Scand. J. Economics 100(3), 643–662, 1998

Do Addicts Behave Rationally?*

Ernst Fehr

University of Zürich, CH-8006 Zürich, Switzerland

Peter K. Zych

Technical University of Vienna, A-1040 Vienna, Austria

Abstract

The theory of rational addiction assumes that addicts' behavior is fully rational. Common sense and psychological introspection suggest, however, that addictive behavior is irrational. Without knowledge of the addicts' preferences this dispute cannot be resolved. This paper reports the results of an experiment in which addictive preferences were induced. It turns out that 'addicts' consume systematically too much compared to the optimal consumption decision. We explain this systematic excess consumption in terms of the psychologically salient features of addictive goods.

I. Introduction

Most people seem to believe that addictive substances are ultimately harmful to the consumers. Moreover, even those who are addicted to certain substances often express the view that their addiction will ultimately hurt them and that, as a consequence, they would prefer to stop consuming or at least reduce their consumption of the addictive good. This divergence between behaviorally revealed preferences and verbally expressed preferences indicates the possibility that addicted people do not behave according to their best interests. Instead, their consumption choices may be affected by behavioral principles other than the maximization of their interests or there may be factors that limit their abilities to do so. A major competitor of any theory that views addiction as irrational behavior is, of course, the theory of rational addiction; see Becker and Murphy (1988). From the viewpoint of Becker and Murphy, addicts exhibit

* This paper was supported by a grant from the Austrian Science Foundation (no. 10136-soz). It has been presented at the IAREP/SABE conference on Integrating Views on Economic Behavior in Rotterdam, 1994, the Summer Institute on Behavioral Economics in Berkeley, 1994, the University of Bielefeld, 1994, the annual meeting of the Verein für Socialpolitik in Jena, 1994 and the 'Behavioral Economics' Seminar at Harvard University. We would like to thank the two referees and the participants for their helpful comments, in particuar Wulf Albers, Friedel Bolle, Colin Camerer, Simon Gächter, Werner Güth, Daniel Kahneman, Michael Kahneman, David Laibson, Matthew Rabin, Reinhard Selten, Dick Thaler and Carsten Vogt.

consistent, forward looking and individually optimal behavior. They know their preferences and the constraints they face and are fully capable of understanding the trade-off between present benefits and future costs.

The purpose of this paper is to directly test the theory of rational addiction. To achieve this purpose we used the induced value approach as developed by experimental economists; cf. e.g. Smith (1976). In our context the major advantage of this approach is that it allows tight control over subjects' true preferences. Therefore, we can directly compare the actual behavior of experimental subjects with the optimal behavior dictated by the experimentally induced preferences.

At present there seem to be two economic-psychological approaches which try to account for addictive behavior as irrationally myopic behavior. One approach is based on the notion of melioration, as in Herrnstein and Prelec (1992), while the other relies on the assumption of time-inconsistent preferences, as in e.g. Ainslie (1992). The melioration hypothesis stipulates that, in general, individuals tend to equalize the *temporary* average returns of different activities. In its pure form, the melioration hypothesis says that subjects neglect the intertemporal effects of present decisions completely. The melioration hypothesis is supported by the results of several animal experiments; for a review see Heyman and Herrnstein (1986). Other animal experiments, however, have reported ambiguous results; see Green *et al.* (1983). The evidence also seems to be ambiguous for human subjects; cf. Herrnstein *et al.* (1986, 1993). A central feature of most melioration experiments has been that subjects did *not* know their (experimentally induced) preferences. Instead, they had to learn their preferences by experiencing the rewards of their decisions in each period.

In our view, the assumption that people have to learn their preferences from the pleasure they experience in consuming different goods is plausible. Yet, it is equally plausible that under such circumstances subjects are unlikely to achieve the maximization outcome that corresponds to the full information situation. Since subjects face an additional constraint if they do not know their preferences and since this constraint is likely to be binding at least initially, proponents of the rational choice view could discount the results of those experiments in which subjects are close to the melioration outcome. There is then a danger that the answer regarding the rationality of addictive behavior would become a matter of belief. Those who adhere to the rational choice view of addiction *believe* that by and large people know their preferences while the others *believe* that people are severely constrained by incomplete information about their preferences.

To avoid such a conflict of beliefs, we decided to design our experiment on the terrain of rational choice theory. Our experimental subjects were fully informed about their (experimentally induced) preferences. In addition, to understand the intertemporal effects of their decisions they received explicit training before the start of the experiment. Therefore, if they deviated significantly and in a systematic way from the optimal consumption path, the results would cast doubts on a rational choice view of addiction.

Irrational consumption decisions may also arise if subjects have time-inconsistent preferences. This has been shown by Strotz (1956) and Yaari (1977), and more recently Laibson (1994).[1] Yet, to give the rational choice view the best chance we refrained from inducing inconsistent preferences. In our experiment, subjects had consistent preferences which were fully known to themselves. The experimental design favored the rational choice view in a further sense: although irrational excess consumption of addictive substances may have much to do with the uncertainties regarding when, how much and how long future costs arise we refrained from introducing such uncertainties.

In total, we eliminated three factors that may well cause irrational behavior in the context of the consumption of addictive substances. Preferences of experimental subjects were (i) consistent, (ii) fully known and (iii) not subject to any uncertainty. In addition, subjects were offered extensive and costless training and learning opportunities, which are described in more detail in Section III. Yet, although the experimental conditions were rather favorable for the confirmation of the rational choice view of addiction, subjects' behavior deviated significantly from the optimal consumption path. Moreover, this deviation is not irregular but very systematic. In almost all periods, subjects consumed too much relative to their true preferences.

The rest of the paper is organized as follows. In Section II we present the model which was implemented experimentally along with the numerical predictions for the optimal consumption path. Section III describes the experimental implementation in more detail and Section IV reports the results. Section V summarizes the results and provides an interpretation of the behavioral facts in terms of the psychologically salient features of the decision problem.

