the return. Stock investing offers risky returns with a delay whereas casino gambling offers risky returns with immediate rewards. The latter is a temptation and agents with a risk loving temptation utility are better off if such lotteries are not available.

6. Infinite horizon

In this section, we illustrate how self-control preferences can be applied to infinite horizon decision problems. The contrast between the self-control model and models of

time-inconsistent agents is most apparent in infinite horizon settings. The self-control model leads to standard dynamic programming problems that can be solved using standard techniques. Models of time-inconsistent decision making must be solved as dynamic games. With an infinite horizon there is typically a folk-theorem type multiplicity of equilibria in these games. In contrast, the self-control model implies that each decision problem can be assigned a unique value.

Infinite horizon decision problems can be described in a simple recursive manner: an infinite horizon decision problem is a (compact) set of lotteries B such that each lottery yields a current consumption and an infinite horizon decision problem B'.[4] As before, the model specifies a preference over decision problems. We assume throughout that these preferences can be represented by a utility function. In the infinite horizon case $W(B)$ denotes the utility of the decision problem B. This utility function is analogous to the value function in standard decision problems.

Assume that—with the exception of (S)—the preference satisfies the standard assumptions that yield separable and recursive preferences. We replace (S) with (SB). Then, the preference can be represented by a W that satisfies

$$W(B) = \max_{p\in B} \int \big(u(c) + \delta W(c, B') + V(c, B')\big)\,\mathrm{d}p - \max_{p\in B} \int V(c, B')\,\mathrm{d}p.$$

The functions u, V are continuous and $\delta \in (0, 1)$ (for details see Gul and Pesendorfer, 2002a).

The equation above defines a standard dynamic programming problem. The problem is recursive and represents a time-consistent preference in the following sense. In every period $t > 0$, the agent makes choices that maximize $u + \delta W + V$. These choices maximize the agent's utility (at any decision date). The preferences are different from standard utility functions in that the agent's utility depends directly on the decision problem because the problem affects his cost of self-control. The cost of self-control when the agent chooses p is the difference between the maximal temptation utility and the temptation utility of p, that is, $\max_{p'\in B} \int V\,\mathrm{d}p' - \int V\,\mathrm{d}p$.

Next, we specialize to the case where the temptation utility V depends only on current consumption. Hence, we assume

$$V(c, B) = v(c).$$

Consider a standard consumption–savings problem without uncertainty. The agent has wealth w_t in period t and faces a constant interest rate r. The individual chooses plans to maximize

$$W(w_t) = \max_{c_t}\big(u(c_t) + v(c_t) + \delta W\big((1 + r)(w_t - c_t)\big)\big) - v(w_t).$$

The first-order necessary conditions for an interior solution are

$$u'(c_t) + v'(c_t) = \delta(1 + r)\big(u'(c_{t+1}) + v'(c_{t+1}) - v'(w_{t+1})\big)$$

[4] For a formal definition of multi-period decision problems, see Gul and Pesendorfer (2002a).

where $w_{t+1} = (1+r)(w_t - c_t)$. Krusell et al. (2003) consider the special case where $v = \lambda u$ and $u = (1-\sigma)^{-1}c^{1-\sigma}$ for $\sigma > 0$. For this isoelastic case, the consumption function is linear, i.e., $c_t = \alpha w_t$ where α is defined as the solution to the equation

$$\frac{1+\lambda}{(1+r)\delta}\alpha^{-\sigma} = \big(\alpha(1-\alpha)(1+r)\big)^{-\sigma} - \lambda\big((1-\alpha)(1+r)\big)^{-\sigma}.$$

The fraction of wealth consumed by the individual depends on the interest rate and on the parameters of the utility function $(\delta, \lambda, \sigma)$. An increase in λ decreases the weight on the temptation utility and leads to a decrease in the individual's (instantaneous) self-control. As a consequence, the agent acts more impatiently and the proportion of wealth consumed increases.

