Experimental Economics, 3:137–152 (2000)

An Experimental Analysis of Intertemporal Allocation Behavior

VITAL ANDERHUB, WERNER GÜTH, WIELAND MÜLLER AND MARTIN STROBEL
Humboldt University at Berlin, Department of Economics, Institute for Economic Theory III, Spandauer Strasse 1, 10178 Berlin, Germany

Abstract

If the future is uncertain, optimal intertemporal decisions rely on anticipating one's own optimal future behavior as is typical in dynamic programming. Our aim is to detect experimentally stylized facts about intertemporal decision making in a rich stochastic environment. Compared to previous experimental studies our experimental design is more complex since the time horizon is uncertain and termination probabilities have to be updated. In particular the decision task is non-stationary as in real life which seriously complicates the task of diagnosing behavioral regularities. In this study we give some illustrative results and provide some general perspectives. Our main result is that subjects' reaction to information about termination probablilities are qualitatively correct.

Keywords: intertemporal allocation behavior

JEL Classification: C91; D81; D90

1. Introduction

When termination probabilities depend on stochastic events during "life" optimal intertemporal decision behavior becomes very difficult to derive since it requires dynamic programming techniques and Bayesian updating. Up to now only a few experimental studies[1] have tried to analyse such behavior (see the survey[2] by Anderhub and Güth, 1999). Unlike the previous experimental studies we do not really hope to find optimal behavior. Our main intention is to detect behavioral regularities and to study the dynamic processes of reasoning and decision making by boundedly rational individuals—a topic that is surprisingly under-researched as far as experimental economics is concerned.

Avoiding Bayesian updating by imposing constant termination probabilities rules out an important aspect in real-life saving behavior. Allowing later termination probabilities to depend on earlier results would enormously complicate life-cycle analysis. We therefore have developed a finite horizon model which can more easily accommodate such phenomena.

In our experiment participants learn more about their individual termination probability during "life". This was implemented by using three dice, representing different termination probabilities, of which one was excluded after the first choice and another after the second choice. The remaining die then represents the constant termination probability after the third choice.

The experimentally observed periodic expenditure choices should reveal how probabilities are updated and how a highly uncertain future is anticipated. It will, however, require many more studies to develop an empirically based behavioral alternative for dynamic programming. The benchmark case of risk neutral intertemporal optimization is derived numerically and compared with the actual decisions.

In the following Section 2 we describe our experimental design. Section 3 is devoted to the optimal decision pattern in case of risk neutrality. The results of our experiments are reported in Section 4. Finally, in Section 5 we discuss our findings.

2. Experimental setup

Assume that one "lives" for exactly T (>1) periods during which one can spend S_1 monetary units in total. Denoting by x_t the expenditure in period $t = 1, \ldots, T$ the optimal intertemporal decision vector $x^* = (x_1^*, \ldots, x_T^*)$ is derived by maximizing the intertemporal utility function $u(x_1, \ldots, x_T)$ subject to constraints $x_t \geq 0$ for $t = 1, \ldots, T$ and $x_1 + \cdots + x_T \leq S_1$. The computational difficulty is similar to determining an optimal consumption vector for T different products with given prices. Behavioral regularities like underestimation (see Fehr and Zych, 1995) or overestimation of future needs (see Johnson et al., 1987) are reported even for deterministic decision environments.

In our experiment we capture the stochastic nature of human life (expectation) and of most intertemporal decision problems by making T, the length of "life", a stochastic variable.[3] More specifically, in our experiment a participant is sure to "live" for at least three periods and no longer than six periods, i.e. $3 \leq T \leq 6$ and $T \in \mathbb{N}$. Whether the subject experiences a 4th, 5th, or 6th period is successively determined by one of three dice of different colors standing for different conditional survival probabilities, namely $\frac{1}{2}$ (red die), $\frac{2}{3}$ (yellow die), and $\frac{5}{6}$ (green die). At the beginning of each round, a subject does not know which of the dice will be applied. Instead of this she is informed that after confirming her expenditure amount x_1 in period 1, one of the three dice (probabilities) is randomly taken out and that after the choice of x_2 in period 2 one of the remaining two dice is also randomly taken out. The last die then determines via successive, independent and identical chance moves whether the participant lives for 3, 4, 5, or 6 periods. Thus, one only knows from the third period on which of the three probabilities will be finally applied.

