ELSEVIER European Economic Review 44 (2000) 1116–1124

EUROPEAN ECONOMIC REVIEW

www.elsevier.com/locate/econbase

Modelling time-inconsistent preferences

Bernard Caillaud[a,b], Bruno Jullien[b,c,*]

[a]*CERAS-ENPC, CNRS URA2036, Paris, France*
[b]*Centre for Economic Policy Research, London, UK*
[c]*GREMAQ, CNRS UMR5604, and IDEI, Toulouse, France*

Abstract

The paper reviews some recent theoretical contributions on the modelization of time-inconsistent preferences, as well as implications for individual behavior. The focus is on the interpretation of the concepts and the link with concepts in psychology. © 2000 Elsevier Science B.V. All rights reserved.

JEL classification: D9

Keywords: Time-inconsistent preferences; Hyperbolic discounting; Self-restraint

1. Introduction

Time consistency is one of the central assumption of traditional utility analysis. Under time-consistent preferences, the marginal rate of substitution (MRS) between consumption at date $t+1$ and date t is constant equal to the discount factor δ.

There is, however, ample evidence that individuals exhibit time-inconsistent preferences (Ainslie, 1992; Lowenstein and Prelec, 1992; Rabin, 1998). Preferences are said to be time-inconsistent when the MRS between consumption at two future dates depends on the date at which it is evaluated. The most common

* Correspondence address: GREMAQ, Université des Sciences Sociales, Place Anatole France, 31042 Toulouse Cedex, France. Tel.: +33-561128561, fax: +33-561225563.
E-mail address: bjullien@cict.fr (B. Jullien)

PII: S0014-2921(99)00061-6

form of time-inconsistency (Strotz, 1956) corresponds to hyperbolic discounting, where this MRS decreases with the horizon. Hyperbolic discounting reflects a salience effect (Akerlof, 1991) of immediate consumption as opposed to delayed consumption. While two streams of consumption starting from tomorrow may be viewed as equivalent today, the one with the most immediate consumption will be the preferred one tomorrow, once time will have elapsed and the prospect of immediate consumption will be salient. This leads to the tendency to choose actions inducing short-run benefits over those inducing long-run benefits.

Research in economics often takes time inconsistency as an intrinsic property of preferences and analyzes the behavior of a 'sophisticated' individual, fully aware of her inconsistency. This behavior involves reflexive conciousness (the individual is able to make an objective assessment of her behavior and of her preferences), so that observed inconsistency need not arise as a sophisticated individual takes into account her future choices. What is meant by 'sophisticated' is critical and captured by a multi-selves game-theoretic approach (Phelps and Pollack, 1968; Laibson, 1997) or a single-self paradigm (Caillaud et al., 1999).[1]

But time-inconsistency can also derive from more basic psychological phenomena. Emotions can result in such inconsistencies (Elster and Loewenstein, 1992; Caplin and Leahy, 1998). They are also systematic biases in beliefs that are incompatible with full awareness (see O'Donoghue and Rabin (2000a,b) for an economic perspective). In particular anticipations of future changes in taste seem to be biased by inertia (Loewenstein et al., 1999), and individuals often overestimate the duration of their affective reactions, which can be affected by an unconscious psychological 'immune system' (Gilbert et al., 1998).

The object of this paper is to propose a short survey of some of these questions.

2. The model and quasi-hyperbolic discounting

An individual must decide upon her consumption of an addictive good at periods $t \in \{0, 1, \dots, \infty\}$.[2] There are two levels of addiction, addicted and non-addicted, and two consumption decisions: the addicted individual can stop (S), in which case she remains addicted with probability p and get disintoxicated with probability $(1-p)$, or consume (P) and remain addicted for sure next period. Once free from addiction, the individual never wants to consume again and stay so. At $t = 0$, the individual is addicted. Utility depends on consumption

[1] See also Thaler and Shefrin (1981) for a model with two selves.