II. The Model

The basis for our experiment is the model of rational addiction as developed by Becker and Murphy (1988). We directly implemented a

[1] There exists rather convincing evidence that animals and human subjects frequently exhibit time-inconsistent behavior; see Ainslie (1975), Thaler (1981), Kagel and Green (1987) and Loewenstein and Thaler (1989).

version of this model in the laboratory. In discrete time our version of the Becker-Murphy (B&M) model has the following structure. The utility of an individual at any time t, u_t, depends on actual consumption c_t and the stock of consumption (addiction) capital S_t:

$$u_t = u(c_t, S_t). \tag{1}$$

Consumption c_t takes place at the beginning of the period and, hence, u_t also accrues at the beginning. The consumption stock S_t is built up by past consumption according to

$$S_t = (1-\delta) S_{t-1} + c_{t-1}, \tag{2}$$

where δ represents the rate of depreciation, while $S_{t\text{-}1}$ and $c_{t\text{-}1}$ denote consumption capital and consumption at t-1, respectively. If T is an individual's lifetime and σ her rate of time preference, lifetime utility as viewed from period T, U_T, is given by[2]

$$U_T = \sum_{t=1}^{T} (1+\sigma)^{T-t+1} u(c_t, S_t). \tag{3}$$

Let A_t be wealth at the beginning of period t, r the interest rate in a perfectly competitive capital market, p_c the price of the consumption good, and w the income per period. As in B&M, we assume $r = \sigma$. In any given period w can be expended on consumption $p_c c_t$, or it can be used to change the stock of wealth by $A_{t+1}/(1+r) - A_t$ which affects consumption opportunities in the future. Since we assume the existence of a perfectly competitive capital market, the intertemporal budget constraint is given by

$$p_c c_t + \frac{A_{t+1}}{1+r} - A_t \leq w, \qquad t = 1, \ldots, T. \tag{4}$$

In our context p_c, $r = \sigma$, w, S_1 and A_1 are exogenously given. A rational individual chooses a consumption sequence $(c_1, c_2, \ldots, c_T)$ to maximize U_T subject to (2) and (4) and the nonnegativity constraints $c_t \geq 0$ and $A_{T+1} \geq 0$.

By imposing further assumptions on the utility function (1) it is possible to capture important aspects of the effects of addictive substances. First of all, and perhaps most importantly, an increase in current consumption c_t

[2] In the experiment we induced a positive rate of time preference σ by paying interest on the amounts of money which accrued in periods $t \leq T$. This was done because we conjectured that it is easier for subjects to understand the concept of compound interest than the concept of discounting.

should give rise to future utility losses, that is, $\partial u_t/\partial S_t \equiv u_s < 0$. A second important assumption of the rational addiction model is $\partial^2 u_t/\partial c_t \partial S_t \equiv u_{cS} > 0$. Addiction is frequently characterized by so-called reinforcement effects, i.e., an increase in past consumption leads to an increase in present and future consumption. B&M show that $u_{cS} > 0$ is a necessary condition for reinforcement.

Numerical Implementation and the Optimal Consumption Path

In our experiment we implemented the following quadratic version of (1):

$$u(c_t, S_t) = a_0 + a_c c_t + \frac{a_{cc}}{2} c_t^2 + a_S S_t + \frac{a_{SS}}{2} S_t^2 + a_{cS} c_t S_t, \tag{5}$$

with $a_0 = 0$, $a_c = 0.6$, $a_{cc} = -0.01$, $a_S = -0.06$, $a_{SS} = -6 \times 10^{-6}$ and $a_{cS} = 6 \times 10^{-5}$. The parameters in (5) ensure that $u(c_t, S_t)$ is strictly concave and that within the relevant range all assumptions of B&M are met. Regarding the other parameters we implemented

$$\delta = 0.1,\ r = \sigma = 0.03,\ T = 30,\ w = 10,\ p_c = 1,\ A_1 = 0 \text{ and } S_1 = 0. \tag{6}$$

Except for the finite time horizon no particular intentions are behind the parameters in (6). It would have been possible to implement an infinite but random time horizon. Yet, apart from the problem of how to make subjects believe that the length of the experiment is randomly selected, there is the difficulty that experimental subjects frequently make the following mistake.[3] If the probability of ending is, let's say, 0.1, they believe that the experiment will end in period 10. And if it does not end in period 10 they are either confused or believe strongly that it must end in the near future. In the context of our experiment, a random horizon would have favored the excess consumption hypothesis. If subjects believe that the experiment will end soon, they will tend to use up their wealth for consumption purposes because — due to the expectation of a small number of remaining periods — expected future costs are curtailed.

With a finite horizon, consumption of addictive substances will also be shifted towards the end because of the smaller number of remaining periods in which costs can accrue. Yet, within a finite horizon framework,

[3] J. D. Hey brought this problem to our attention. He has conducted random horizon experiments in which this phenomenon occurred.

such a decision is rational. For our parameters the optimal consumption path is given by

$$c_t = -0.025 \times 0.904^t + 0.932 \times 1.140^t + 0.182. \qquad (7)$$

This path is monotonically increasing (see also Figure 1).[4]

The Conditionally Optimal Consumption Path

In the theory of rational addiction subjects do not deviate from the optimal path. Therefore, only the optimal consumption path is of interest and relevance. Yet, in an experiment we cannot rule out that subjects choose nonoptimal consumption levels. In case that an individual deviates, for whatever reason, from the optimal consumption path in some period $t \in [1, \tau]$, her stock of consumption capital $S_{\tau+1}$ and her wealth $A_{\tau+1}$ are affected. As a consequence, her optimal consumption path from period $\tau+1$ onwards is also affected. We call the 'new' optimal consumption path 'conditionally optimal' because it is optimal for the actually given values of $S_{\tau+1}$ and $A_{\tau+1}$. If an individual has consumed too much relative to the optimal consumption path in $t \in [1, \tau]$, $S_{\tau+1}$ is larger and $A_{\tau+1}$ is smaller compared to the optimal stock levels. As a consequence, the optimal consumption level $c_{\tau+1}$ conditional on $S_{\tau+1}$ and $A_{\tau+1}$ is lower than the optimal level in the absence of past errors. Since it is not impossible that subjects learn how to make optimal consumption decisions during the course of a 30-period experiment, the relevant standard of comparison is the conditionally optimal level c_{co}.