Participants (mainly students of economics or business administration at Humboldt University Berlin) were invited by leaflets to register for the experiments which were predicted to last at most two hours. In the lab participants were seated at isolated terminals where they found typed German instructions which were also available on the screen. (See the Appendix for an English translation of the instructions used in the Π-treatment. Except for a few changes the same instructions were used in the Σ-treatment.)

The instructions informed participants that the experiment consists of 12 rounds and that in each round their task is to allocate a given monetary amount ($S_1 = 11.92$ ECU)[4] over an unknown number of periods (at least 3, at most 6). It was explained that after period 3, 4 and 5 it will be periodically decided by one of the three (red, yellow, green) dice with different termination numbers ({1, 2, 3} for red, {1, 2} for yellow, {1} for green) whether or not there will be a further period. Let S_t $(t = 2, \ldots, 6)$ denote what is left of S_1 before deciding in period t, i.e. for $0 \leq x_t \leq S_t$ one has $S_{t+1} = S_t - x_t$. Participants were informed that after their choice of x_1 one of the three dice is excluded and that one of the two remaining dice is excluded after choosing x_2.

Our experimental treatment variable concerns how the monetary win depends on the allocation pattern $x_1, \ldots, x_T$. We distinguish the payoff function

$$U_\Pi = \prod_{t=1}^{T} x_t = x_1 \cdot x_2 \cdot \cdots \cdot x_T$$

named the Π-treatment and the payoff function

$$U_\Sigma = \sum_{t=1}^{T} \sqrt{x_t} = \sqrt{x_1} + \sqrt{x_2} + \cdots + \sqrt{x_T}$$

named the Σ-treatment.

The software of the computerized experiment offered access to a calculator so that participants could easily check the numerical consequences of certain choices.[5] Before playing the 12 rounds of the game based on U_Π or U_Σ participants are asked to fill out the 16PA-personality questionnaire (Brandstätter, 1988), including personal characteristics like age, gender, and subject of study. Having played the game for 12 rounds they are debriefed by asking them to rate the experimental situation. Altogether 50 participants played the Π-game and 50 the Σ-game. Without giving any time restrictions sessions needed 90 minutes on average.

Altogether there are 6 possible sequences of initial chance moves (three possible dice for the first, two possible dice for the second chance move). Each participant played all six sequences in a random order before they were repeated in another random order.[6] In the following we refer to the first random order of the six possible sequences as the "first cycle", including rounds 1 to 6, and to the second random order as the "second cycle". Experience effects can be explored by comparing the behavior for the same sequence in the first and second cycle, but also within a cycle (within a cycle two of the six sequences rely on the same first chance move).

Before the experiment, participants are told how $U \in \{U_\Pi, U_\Sigma\}$ is related to their monetary win in DM (German Mark) which was DM 1.00 = ECU 1.00 for the Π-treatment and DM 5.00 = ECU 1.00 for $U = U_\Sigma$. Instead of imposing that the actual win of a participant is the average win of all 12 rounds or of a randomly selected round, participants are asked to decide this for themselves. This decision could indicate personal attitudes towards risk.

3. Optimal consumption behavior

Consider the case when consumption choices are evaluated by U_Π. If the length of "life" T with $T \in \{3, 4, 5, 6\}$ were known, the optimal consumption pattern could easily be derived as $x_t^* = S_1/T$ for $t = 1, \ldots, T$. However, for an experimental subject the optimization task