[2] See Boyer (1983) and Becker and Murphy (1988) for rational models of addiction.

and on the level of addiction. The per-period utility of consuming while addicted is normalized to 0. H denotes the per-period utility when non addicted, and C is the per-period cost of disintoxication (choosing S while still addicted).

In the standard theory of time-consistency, the individual uses a discount factor $\delta \in (0,1)$ to evaluate a flow of utilities. A strategy indicates which action to take at each period when still addicted. The constant strategy 'S always' yields

$$J^C = \frac{(1-p)\delta H/(1-\delta) - C}{1-\delta p},$$

which is assumed positive (that is, $X \equiv (1-\delta)C/\delta(1-p)H < 1$). The strategy 'P always' yields intertemporal welfare equal to 0.

Because the problem is stationary, the optimal strategy that maximizes the individual's intertemporal utility is also stationary: it is the strategy 'S always'. Note that, planning 'S always' involves a decision that could a priori be reversed in the future. But if disintoxication is the best choice at time t, it is also the best choice any other time $s > t$. In other words the strategy 'S always' is time-consistent.

Quasi-hyperbolic discounting, as developed by Phelps and Pollack (1968), corresponds to one possible specification of time-inconsistent preferences,

$$J_t = u_t + \beta\left(\sum_{s=1}^{\infty} \delta^s u_{t+s}\right) \quad \text{with } \beta \in (0,1).$$

The salience effect of present consumption implies that the individual does not evaluate future utility streams in the same manner at different dates. The individual may then rationally intend to follow a path of consumption that is different from the one she would have chosen if a commitment device were available. It does not, however, imply that there is a contradiction between what the individual plans to do at date t and what she actually does when t has come, since an individual who is aware of her time inconsistency will anticipate her future choices and choose consequently. A theory of rational decision under time inconsistency must thus describe how the individual anticipates her future choices.

3. Multi-selves: Game theory

The prevalent approach (Phelps and Pollack, 1968; Laibson, 1997) views the individual as a collection of different selves. The different incarnations of the individual at different dates are different players in a dynamic game, player t choosing at t according to utility J_t.

Focusing on subgame perfect equilibria (PE) of the game between the infinite number of selves formalizes the absence of commitment. The possibility of coordination with future selves leads one to focus on 'symmetric' PE, that is on

PE that yield the same value index for every continuation path (on and off the equilibrium path, see Kocherlakota (1996)). In our context, symmetric PE are identical to stationary Markov perfect equilibrium (SMPE), that is on PE where the choice at t does not depend on payoff-irrelevant history and is constant over time.

In our setting, an SMPE is thus a constant action a^M such that choosing a^M today is optimal when the current self anticipates that all future addicted selves will choose a^M, irrespective of what she is currently doing. Before turning to the solution, note that the strategy 'P always' yields 0 payoffs and the strategy 'S always' yields

$$J_S = -C + \beta(1-p)\frac{\delta H}{1-\delta} + \beta\delta p\, J^C.$$

Note also that when $\beta < X$,

$$J_P = 0 > -C + \beta\delta p(0) + \beta(1-p)\frac{\delta H}{1-\delta},$$

that is, it is not worthwhile trying one-shot of disintoxication, if it is not followed by any other efforts. So:

Proposition 1. When $\beta < X$, 'P always' is an SMPE of the multi-selves game.

Note that if $\beta\delta J^C > J_S$, or equivalently:

$$\beta < Z \equiv X(1-\delta p)/1-\delta-\delta(1-p)X,$$

the individual prefers to procrastinate one period (P) than to play 'S always' because of the salience cost of stopping now. Since $X < Z$, this is the case when $\beta < X$. When, however, $\beta > Z$, then the strategy 'S always' is an SMPE and the individual makes the same choices as with time-consistent preferences. Finally, for $X < \beta < Z$, there is no pure strategy SMPE, but there exists a unique mixed strategy SMPE.