As Krusell et al. (2003) point out, the steady-state interest rate (that is, the interest rate at which consumption is constant over time) satisfies

$$\frac{1+\lambda}{(1+r)\delta}\alpha^{-\sigma} = 1 - \lambda\alpha^{\sigma}$$

and therefore depends on the parameter σ. In particular, the steady-state interest rate is *decreasing* in σ. The reason is that σ measures the curvature of the utility function. At a steady state, a higher σ means a lower marginal cost of self-control.

The time-consistency approach in the familiar $\beta-\delta$ parametrization generates a related consumption function for the isoelastic consumption savings problem. (See, for example, Krusell et al., 2002.) The discounting parameter β plays a similar role as the self-control parameter λ in our model. However, this close connection only holds superficially. The time-inconsistency model permits many equilibria. The comparative statics prediction are therefore specific to the particular equilibrium selected. For example, at the equilibrium that is most preferred by the period-0 self, if δ sufficiently high the parameter β has no influence on the savings behavior of the agent. In this equilibrium (for high δ), the agent picks the optimal plan for δ and punishes deviations by reverting to a sufficiently "bad" equilibrium. As a second example, consider the equilibrium with the highest savings in a given period. For high δ it is straightforward to construct equilibria where the agent saves more for $\beta < 1$ than the corresponding agent with $\beta = 1$. Hence, in general the effect of a change in β on the savings behavior remains ambiguous in the time-inconsistency model.

Time-inconsistent models specify a dynamic game and hence the multiplicity of outcomes (and values) is one aspect of the more general phenomenon of multiplicity of equilibria in dynamic games. However, there is a key difference between a multi-person setting and a single person decision problem with a time-inconsistent agent. In a multi-person context, subgame perfect Nash equilibrium is meant to capture a rest point of the player's expectations and strategizing. The multiplicity of equilibria reflects the fact that there is no communication between players and no single player can coordinate play.

In a time-inconsistent decision problem, it seems straightforward to coordinate with one's future self or to renegotiate one's self out of an unattractive continuation equilibrium.[5] Therefore, the standard argument for why multiplicity should be expected

[5] The argument that renegotiation is particularly plausible in time-inconsistent decision problems is due to Kotcherlakota (1996).

has no force in this context. More generally, the appropriateness of (subgame perfect) Nash equilibrium often rests on the assumption of "independent" behavior and absence of communication. But clearly, the agent at time t and his slightly modified self at time $t+1$ are able to communicate and coordinate and hence Nash equilibrium may not be the appropriate solution concept.[6]

7. Interpretation and welfare

Our model and time-inconsistent models have in common the feature that an agent may choose a different period-t consumption depending on when he chooses. The time-inconsistency approach interprets this divergence as a change in preference. Our model permits the agent's preference over consumption in period t to depend on the choice set from which period-t consumption is chosen. If period-t consumption is chosen in period 0 then this choice is unaffected by temptation. If period-t consumption is chosen in period t then temptations may distort the agent's consumption choice. This distortion of behavior is an optimal response to the presence of temptations taking into account the cost of self-control.

Suppose the agent chooses to commit at date 0 not to smoke at any future date. At date $t > 0$ after unexpectedly being offered a cigarette the agent begins to smoke. Similar examples are typically described as proof of dynamic inconsistency. However, such reversals may arise even with consistent preferences.[7] The agent's period-t choice only reveals that when cigarettes are available he prefers to smoke. Hence acceptance of the cigarette in period t is consistent with a preference for a situation where cigarettes are unavailable (as his previous choice indicates). Conversely, the period-0 choice is not in conflict with his period-t behavior since in period 0 the agent did not expect that the promise of commitment would be broken in period t.

Our model differs from the time-inconsistent model in two ways. First, it permits self-control. Second, it has clear, testable welfare implications. In terms of welfare implications our model follows the revealed preference tradition of standard economic models: if the agent chooses one alternative/policy over another then he is better off with that choice. However, since our agents value commitment, it is essential to include the timing of choices in the specification of the agents' alternatives.

For example, we deem option a to be better for the agent than option b if the agent chooses a over b. Since the agents well-being is influenced by both his consumption and the temptation he suffers, in our models a satisfactory description of the agents decisions entails both the consumptions he enjoys and the temptations he suffers.