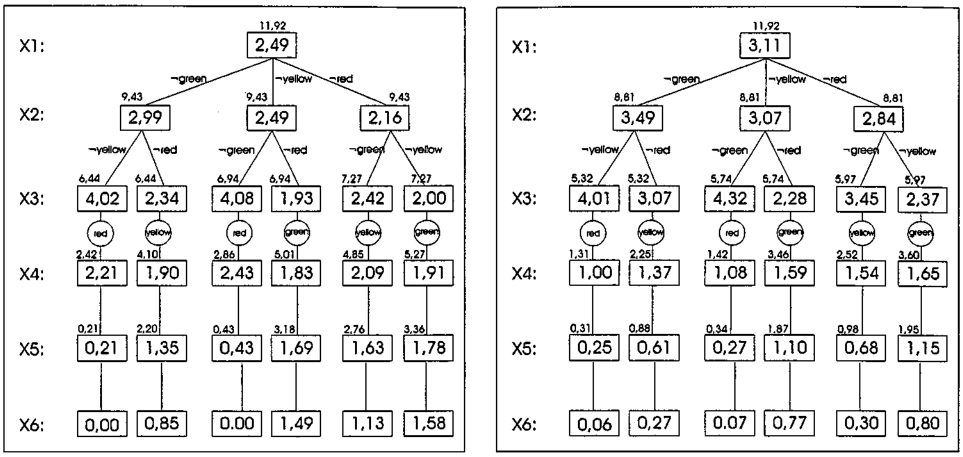

Figure 1. Optimal consumption behavior (Π left, Σ right).

becomes prohibitively difficult when T is a stochastic variable with values $T \in \{3, 4, 5, 6\}$. The fact that the survival probability can assume three different levels illustrates the practical impossibility of deriving the optimal behavior, at least in the course of the experiment.[7]

In fact we ourselves found it difficult to derive the optimal behavior of risk neutral decision makers for U_Π and U_Σ. However, relying on numerical methods we were able to compute the optimal consumption paths listed in figure 1 which will be used as the benchmark solutions (U_Π left and U_Σ right). Here "¬green" means, for instance, that the green die with termination probability of $\frac{1}{6}$ has been excluded. Whereas the boxes of figure 1 contain the optimal choices x_t^*, the resulting residual funds S_t^* are given above the boxes.

For U_Π only two paths imply $x_t = 0$, namely $x_6^* = 0$ in case of the two sequences with ¬green and ¬yellow that offer the least chances for a long "life". If in spite of the low continuation probability of $\frac{1}{2}$, implied by the red die, one lives for six periods ($T = 6$), U_Π would be 0 due to $x_6^* = 0$.[8] Of course, other consumption paths rely on similar forms of gambling, but they never prescribe $x_t^* = 0$ for some period t. For U_Σ an extreme choice $x_t^* = 0$ can never be optimal since the marginal utility of x_t goes to $+\infty$ when x_t approaches 0. For both payoff functions consumption increases after "bad news" (¬green) whereas consumption decreases after "good news" (¬red).

In general, the situation is more risky for U_Π than for U_Σ where $U_\Sigma = 0$ is possible only in the case of the absurd behavior $x_1 = x_2 = \cdots = x_T = 0$. For the optimal consumption behavior described in figure 1 the expected value μ of U_Π is 35.16 and the standard deviation σ is 18.25 whereas for U_Σ the corresponding values are 6.75 and 1.10.

Allowing for risk aversion (see Müller, in press) does not change the decisive qualitative properties of the risk neutral benchmark solution such as increasing/decreasing consumption after "bad"/"good" news or strictly monotonically decreasing consumption from the third period on. And since we are more interested in the qualitative characteristics of optimal behavior than in exact numbers we will use the solution that assumes risk neutrality as a reference.[9]

4. Experimental results

4.1. Average observed behavior and efficiency

The average profit of the $600 = 50 \cdot 12$ (participants $\cdot$ rounds) observed plays was 27.62 ECU (1 ECU $=$ 1 DM) for the Π-treatment and 6.50 ECU (1 ECU $=$ 5 DM) for the Σ-treatment.

To give a first impression of the decision data[10] figure 2 shows the mean, minimum, maximum and variance values for every possible decision node. Above each box the number of cases is given. The line above or below the entries of each box indicates whether the observed mean is above or below the benchmark value (figure 1). Note that in the first three (certain) periods of the Σ-treatment subjects tend to oversave. This is not the case for the Π-treatment.