3.1. Pareto efficiency

Given the game-theoretic nature of the approach, there is no reason for the outcome to be efficient. In a multi-selves context, Pareto efficiency of an SMPE a^M with respect to the complete collection of selves requires that there is no sequence of actions $(a_0, a_1, \ldots)$ such that at all dates, the t-self prefers the sequence $(a_t, a_{t+1}, \ldots)$ to the SMPE. As long as $\beta > Y \equiv (X - p\delta X)/(1 - p\delta X)$, where $Y < X$, $J_S > J_P$ so that 'S always' Pareto dominates 'P always'. Hence:

Proposition 2. For $Y < \beta < X$, the SMPE is Pareto dominated by 'S always'.

The problem with the Pareto superior strategy 'S always' is that the option to procrastinate, 'wait one period, then stop', annihilates the credibility of the strategy 'S always'.

3.2. Private side bets and self-confidence

In his discussion of self-control, Ainslie (1992) refers to the concept of private side bets. The individual 'bets' on his behavior and looses the future reward of stopping if she decides to consume today. While this idea creates conceptual difficulties for a single individual, it fits with the way game theory links anticipations of future actions to current and past actions in history-dependent, or bootstrap, strategies: PE strategies may depend on past behavior at t, although past behavior is payoff-irrelevant, because behavior at $s > t$ also depends on past behavior.

Indeed, unlike the time consistent case, there is a multiplicity of PE of the multi-selves game, some of them having a natural interpretation. In our setting, consider the following strategy profile: for t-self, '*S if S has always been played before*; *P otherwise*'. When $Y < \beta < X$, it constitutes a PE that leads the individual to stop immediately![3] This outcome can be interpreted as a private side bet where the individual loses her internal bet if she deviates from a given plan of actions: she then gets low payoffs indefinitely as she will always make the easy choice in the future and lose any self-control benefit. Another way to look at it is to attach a psychological state of self-confidence to histories, interpreted as the degree of faith in the future ability to maintain a commitment not to consume, which could be annihilated if the individual ever consumes. Consumption is then perceived as involving a cost of $\beta\delta J^C$ under self-confidence.

3.3. Information acquisition and memory

Carillo and Mariotti (1997) and Brocas and Carillo (1999a,b) point out that the multi-selves approach implies that information, shared among the different selves, may have negative value; this is a consequence of the game-theoretical approach. An individual may then rationally refuse to collect some information even if it is free.

Suppose that H is unknown at date 0, with prior mean $H_0 = E(H)$, generating $Z_0 < \beta$, so that the agent would choose 'S always'. Suppose also that, at period 0, the agent has free access to a signal that leads to posterior H_1 with probability q_1, to posterior H_2 with probability q_2 and to no new information otherwise. If $Y_2 < \beta < X_2$, the individual will choose 'P always' under posterior H_2 while

[3] Deviating at t by consuming leads to 'P always' which yields a lower utility to t-self than the continuation 'S always'.

she would be better off choosing 'S always' (and 'S always' under H_1). Information leading to H_2 is thus self-defeating, the value of information is negative.

In a similar vein, cognitive dissonance, as developed by Festinger (1957), arises 'when cognitive elements (…) in an inconsistent relationship (…) create negative psychological tension, motivating the person to resolve the inconsistency so as to reduce the tension' (Pittman, 1997). Among the various mechanisms involved, such as changing preferences or manipulation of information, the individual can 'forget' self-defeating information. In such a case, there can be 'optimal' selective memory, eventhough the self knows that it is possible that she forgot selectively (Benabou and Tirole, 1999).

In the previous example, assume that the agent can decide to forget information H_2 and that she learned it. Because she does not know whether she learned and forgot H_2 or did not learn at all, her posterior when uninformed lies between H_2 and H_0. For q_1 and q_2 small, it is close to H_0: selectively erasing information H_2 leads the agent to the Pareto efficient strategy 'S always', and can be seen as a rational optimal mechanism.