Consider, for example, a tax policy. Assume that in every period the agent must choose (deterministic) consumption from a fixed budget that depends on the tax τ. Let $B(\tau)$ denote

[6] For a related critique of the use of Nash equilibrium to model a (different) departure from fully rational behavior, see Piccione and Rubinstein (1997).

[7] Our time-consistent interpretation is in the spirit of Machina (1989), who provides a similar resolution to the dynamic inconsistency that arises from the failure of the independence axiom.

the decision problem when the tax is τ in every period. Assume that changes in the tax rate are permanent.

The agent's utility is given by

$$W\big(B(\tau)\big) = \arg\max_{c\in B(\tau)} u(c) + \delta W\big(B(\tau)\big) + V\big(c, B(\tau)\big) - \arg\max_{c\in B(\tau)} V\big(c, B(\tau)\big).$$

Suppose that we must determine the welfare consequences of a change in the tax to τ'. First, suppose this change is considered in period 0. If $W(B(\tau')) - W(B(\tau)) > 0$ then the change to τ' in period 0 improves the agent's welfare.

Now suppose that the tax change is considered in period $t > 0$. In a dynamically consistent model the agent's welfare corresponds to his choice behavior. That means that we need to determine whether the agent would choose the tax increase in period $t > 0$. In period t, the agent cares about period-t consumption and the continuation problem. Choosing between the two tax rates therefore is equivalent to choosing the period-t alternative that maximizes $u + \delta W + V$ in the feasible set of alternatives compatible with the two tax rates. Hence, a tax change in period t improves the agent's welfare if

$$\max_{c\in B(\tau')} E_p\big(u(c) + \delta W\big(B(\tau')\big) + V\big(c, B(\tau')\big)\big) > \max_{c\in B(\tau)} E_p\big(u(c) + \delta W\big(B(\tau)\big) + V\big(c, B(\tau)\big)\big).$$

When the tax change is made in period t, the agent is not committed to $B(\tau')$ in period t. Rather, he contemplates the *choice* between the two tax rates. This means that the agent must evaluate the optimal period-t consumption and continuation problem consistent with either tax rate. In contrast, when the tax change is proposed in period 0, the agent chooses between the two choice problems $B(\tau)$ and $B(\tau')$. The period-0 choice is unaffected by temptation. Hence, choosing the tax rate τ' in period 0 and choosing tax rate τ' in period 1 are different outcomes for agent with self-control preferences. Both the agent's behavior and our welfare criterion reflects this difference.

The key property of our (and any) dynamically consistent model is that any two policies that are feasible at time t can be evaluated using information that is available at time t. If policy a is welfare improving then it must be the case that the agent is willing to choose that policy over the status quo alternative b. Hence, the theory does not allow paternalistic welfare statements where the agent rejects a policy even though the model describes it as welfare improving. For example, the tax policy analysis above establishes that for the specified range of parameters, the agent's first choice is to commit in period t to the tax rate τ' (to be implemented in period $t + 1$). If such a commitment is unavailable, and the agent has to choose the tax rate in period t that is to be applied starting with period t, then he prefers the rate τ to τ'. Note that once period t is reached commitment to a tax policy that is effective in period t is no longer an option. Therefore the question of whether he would have liked such a commitment is irrelevant.

In contrast, the time-inconsistency literature takes the view that the agent in period t is a different person than the agent in period $t + 1$. Policy experiments will often increase the utility of some self and at the same time decrease the utility of other selves. In the time-inconsistency literature, the utility of the period-0 self is sometimes used as the welfare criterion. Since the individuals' preferences are changing over time, there is no way to find

out which tax policy the agent would have preferred at time 0, once we reach period t. Therefore, policy analysis with this criterion boils down to the modelers assessment of what would have made the period-0 self—an agent who is long dead—better off. Even if the preferences of the period-0 self happen to be known, it does not seem reasonable for a social planner to impose a policy option on the grounds that it improves the welfare of an agent who is no longer present.