Recall that according to optimal behavior (see figure 1) the exclusion of a die requires updating and adjustment of consumption levels. Almost all means in figure 2 display the same ordinal relations to each other as the corresponding optimal choices shown in figure 1. The benchmark solution always prescribes decreasing expenditures after the third period which is also true for the average behavior. Naturally, the distance between maximum and

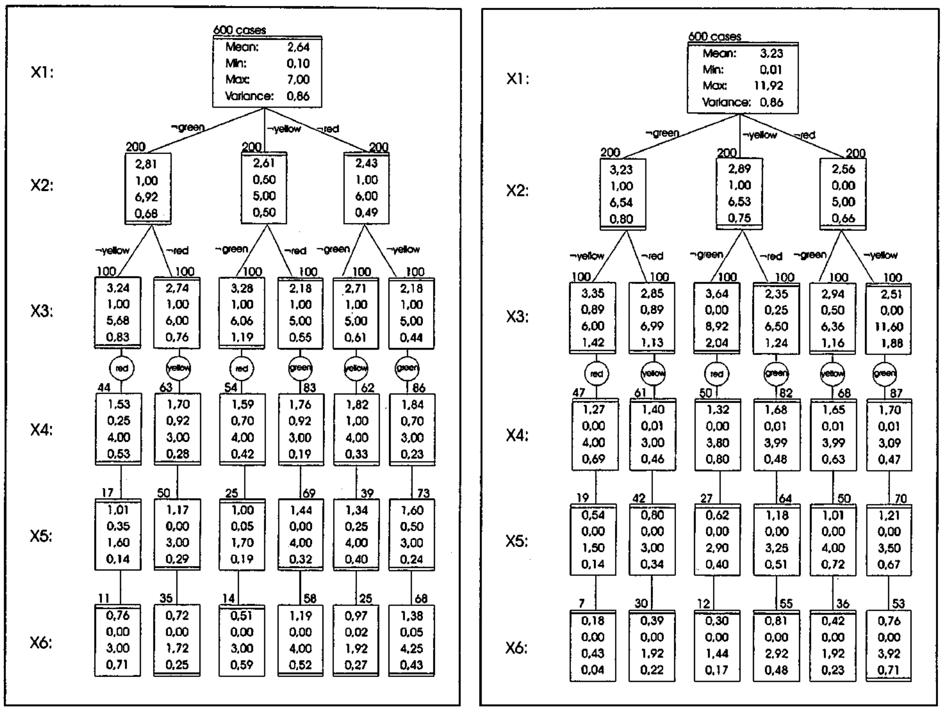

Figure 2. Average behavior: Mean, maximal, minimal x_t nad variance for each cell (Π-treatment left, Σ-treatment right).

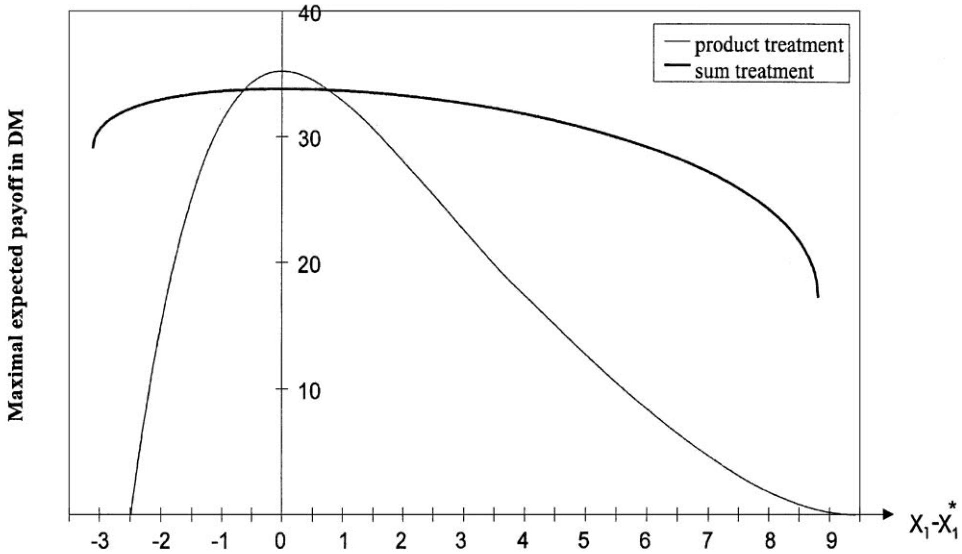

Figure 3. Maximal obtainable payoff for given deviation from x_1^*.

minimum (see figure 2) becomes usually smaller in later periods when the available funds are smaller. Thus we state

Observation 1. Average observed behavior is qualitatively similar to optimal behavior.