4. Single self: Self-restraint

It has been however documented that individuals very often follow rules and principles. This leads to an alternative approach that attempts to restore the unity of the self and to consider the individual as one single self. An individual is considered as choosing a *plan* of action, as opposed to just an action now, on which she can expect to build upon not only today but in the future as well. Such a plan must then be based on logical grounds, taking into account the fact that the individual will face the same problem in the future.

Caillaud et al. (1999) formalizes this approach by letting the individual internalize logical inconsistencies caused by the possibility of her re-initializing a plan of action later on.[4] Hence, the individual disregards inconsistent plans, such as for example 'P today and S from tomorrow on', as they would be restarted again and again, or alternatively, as they can be seen as imposing a lower welfare for future selves although they face the same decision problem and could presumably make the same choice of plan. Conscious of her own unity, the individual takes this into account and therefore is assumed to restrict attention to internally consistent (IC) plans defined as follows. For any plan let J_t be the value evaluated with date-t preferences of continuing to follow the plan.

Definition 1. A plan is IC if $J_0 \leq J_t$ for all $t \geq 0$.

[4] For a related view based on cooperative game theory concepts, see Asheim (1997).

At any point in time continuing the plan is preferable to starting again the plan from its beginning (in a general stochastic context this should be true under any contingency that can occur with a positive probability given the initial plan). In our setting, IC plans have the following immediate property:

Lemma 1. If $\beta > Y$, for any IC plan, $J_0 \leq J_S$.

Now, consider date 0-self and her choice of an internal plan. The previous result shows that the strategy plan 'S always' is an acceptable choice. It is clearly IC and, as it delivers welfare equal to J_S, it (weakly) dominates any other IC plan available to the individual.

The strategy 'S always' is called a *self-restrained strategy.* Once it is planned, there is never an incentive to change it when the individual takes into account her unity and uses a logical mental process that makes her aware of the necessity to plan in an internally consistent way. There may exist other, more complicated, self-restrained strategies, but they are payoff-equivalent to the strategy 'S always'; more generally, they must be Pareto optimal in the multi-selves game.

This approach to self-control thus restores the unity of the self and results in a model of rational rules of behavior.

5. Psychological expected utility and time consistency

Psychologists have given ample evidence that the well-being of an individual not only depends upon physical consumption, but also upon various psychological elements of her state of mind. In an exponentially discounted utility framework, hence without postulating hyperbolic discounting, Caplin and Leahy (1998) proposes an axiomatic approach where anticipatory feelings about future uncertain events, such as anxiety, enter the utility function. This causes time-inconsistency since anxiety at t with respect to date $t+1$ decision disappears once date $t+1$ is reached.

To illustrate how the model generates time-inconsistent preferences, and keeping in mind that this is just an example, consider the previous setting with $\beta = 1$ and $J^C > 0$. Planning P tomorrow involves no anxiety today, while planning S tomorrow implies the resolution of a lottery at $t+1$ that causes a psychological state of mind characterized by anxiety today. Anxiety at t relates only to the next period uncertainty, but not to more distant dates. Let K denote the anxiety cost of planning S tomorrow (for simplicity assume it doesn't depend on the current action S or P).

The individual internalizes the future utility losses due to anxiety and maximizes expected utility at each date. In this framework, 'S always' yields intertemporal utility equal to $J_S^K = J^C - K/(1-\delta p)$. If K is large, this strategy may yield negative payoffs: anxiety is so costly compared to the benefit of non-addiction that it is not worthwhile getting anxious all the time by trying to stop. This strategy 'S always' is however an SPME, in a multi-selves approach, if

$$-C-K+(1-p)\frac{\delta H}{1-\delta}+p\delta J_{\mathrm{S}}^{K}>-K+\delta J_{\mathrm{S}}^{K},$$

which is always the case. If the individual expects S in the future, she reduces the probability to suffer from anxiety tomorrow by trying to stop today which raises her incentives to stop, compared to the case without anxiety. When K is large, it would be preferable to plan 'P always', but at each date, the individual prefers S since there is no anxiety cost for immediate actions and $C<(1-p)\,\delta H/(1-\delta)$.