III. The Experimental Design

Structure

In our experiment we implemented the model described in the preceding section in a computerized laboratory. To avoid a loss of control over subjects' preferences we did not frame the experiment as an addiction problem. Instead it was framed as a problem of buying goods in different time periods with a limited amount of resources. By buying goods subjects

[4] Equation (7) gives the optimal path if c_t can be any nonnegative real number. In the experiment c_t was restricted to nonnegative integers. This implies only a negligible efficiency loss of less than one tenth of a percent of the maximum utility in the absence of the integer constraint.

could earn experimental money which was transformed into real money at the end of the experiment. The endowment w and wealth A_t were expressed in points, consumption c_t and addiction capital S_t in units of goods, whereas utility $u(c_t)$ accrued in terms of experimental money. Experimental money was called Guilders and the exchange rate between real money (= Austrian Schillings) and Guilders was 1:1. This exchange rate was public information. In each period subjects received an endowment of $w = 10$ points. They could use these points for the purchase of goods, that is, for consumption purposes ($c_t > 0$) and/or they could save them ($c_t < 10$). Positive consumption ($c_t > 0$) led to a gain in terms of Guilders according to the prespecified utility function (5). When subjects had made their consumption decision in a given period, the gain in terms of Guilders was added to a Guilders account. Paying three percent interest on this account in each period induced the time preference $\sigma = 0.03$.

If – in a given period – subjects consumed less than $w = 10$, the difference $(w - c_t)$ was added to a savings account[5] which also earned interest of $r = 0.03$. If, on the other hand, $(w - c_t)$ was negative, subjects dissaved or, in case the savings account became negative, raised a loan. The savings account (credit account) mimicked a perfect capital market. In any given period subjects had the opportunity to use up their whole discounted future endowments plus their savings account plus their present endowment for consumption. In each period the discounted value of total resources available for consumption in that period was shown on the screen.

Subjects made their decisions by choosing among the numbers of a consumption schedule on their computer screen (see Screen A1 in the Appendix). The consumption schedule was supplemented by a Guilder schedule that indicated how many Guilders (utility) were associated with each consumption level. Thus, in each period, subjects had the information about the temporal utility function that was relevant in that period, $u(c_t, \cdot)$, on the (right hand side of the) computer screen. Before making a decision they could scroll this utility schedule up and down as they liked.

Subjects' consumption decisions and the depreciation of addiction capital led to changes in S_t. Hence, in general the temporal utility function $u(c_t, \cdot)$ changes from one period to the next. While the computer automatically informed subjects about their relevant temporal utility functions, information about the total utility function $u(c_t, S_t)$ was given to subjects by a large matrix (see Table A1 in the Appendix). In this matrix, which was called 'transformation scheme' in the experiment, S_t is termed 'current

[5] Since the term 'saving' might evoke positive (or negative) connotations, we called the savings account 'point account' in the experiment.

stock of goods'. For discrete levels of the 'current stock of goods' S_t the matrix displayed the associated temporal utility function. For example, if a subject had accumulated an addiction capital of $S = 60$ and consumed 18 units, the temporal utility gain was given by 5.6 Guilders.

The transformation scheme provided subjects – for the relevant S-interval – with the information about (5). As one can see, it is very easy to calculate temporal utility gains from c_t whereas the total amount of discounted utility losses which accrues in future periods due to the present consumption level is much more difficult to grasp.

Our design ensures that the experiment is fully isomorphic to the model presented in Section II. The maximization of compounded (or discounted) income is tantamount to the maximization of the intertemporal utility function (5). The main incentive for subjects to participate was our promise that they could earn money in a decision-making experiment. Therefore, it is safe to assume that they were primarily interested in the maximization of Guilders. This presumption is reinforced by two additional considerations. Subjects were undergraduate students of business administration for whom the maximization of monetary returns is a rather familiar and subjectively approved concept. Moreover, they were challenged by the problem per se, that is, they took their Guilders score as an indication of how well they had solved the problem.

It is worthwhile to stress that from the perspective of testing the theory of rational addiction and of investigating the link between true preferences and actual behavior, this motivation is highly desirable because it allows an isomorphic implementation of the rational addiction model and an unambiguous determination of the optimal consumption sequence. If, instead, subjects would have had a different motivation, we would have lost control over subjects' preferences and, as a consequence, we would have been unable to determine the optimal and the conditionally optimal consumption paths. Of course, without knowledge of these paths, it is no longer possible to decide whether a particular consumption sequence exhibits myopic excess consumption.

Thus, there are obvious benefits if money earnings are the sole motivator. On the other hand, several authors have argued that addiction is characterized by motivational conflicts ('multiple selves') *within* the person and that these conflicts give rise to excess consumption. In our experiment such conflicts did not arise because subjects had a clear, one-dimensional motivation. In our view, however, the absence of motivational conflicts is an advantage because it allows us to find out whether there are other relevant factors, such as cognitive limitations, which contribute to excess consumption. If, instead, we had introduced motivational conflicts (e.g. by inducing time-inconsistent preferences) the effects of these other factors would have been confounded.

Learning and Training Opportunities

Before the start of the experiment subjects received written instructions.[6] Besides a description of the decision problem, the instructions contained many examples. The purpose of these examples was to inform subjects as well as possible about the decision opportunities and consequences of alternative decisions. In particular we stressed the future costs of current consumption decisions. When we introduced the transformation scheme (utility function) in the instructions we wrote:

> 'The more goods you buy in the current period the higher will be the stock of goods in future periods and the smaller are the Guilder receipts you can earn in future periods from the purchase of the same amount of goods.'

This remark was then illustrated by *examples* which showed that the purchase of a given amount of goods creates lower Guilder earnings at higher stocks of the good. When subjects had finished reading the instructions they had to solve several *exercises*.[7] The purpose of these exercises was again to train their comprehension of the intertemporal consequences of current decisions. No subject could start with the experiment until all exercises had been solved correctly.[8]

At the beginning of each period, before the consumption decision was made, subjects had been informed about the current stocks on (i) the savings account and (ii) the Guilder account. In addition, the computer screen (see Screen A1 in the Appendix) showed them (iii) the total resources available for consumption and (iv) the current stock of addiction capital ('current stock of goods'). When a consumption decision had been made, the computer showed subjects explicitly on a new screen how this decision affected these four variables (see Screen A2 in the Appendix). Subjects then had the possibility to confirm their decision or to change it. Only in case of confirmation was their decision implemented. In principle, subjects could go back and forth in a given period as often as they wanted. They could make as many preliminary decisions as they liked and examine their consequences without implementing the decisions.