Instead of choosing the period-0 utility function as a welfare criterion, some authors (see Laibson, 1997) have proposed the stronger requirement that *all* selves be made better off by a welfare improving policy. Unlike the earlier criterion, this one does not impose policy recommendations based on the proported preferences of an economic agent that is no longer in existence. However, it uses such preferences to *block* policy alternatives. For example, it may be the case that at time T all of the period-t selves for $t \geqslant T$ prefer an option a to the status quo alternative b but the policy-maker concludes that a does not dominate b since one or more of the selves prior to time T would have preferred b. Note that any attempt to remedy this problem will lead to a welfare criterion that is itself dynamically inconsistent. Hence any welfare criterion for time-inconsistent agents is either inconsistent itself or has the planner forever guarding the perceived interests of nonexistent former selves.

Acknowledgments

This research was supported by grants SES9911177, SES0236882, SES9905178 and SES0214050 from the National Science Foundation.

References

Ariely, D., Wertenbroch, K., 2002. Procrastination, deadlines and performance: Using precommitment to regulate one's behavior. Psychological Science 13 (3), 219–224.

Dekel, E., Lipman, B., Rustichini, A., 2001. Representing preferences with a unique subjective state space. Econometrica 69, 891–935.

Frederick, S., Loewenstein, G., O'Donoghue, T., 2002. Time discounting and time preference: a critical review. Journal of Economic Literature XL (2), 351–402.

Green, L., Rachlin, H., 1996. Commitment using punishment. Journal of the Experimental Analysis of Behavior 65, 593–601.

Gul, F., Pesendorfer, W., 2001. Temptation and self-control. Econometrica 69 (6), 1403–1436.

Gul, F., Pesendorfer, W., 2002a. Self-control and the theory of consumption. Mimeo. Econometrica, in press.

Gul, F., Pesendorfer, W., 2002b. The revealed preference theory of changing tastes. Mimeo. Review of Economic Studies, in press.

Imrohoroglu, I., Imrohoroglu, S., Joines, D., 2002. Time inconsistent preferences and social security. Quarterly Journal of Economics. In press.

Kocherlakota, N.R., 1996. Reconsideration-proofness: A refinement for infinite horizon time inconsistency. Games and Economic Behavior 15 (1), 33–54.

Kocherlakota, N.R., 2001. Looking for evidence of time-inconsistent preferences in asset market data. Federal Reserve Bank of Minneapolis Quarterly Review 25 (3), 13–24.

Kreps, D., 1979. A preference for flexibility. Econometrica 47, 565–576.

Krusell, P., Smith, A., 2003. Consumption–savings decisions with quasi-geometric discounting. Econometrica 71, 365–377.

Krusell, P., Kuruscu, B., Smith Jr., T., 2002. Equilibrium welfare and government policy with quasi-geometric discounting. Journal of Economic Theory 105, 42–72.

Krusell, B., Kuruscu, B., Smith Jr., T., 2003. Temptation and taxation. Mimeo.

Laibson, D., 1997. Golden eggs and hyperbolic discounting. Quarterly Journal of Economics 112, 443–477.

Loewenstein, G., 1996. Out of control: Visceral influences on behavior. Organizational Behavior and Human Decision Processes 62, 272–292.

Machina, M.J., 1989. Dynamic consistency and non-expected utility models of choice under uncertainty. Journal of Economic Literature 27, 1622–1668.

O'Donoghue, T., Rabin, M., 1999. Doing it now or later. American Economic Review 89 (1), 103–124.

Peleg, M., Yaari, M.E., 1973. On the existence of a consistent course of action when tastes are changing. Review of Economic Studies 40, 391–401.

Phelps, E.S., Pollak, R.A., 1968. On second-best national saving and game-equilibrium growth. Review of Economic Studies 35, 185–199.

Piccione, M., Rubinstein, A., 1997. On the interpretation of decision problems with imperfect recall. Games and Economic Behavior 20 (1), 3–24.

Strotz, R.H., 1956. Myopia and inconsistency in dynamic utility maximization. Review of Economic Studies 23 (3), 165–180.

Wertenbroch, K., 1998. Consumption and self-control by rationing purchase quantities of virtue and vice. Marketing Science 17, 317–337.