As a slightly different point, the model exhibits also strict preferences on the timing of resolution of uncertainty, as in Kreps–Porteus theory, a phenomenon that is well-documented in the psychological literature. The individual tries to speed up the resolution of uncertainty by trying to stop today because it generates less anxiety. Note that while negative emotions like anxiety can explain the fact that individual may speed up costly and uncertain prospects, positive emotions can explain the delaying of some rewarding actions (and thus procrastination on rewarding acts), very much like a gourmet who may choose to keep the best of a meal to eat last.

6. Concluding remarks

A more fundamental approach to the formalization of time-inconsistent preferences is to use an axiomatic approach based solely on revealed preferences. Time-inconsistency is characterized by a positive value of commitment. Consequently, one can start from the revealed preferences of an individual over the restrictions of choices available to her in the future. Gul and Pesendorfer (1999) provides results to characterize the form of time-inconsistent preferences in a finite horizon environment. The complete axiomatic analysis remains to be done for the simple models laid in the present paper.

References

Akerlof, G., 1991. Procrastination and obedience, Richard T. Ely Lecture. American Economic Review 81 (2), 1–19.

Ainslie, G., 1992. Picoeconomics. Cambridge University Press, Cambridge.

Asheim, G., 1997. Individual and collective time-consistency. Review of Economics Studies 64, 427–443.

Becker, G., Murphy, K., 1988. A theory of rational addiction. Journal of Political Economy 96, 675–700.

Benabou, R., Tirole, J., 1999. Self-confidence: intrapersonal strategies. Working paper.

Boyer, M., 1983. Rational demand and expenditures patterns under habit formation. Journal of Economic Theory XXXVI, 27–53.

Brocas, I., Carillo, J., 1999a. Entry mistakes, entrepreneurial boldness and optimism. Working paper. ECARE, Brussels.

Brocas, I., Carillo, J., 1999b. On rush and procrastination. Working paper. ECARE, Brussels.

Caillaud, B., Cohen, D., Jullien, B., 1999. Towards a theory of self-restraint. Working paper.

Caplin, A., Leahy, J., 1998. Psychological expected utility and anticipatory feelings. Working paper, New York University.

Carillo, J., Mariotti, T., 1997. Strategic ignorance as a self-disciplining device. Working paper. ECARE, Brussels.

Elster, J., Loewenstein, G., 1992. Utility from memory and from anticipation. In: Loewenstein, G., Elster, J. (Eds.), Choice over Time. Russell Sage Foundation, New York.

Festinger, L., 1957. A Theory of Cognitive Dissonance. Row, Peterson, Evanston, IL.

Gilbert, D., Pinel, E., Wilson, T., Blumberg, S., Wheatley, T., 1998. Immune neglect: A source of durability bias in effective forecasting. Journal of Personality and Social Psychology 75, 617–638.

Gul, F., Pesendorfer, W., 1999. Dynamic inconsistency and self-control. Working paper.

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

Laibson, D., 1997. Golden eggs and hyperbolic discounting. Quarterly Journal of Economics 112 (2), 443–477.

Loewenstein, G., O'Donoghue, T., Rabin, M., 1999. Projection bias in predicting future utility. Working paper.

Lowenstein, G., Prelec, D., 1992. Anomalies in intertemporal choice: Evidence and an interpretation. Quarterly Journal of Economics 107 (2), 573–597.

O'Donoghue, T., Rabin, M., 2000a. Doing it now or later, American Economic Review, in preparation.

O'Donoghue, T., Rabin, M., 2000b. Incentives for procrastinators. Quarterly Journal of Economics, in preparation.

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

Pittman, T., 1997. Motivation. In: Gilbert, D., Fiske, S., Lindzey, G. (Eds.), The Handbook of Social Psychology. McGraw-Hill, New York.

Rabin, M., 1998. Psychology and economics. Journal of Economic Perspective XXXVI, 11–46.

Thaler, R., Shefrin, H., 1981. An economic theory of self-control. Journal of Political Economy 89, 392–406.

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