Before the start of the 30-period decision problem we allowed for two further training opportunities. There was a 3-period decision problem and a 30-minute training round. After the 30-minute training round, two

[6] Experimental instructions are available on request.
[7] Subjects used calculators to solve the exercises as well as during the experiment.
[8] To avoid embarrassment in the case of mistakes, we required subjects to write down the whole computation process. If a subject made a mistake we indicated the error by encircling the numbers in the computation which were wrong. Correct computations were marked differently.

experimental rounds, which we called ‘main rounds’, took place. Both in the 3-period problem and during the training round we implemented the same parameters as in the subsequent main rounds. Contrary to the main rounds, however, there was no money at stake. Except for effort and time costs training was, therefore, costless. The intention of the 3-period problem was to make subjects familiar with the computer. They could learn how to make and revise their decisions on the PC. The 30-minute round, instead, gave them the opportunity to experiment with different consumption paths. They could run as many 30-period problems as they liked during this training round.[9]

After the training round the first main round began. Subjects did not know that there was a second main round. They were informed about it only after they had finished the first main round. After the second main round subjects had to describe their strategies in a questionnaire. In addition, they were asked about what decisions they would implement if they had another opportunity to solve the same 30-period problem. They answered this question by drawing a consumption path into a table with a fine grid. In the following this hypothetical consumption path is dubbed ‘main round three’.[10] Using data from three consecutive rounds enables us to examine whether learning processes push consumption decisions closer to the optimum or whether excess consumption is a persistent phenomenon.

IV. Results[11]

Twenty-five undergraduate students of business administration from the University of Vienna participated in this experiment which lasted approximately two hours and 15 minutes.[12] Most subjects were between 19 and 21 years old. Roughly one third of the subjects were female, two thirds were

[9] In addition, once started, they did not have to finish a 30-period problem. They could stop at any period and start the 30-period problem anew. Subjects also had the opportunity to end the whole training round before 30 minutes were over. Yet, they made extensive use of training opportunities. All of them made complete use of the available 30 minutes. During the training round, most subjects tried to find the best possible consumption path by trial and error. Explicit computation of the optimal path was not possible because the subjects had no knowledge of dynamic optimization techniques.

[10] In general, subjects drew a consumption path that did not obey the intertemporal budget constraint. To meet this constraint we rescaled each individual’s path by proportionally increasing or decredasing consumption levels in all periods. Our results below hold irrespective of whether we include the third main round in the statistical analysis.

[11] Our data are available on request.

[12] During the first 45–60 minutes, subjects read the instructions, solved the exercises and experimented with different consumption paths in the training round. The remaining 75–90 minutes were spent on consumption decisions in the two main rounds.

male students. On average subjects earned 245 Austrian Schillings (ATS).[13] The lowest income achieved was ATS 175 while the highest income was ATS 288. This indicates that there were substantial differences in individual performance. The nature of the problem and the monetary incentives gave rise to highly motivated subjects.

Actual and Optimal Consumption Paths

RESULT 1: *Actual consumption paths are characterized by suboptimal excess consumption in early periods and suboptimal underconsumption during final periods.*

Figure 1 shows the average actual consumption path taken over all three main rounds c_{all} in comparison with the optimal path. Until period 15, subjects consume on average too much and from period 19 onwards they consume too little. It is worthwhile to stress that this pattern occurs not only when taking the average over all three main rounds. The same pattern can be observed in each of the three main rounds. T-tests that compare average consumption in main round one (c^1), main round two (c^2) and main round three (c^3) for each period indicate that there are no differences between the main rounds. Pairwise comparisons of c^1, c^2 and c^3 reveal that the null hypothesis of equal average consumption levels cannot be rejected at the 5 percent level for any period in the interval $t \in [1, 28]$. Only in period 29 does c^1 differ slightly from c^2 and in period 30 c^1 differs somewhat from c^3. This similarity between all main rounds indicates that the forces which give rise to the deviation from the optimal path are persistent, that is, they are not removed over time by subjects' learning processes. It seems that no significant learning took place after the end of the 30-minute training round. Our instructions and the many learning opportunities offered by our design seem to have allowed subjects to learn what is within their cognitive limits.

Figure 1 provides a misleading picture of the deviations from the optimal path, as it does not take into account that the optimal path changes as a result of past errors. The essential question, therefore, is whether excess consumption is observed relative to the conditionally optimal consumption path c_{co}.

Actual and Conditionally Optimal Consumption Path

RESULT 2: *Actual average consumption c_{all} is characterized by permanent but eventually declining excess consumption relative to the conditionally*

[13] This number also includes the show-up fee of ATS 70 (ATS 100 ≈ US$ 10).

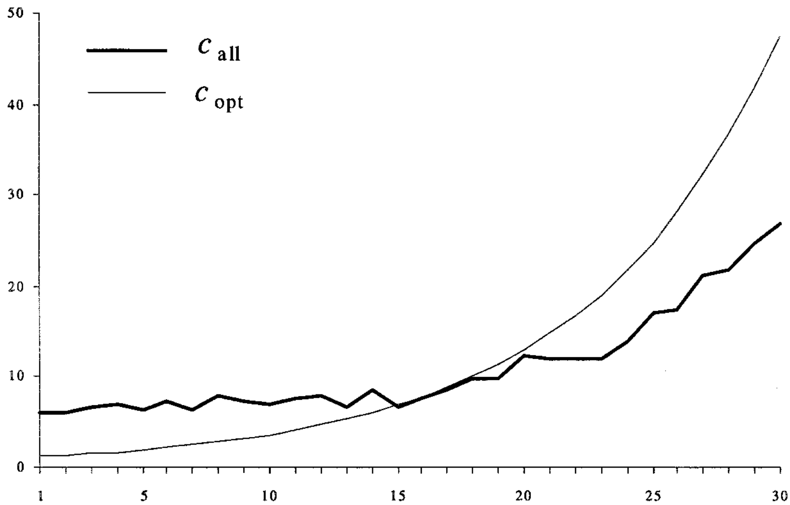

Fig. 1. Average of actual paths for all main rounds and optimal path

optimal consumption path. In the final and the next to final period c_{all} reaches the conditionally optimal path.

Figure 2 provides support for R2. In Figure 2 we see the difference between c_{all} and the average conditionally optimal path c_{co} together with

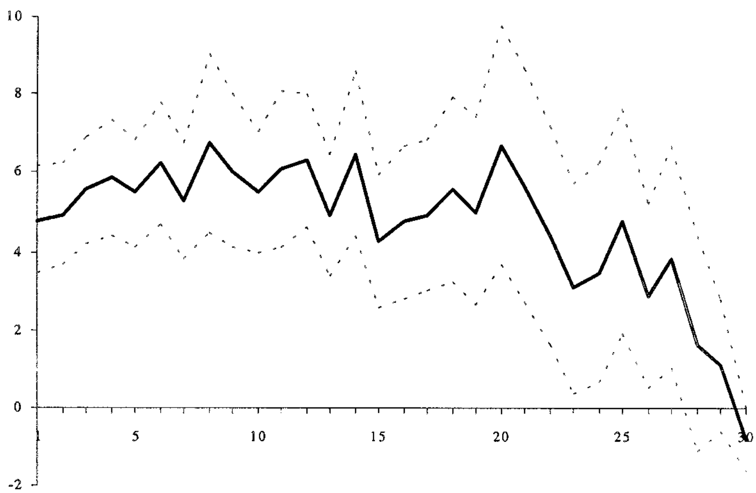

Fig. 2. Difference between c_{all} and c_{co} with 95% confidence interval

the 95% confidence interval.[14] The figure reveals that excess consumption has several ups and downs between period 1 and period 20, but in general there is neither an increasing nor a decreasing tendency. However, from period 20 onwards, excess consumption gradually approaches zero, although again with several ups and downs. From period 28 onwards we can no longer reject the hypothesis that subjects were along the conditionally optimal consumption path.[15] Thus, R2 is well supported by the data.

How much suboptimality in terms of foregone overall earnings did excess consumption cause? To measure the degree of inefficiency we define

$$\eta = \frac{U_T^{\max} - U_T}{U_T^{\max}}, \tag{8}$$

where $U_T^{\max}$ denotes the maximum overall payoff. η tells us by how many percent actual overall earnings fall short of maximum overall earnings. The average η is given by 27.9 percent. Thus, although each consumption decision may only cause small overall losses, the sum of all losses due to suboptimal consumption decisions was, on average, quite substantial.

Evaluation of Questionnaires

After the experiment we asked subjects on what basis they had made their decisions and whether they changed their strategies between round 1 and round 2. We already know that there was no significant behavioral difference between these rounds. This result is reinforced by the answers to the questionnaire. Most subjects said that they tried to implement roughly the same pattern of consumption decisions in both rounds. Yet, within the limits of the globally fixed strategy, they tried to achieve some fine-tuning. According to subjects' answers we can identify several broad patterns:

A: To consume just the endowment with maximal deviations of ± 10.
B: To consume very little at the beginning such that there is an accumulation of wealth A_t. This wealth is used towards the end for very high consumption levels.

---

[14] To get c_{co} we computed the conditionally optimal path for each individual in each period and each session. We then took – for any period between 1 and 30 – the average over each individual and each session. Thus one point represents the average over 75 values.
[15] We also conducted a nonparametric chi-square test of the null hypothesis that excess consumption is equally likely as underconsumption in a given period. Except for period 30 – where, by definition, there is no excess consumption – we can reject the null hypothesis in favor of the alternative hypothesis that excess consumption is more likely ($p < 0.002$).

C: Start with relatively little consumption and increase c_t gradually.
D: A more or less cyclical consumption pattern. High consumption at a particular period is followed by several periods of rather low consumption in order to allow for a depreciation of addiction capital. After that there is again a period of high consumption, etc.
E: Compute the maximal temporary payoff per unit of consumption $u(c_t)/c_t$. Since the maximization of $u(c_t)/c_t$ implies rapidly increasing consumption levels subjects who intentionally tried to implement this behavioral pattern over some time either saved a lot during initial periods or implemented a cyclical consumption pattern with an approximate cycle length of four periods.

In Table 1 we show how many subjects can be associated with each behavioral pattern. In addition we report the average degree of inefficiency, η, for each group. As one can see, behavioral pattern C, which qualitatively resembles the optimal consumption path, exhibits the lowest degree of inefficiency (14.7%). Yet, only five subjects deliberately followed this strategy. Other subjects intended to implement strategies which do not even qualitatively resemble the optimal path. As a consequence, their consumption sequences caused considerably higher degrees of suboptimality.

V. Interpretation and Potential Policy Implications

The data of our experiment show that actual consumption paths significantly deviate from the optimal path as well as from the conditionally optimal path. Deviations are in no way irregular but follow a systematic pattern. In relation to the conditionally optimal path, subjects consume too much in almost all periods. During the first 20 periods, excess consumption has neither an increasing nor a decreasing trend but after period 20 it tends to decrease until it vanishes in the last three periods.

How can we interpret this behavioral pattern? What are the factors that contribute to systematic and persistent deviation from the optimal consumption path? The starting point for our interpretation of the data is the assumption that, in general, people have cognitive limitations.[16] In some circumstances these limitations represent binding constraints, while in other circumstances they have no effects on behavior. In the context of our experiment they are likely to be binding because it is a rather difficult task to compute the present value of future per-period costs and aggregate it correctly. Yet, if subjects are unable to compute the aggregate present value of future costs which is associated with each level of c_t they face a form of *subjective* uncertainty. Although objectively future costs are given with certainty, subjects' cognitive limitations generate a situation that is characterized by subjective uncertainty. If they consume, say, $\tilde{c}_t$, does the true aggregate present value of future costs exceed or fall short of the present benefits of $\tilde{c}_t$? We hypothesize that in the face of this subjective uncertainty about what subjects should do, the psychologically relevant features of the decision problem play an important role. In the context of our intertemporal decision problem the fact that the present benefits of c_t are relatively large, immediately available and, thus, unambiguously given is likely to be of psychological relevance because it renders present benefits highly salient. In contrast, the fact that future costs per period are relatively small, distributed over time and, thus, subjectively of ambiguous size renders them much less salient. In the face of subjective uncertainty, the combination of these psychological features of the decision problem is likely to create an upwards bias of subjects' perception of present benefits in relation to the aggregate present value of future costs.[17]

The above reasoning provides a plausible explanation of persistent excess consumption during the first 20 periods. However, how can we account for the decrease in excess consumption during the last ten periods? In our view it is plausible that the drive towards excess consumption loses its force towards the end because costs are distributed over a smaller number of future periods and are, therefore, more easily aggregated to a correct number. Moreover, in the final period and in the next to final period it is relatively easy to find the optimal decision because there is no longer any distributivity of future costs. In the final period this is trivially true because there are no future costs, whereas the decision in the next to final period can be based on *one* simple cost item: the decrease in u_T that is due to c_{T-1}. There is thus a plausible explanation of the behavioral facts if cognitive limitations and the psychologically salient features of the decision problem are taken into account.

---

[16] For an early analysis of the impact of cognitive limitations on individual and organizational behavior, see Simon (1957).

[17] Note that the distributivity of future costs enters our argument twice. First, combined with subjects' cognitive limitations, it creates a state of subjective uncertainty. Second, it affects subjects' perception of present benefits and future costs.

In our view the message of our results is twofold. Since we have conducted an isomorphic test of the theory of rational addiction and since the predictions are not met, the data cast doubts on the validity of this theory. In addition, we think that we have identified a psychological factor that contributes to the existence of irrationally myopic excess consumption. Although we offered subjects a host of training and learning opportunities which were used readily, although subjects were highly motivated and although they were highly skilled relative to the average skills of the population, their choices reveal a systematic undervaluation of future costs. It seems that their cognitive abilities did not allow them to integrate distributed future costs into one correct number. This interpretation is also supported by the fact that after period 20 subjects slowly approached the conditionally optimal path. Since the number of cost items which has to be taken into account declines as one approaches the last period, it becomes easier to integrate future costs into a correct figure. Therefore, the cognitive limitations of subjects have less impact towards the end.

Does this tell us anything about real-world addictions or real-world consumption patterns? Well, if cognitive limits exert this sort of impact in the laboratory, why should they be inoperative in the real world? Why should an addict be less prone to cognitively undervalue distributed future costs than an undergraduate student of business administration? If anything, the opposite seems to be more likely. For some addictions it may even be the case that the consumption of the addictive substance inhibits the cognitive abilities of the addict. Therefore, if undervaluation of future costs is really driven by cognitive limits the tendency to undervalue may even increase over time.

The relevance of our results is not necessarily restricted to the case of addictive behavior in a narrow sense. The starting point of our analysis is that addictions obey certain qualitative hedonic regularities like saliency of present benefits and distributivity of future costs. However, this does not mean that these regularities are completely absent in other consumption activities. In fact, addictive behavior may be defined as those activities in which such hedonic regularities are particularly strong. But to the extent that such regularities are also present in the consumption process of other goods, our results will also have some relevance for these goods. Our results may be particularly relevant for habit-forming goods because addiction can be viewed as a special case of habit formation. Take for example the case where an activity involves high present costs while the benefits are distributed over many future periods.[18] If we apply our reasoning to this situation, we would predict that individuals underinvest in such activities.

[18] The acquisition of human capital and many health activities seem to be characterized by these hedonic regularities.

This might provide a justification for public interventions that alleviate this underinvestment.

The previous argument leads us directly to the potential policy implications of our results. In our experiment too many resources were used for consumption in early periods. In reality, however, it is also possible that overconsumption of one good is financed by the suboptimal underconsumption of other goods. Rationality failures may, therefore, generate suboptimal consumption structures. It may be that those goods, which exhibit the hedonic regularities of what we have called addictive goods, are overconsumed at the expense of those goods where these hedonic regularities are absent. By means of differential goods taxation it could be possible to correct for these efficiency-reducing distortions. In this way 'sin taxes' may be justified on efficiency grounds even in the absence of any externalities.

Appendix

Due to space limitations we do not present the whole set of instructions in this appendix. Instead we only show the decision screen, the confirmation screen, and the utility matrix (transformation scheme) of the experiment.

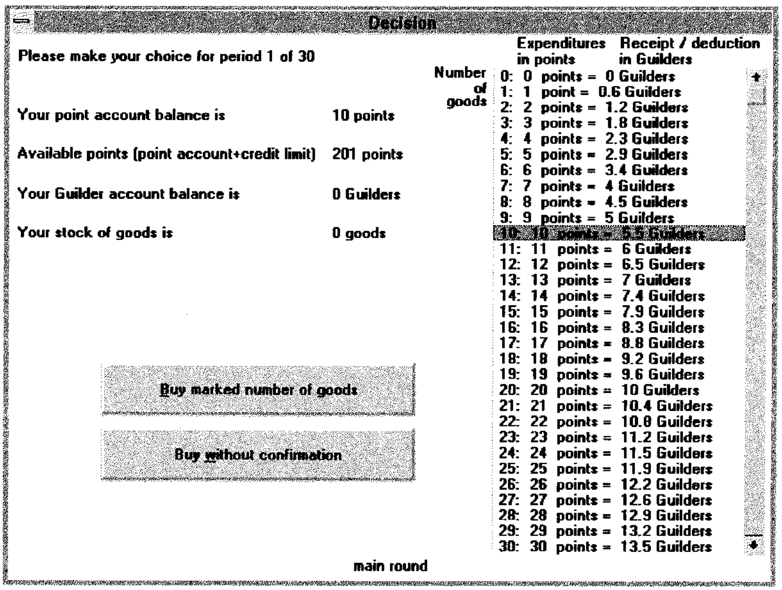

Screen A1. Decision screen

Table A1. *Transformation scheme*

No. of goods bought in the current period	Current stock of goods															
	0	10	20	30	40	50	60	70	80	90	100	125	150	200	250	300
0	0.0	−0.6	−1.2	−1.8	−2.4	−3.0	−3.6	−4.2	−4.8	−5.4	−6.0	−7.5	−9.1	−12.1	−15.2	−18.3
1	0.6	0.0	−0.6	−1.2	−1.8	−2.4	−3.0	−3.6	−4.2	−4.8	−5.4	−6.9	−8.5	−11.5	−14.6	−17.7
2	1.2	0.6	0.0	−0.6	−1.2	−1.8	−2.4	−3.0	−3.6	−4.2	−4.8	−6.4	−7.9	−10.9	−14.0	−17.1
3	1.8	1.2	0.6	0.0	−0.6	−1.2	−1.8	−2.4	−3.0	−3.7	−4.3	−5.8	−7.3	−10.3	−13.4	−16.5
4	2.3	1.7	1.1	0.5	−0.1	−0.7	−1.3	−1.9	−2.5	−3.1	−3.7	−5.2	−6.7	−9.8	−12.8	−15.9
5	2.9	2.3	1.7	1.1	0.5	−0.1	−0.7	−1.3	−1.9	−2.5	−3.1	−4.6	−6.1	−9.2	−12.2	−15.3
6	3.4	2.8	2.2	1.6	1.0	0.4	−0.2	−0.8	−1.4	−2.0	−2.6	−4.1	−5.6	−8.6	−11.7	−14.7
7	4.0	3.4	2.8	2.2	1.6	1.0	0.4	−0.2	−0.8	−1.4	−2.0	−3.5	−5.0	−8.1	−11.1	−14.2
8	4.5	3.9	3.3	2.7	2.1	1.5	0.9	0.3	−0.3	−0.9	−1.5	−3.0	−4.5	−7.5	−10.6	−13.8
9	5.0	4.4	3.8	3.2	2.6	2.0	1.4	0.8	0.2	−0.4	−1.0	−2.5	−4.0	−7.0	−10.1	−13.1
10	5.5	4.9	4.3	3.7	3.1	2.5	1.9	1.3	0.7	0.1	−0.5	−2.0	−3.5	−6.5	−9.5	−12.6
11	6.0	5.4	4.8	4.2	3.6	3.0	2.4	1.8	1.2	0.6	0.0	−1.5	−3.0	−6.0	−9.0	−12.1
12	6.5	5.9	5.3	4.7	4.1	3.5	2.9	2.3	1.7	1.1	0.5	−1.0	−2.5	−5.5	−8.5	−11.6
13	7.0	6.4	5.8	5.2	4.6	4.0	3.4	2.8	2.2	1.6	1.0	−0.5	−2.0	−5.0	−8.0	−11.1
14	7.4	6.8	6.2	5.6	5.0	4.5	3.9	3.3	2.7	2.1	1.5	0.0	−1.5	−4.5	−7.6	−10.6
15	7.9	7.3	6.7	6.1	5.5	4.9	4.3	3.7	3.1	2.5	1.9	0.4	−1.1	−4.1	−7.1	−10.1
16	8.3	7.7	7.1	6.5	6.0	5.4	4.8	4.2	3.6	3.0	2.4	0.9	−0.6	−3.6	−6.6	−9.7
17	8.8	8.2	7.6	7.0	6.4	5.8	5.2	4.6	4.0	3.4	2.8	1.3	−0.2	−3.2	−6.2	−9.2
18	9.2	8.6	8.0	7.4	6.8	6.2	5.6	5.0	4.4	3.9	3.3	1.8	0.3	−2.7	−5.7	−8.8
19	9.6	9.0	8.4	7.8	7.2	6.6	6.1	5.5	4.9	4.3	3.7	2.2	0.7	−2.3	−5.3	−8.3
20	10.0	9.4	8.8	8.2	7.6	7.1	6.5	5.9	5.3	4.7	4.1	2.6	1.1	−1.9	−4.9	−7.9
21	10.4	9.8	9.2	8.6	8.0	7.5	6.9	6.3	5.7	5.1	4.5	3.0	1.5	−1.5	−4.5	−7.5
22	10.8	10.2	9.6	9.0	8.4	7.8	7.2	6.7	6.1	5.5	4.9	3.4	1.9	−1.1	−4.1	−7.1
23	11.2	10.6	10.0	9.4	8.8	8.2	7.6	7.0	6.4	5.9	5.3	3.8	2.3	−0.7	−3.7	−6.7
24	11.5	10.9	10.3	9.8	9.2	8.6	8.0	7.4	6.8	6.2	5.6	4.2	2.7	−0.3	−3.3	−6.3
25	11.9	11.3	10.7	10.1	9.5	8.9	8.4	7.8	7.2	6.6	6.0	4.5	3.0	0.1	−2.9	−5.9
26	12.2	11.6	11.1	10.5	9.9	9.3	8.7	8.1	7.5	6.9	6.3	4.9	3.4	0.4	−2.6	−5.6
27	12.6	12.0	11.4	10.8	10.2	9.6	9.0	8.5	7.9	7.3	6.7	5.2	3.7	0.8	−2.2	−5.2
28	12.9	12.3	11.7	11.1	10.5	10.0	9.4	8.8	8.2	7.6	7.0	5.5	4.1	1.1	−1.9	−4.9
29	13.2	12.6	12.0	11.4	10.9	10.3	9.7	9.1	8.5	7.9	7.3	5.9	4.4	1.4	−1.6	−4.6
30	13.5	12.9	12.3	11.8	11.2	10.6	10.0	9.4	8.8	8.2	7.7	6.2	4.7	1.7	−1.2	−4.2
31	13.8	13.2	12.6	12.0	11.5	10.9	10.3	9.7	9.1	8.5	8.0	6.5	5.0	2.0	−0.9	−3.9
32	14.1	13.5	12.9	12.3	11.8	11.2	10.6	10.0	9.4	8.8	8.2	6.8	5.3	2.3	−0.6	−3.6
33	14.4	13.8	13.2	12.6	12.0	11.4	10.9	10.3	9.7	9.1	8.5	7.1	5.6	2.6	−0.3	−3.3
34	14.6	14.0	13.5	12.9	12.3	11.7	11.1	10.5	10.0	9.4	8.8	7.3	5.9	2.9	−0.1	−3.0
35	14.9	14.3	13.7	13.1	12.6	12.0	11.4	10.8	10.2	9.6	9.1	7.6	6.1	3.2	0.2	−2.8
40	16.0	15.4	14.8	14.3	13.7	13.1	12.5	12.0	11.4	10.8	10.2	8.8	7.3	4.4	1.4	−1.6
45	16.9	16.3	15.7	15.2	14.6	14.0	13.4	12.8	12.3	11.7	11.1	9.7	8.2	5.3	2.4	−0.6
50	17.5	16.9	16.4	15.8	15.2	14.6	14.1	13.5	12.9	12.3	11.8	10.3	8.9	6.0	3.1	0.1
55	17.9	17.3	16.7	16.2	15.6	15.0	14.5	13.9	13.3	12.7	12.2	10.7	9.3	6.4	3.5	0.6
60	18.0	17.4	16.9	16.3	15.7	15.2	14.6	14.0	13.5	12.9	12.3	10.9	9.5	6.6	3.7	0.8
65	17.9	17.3	16.8	16.2	15.6	15.1	14.5	13.9	13.4	12.8	12.2	10.8	9.4	6.5	3.7	0.8
70	17.5	16.9	16.4	15.8	15.3	14.7	14.1	13.6	13.0	12.5	11.9	10.5	9.1	6.2	3.4	0.5
75	16.9	16.3	15.8	15.2	14.7	14.1	13.5	13.0	12.4	11.9	11.3	9.9	8.5	5.7	2.8	0.0
80	16.0	15.4	14.9	14.3	13.8	13.2	12.7	12.1	11.6	11.0	10.5	9.1	7.7	4.8	2.0	−0.8
85	14.9	14.3	13.8	13.2	12.7	12.1	11.6	11.0	10.5	9.9	9.4	8.0	6.6	3.8	1.0	−1.9
90	13.5	13.0	12.4	11.9	11.3	10.8	10.2	9.7	9.1	8.6	8.0	6.6	5.2	2.5	−0.3	−3.2
95	11.9	11.3	10.8	10.2	9.7	9.2	8.6	8.1	7.5	7.0	6.4	5.0	3.7	0.9	−1.9	−4.7
100	10.0	9.5	8.9	8.4	7.8	7.3	6.7	6.2	5.7	5.1	4.6	3.2	1.8	−0.9	−3.7	−6.5
105	7.9	7.3	6.8	6.3	5.7	5.2	4.6	4.1	3.6	3.0	2.5	1.1	−0.2	−3.0	−5.7	−8.5
110	5.5	5.0	4.4	3.9	3.4	2.8	2.3	1.7	1.2	0.7	0.1	−1.2	−2.6	−5.3	−8.0	−10.8
115	2.9	2.3	1.8	1.3	0.7	0.2	−0.3	−0.9	−1.4	−1.9	−2.5	−3.8	−5.2	−7.9	−10.6	−13.3
120	0.0	−0.5	−1.1	−1.6	−2.1	−2.6	−3.2	−3.7	−4.2	−4.8	−5.3	−6.6	−8.0	−10.7	−13.4	−16.1

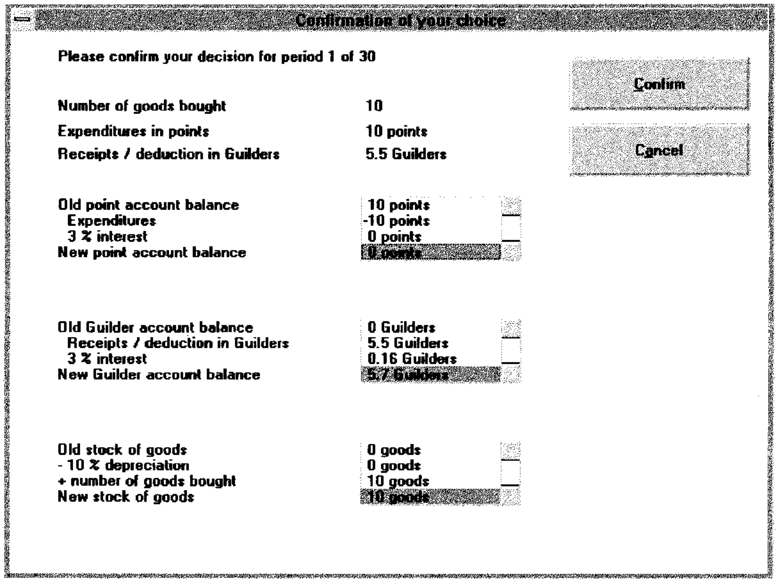

Screen A2. Confirmation screen

References

Ainslie, G.: Specious reward: A behavioral theory of impulsiveness and impulse control. *Psychological Bulletin 82*, 463–509, 1975.

Ainslie, G.: *Picoeconomics. The strategic interaction of successive motivational states within the person*. Cambridge University Press, Cambridge, 1992.

Becker, G.S. and Murphy, K.M.: A theory of rational addiction. *Journal of Political Economy 96*, 675–700, 1988.

Green, L., Rachlin, H. and Hanson, J.: Matching and maximizing with concurrent ratio-interval schedules. *Journal of the Experimental Analysis of Behavior 40*, 217–24, 1983.

Herrnstein, R.J., Prelec, D. and Vaughan, W.: An intra-personal prisoners' dilemma. Paper presented at the IX Symposium on the Quantitative Analysis of Behavior, Harvard University, 1986.

Herrnstein, R.J. and Prelec, D.: A theory of addiction. In G.F. Loewenstein and J. Elster (eds.), *Choice over Time.* Russell Sage Foundation, New York, 331–60, 1992.

Herrnstein, R.J., Loewenstein, G.F., Prelec, D. and Vaughan, W.: Utility maximization and melioration: Internalities in individual choice. *Journal of Behavioral Decision Making 6*, 149–85, 1993.

Heyman, G. and Herrnstein, R.J.: More on concurrent interval-ratio schedules: A replication and review. *Journal of the Experimental Analysis of Behavior 46*, 331–51, 1986.

Kagel, J.H. and Green, L.: Intertemporal choice behavior: Evaluation of economic and psychological models. In L. Green and J.H. Kagel (eds.), *Advances in Behavioral Economics*. Volume 1, Ablex, Norwood, 166–84, 1987.

Laibson, D.: Self-control and saving. Mimeo, MIT, 1994.
Loewenstein, G. and Thaler, R.H.: Anomalies. Intertemporal choice. *Journal of Economic Perspectives 3*, 181–93, 1989.
Simon, H.A.: *Models of Man*. John Wiley & Sons, New York, 1957.
Smith, V.L.: Experimental economics: Induced value theory. *American Economic Review Papers and Proceedings 66*, 274–9, 1976.
Strotz, R. H.: Myopia and inconsistency in dynamic utility maximization. *Review of Economic Studies 23*, 165–80, 1956.
Thaler, R.H.: Some empirical evidence on dynamic inconsistency. *Economics Letters 8*, 201–7, 1981.
Yaari, M.E.: Consistent utilization of an exhaustible resource, or: How to eat an appetite-arousing cake. Unpublished paper; Walras-Bowley Lecture, North-American meeting of the Econometric Society, Ottawa, June 1977.

First version submitted August 1995;
final version received October 1997.