Let U_k with $k \in \{1, 2\}$ denote the average payoff in cycle k and let U^* denote the optimal expected payoff for the Π- and Σ-treatment. The efficiency rate U_k/U^* for the Π-treatment is .79 for the first cycle and .78 for the second cycle, for the Σ-treatment these values are .96 for the first cycle and .97 for the second cycle. Thus the average efficiency rate essentially does not change with experience.

The efficiency rate in the Σ-treatment is consistently higher for U_Σ than for U_Π since U_Σ is more "flat" — in the sense that the same relative deviation from the optimal choice x_t^* implies a smaller loss.[11] Figure 3 displays the maximal obtainable payoff after a given choice of x_1, i.e. under the assumption that after the choice of x_1 a player's behavior is conditionally optimal.[12] Since the relative deviations from the benchmark behavior in the Σ-treatment are at least as large as in the Π-treatment, the lower efficiency for the Π-treatment is due to its more reactive monetary incentives.

4.2. *Initial consumption*

In this subsection we investigate consumption within the certain periods $t = 1, 2, 3$. For that purpose define $\mu_k := (x_1^k + x_2^k + x_3^k)/11.92$ where $k = 1, 2, \ldots, 6$ refers to the six possible orders of initial chance moves. Thus μ_k measures the relative amount consumed

Table 1. Ranked sequences of initial consumption (optimal values in brackets).

Sequence	1st chance move	2nd chance move	Π-treatment	Σ-treatment
1	¬green	¬yellow	.73 (.80)	.83 (.89)
2	¬yellow	¬green	.72 (.76)	.82 (.88)
3	¬green	¬red	.69 (.66)	.78 (.81)
4	¬red	¬green	.65 (.59)	.74 (.79)
5	¬yellow	¬red	.63 (.58)	.71 (.71)
6	¬red	¬yellow	.60 (.56)	.70 (.70)

during the three certain periods. We can answer the question of whether there is over- or undersaving in the certain periods by comparing the average observed value $\bar{\mu}_k$ with the value implied by the benchmark solution. Of course, it is also interesting to compare the measures $\mu(\Sigma)$ for U_Σ with the measures $\mu(\Pi)$ for U_Π.[13]

In Table 1 the 6 possible sequences of initial chance moves are ranked from $k = 1$ to $k = 6$ according to the level $\mu_k^* = \frac{(x_1^k)^* + (x_2^k)^* + (x_3^k)^*}{11.92}$ which the benchmark choices $(x_1^k)^*$, $(x_2^k)^*$ and $(x_3^k)^*$ imply. (The values $(x_t^k)^*$ can be inferred from figure 1.) Table 1 shows the means $\bar{\mu}_k$ $(k = 1, \ldots, 6)$ of the observed values μ_k for each treatment and sequence separately. Furthermore, the optimal values μ_k^* are given in parentheses. Inspection of Table 1 reveals that the observed means are ranked in the same order as the optimal values, i.e. on average subjects' updating of termination probabilities is qualitatively correct. Furthermore, the observed means $\bar{\mu}_k$ of the Σ-treatment are on average about 10% higher than in the Π-treatment. Our main conclusion is that subjects' reactions to information about termination probabilities are qualitatively correct.

This conclusion can be statistically validated: Let us divide the six sequences into two groups. Group 1 contains sequence 1, 2 and 3 and group 2 all others, i.e. we combine the sequences with higher initial consumption in group 1 and those with lower initial consumption in group 2. For each round we can assign each participant to one of the two groups. Using the Mann–Whitney–U test we have checked whether consumption in these groups (see Table 2) can be viewed as being significantly different. The null-hypothesis of equal initial consumption in group 1 and 2 can usually be rejected.[14] Note that this test already indicates significant differences in the first round of both treatments. Only for the 5$^{\text{th}}$ round of the Σ-treatment we could not reject the null-hypothesis. Thus, participants'

Table 2. Average initial consumption shares within the two groups.

Round	1	2	3	4	5	6	7	8	9	10	11	12
Π, group 1	.70	.70	.73	.74	.71	.69	.69	.71	.72	.71	.72	.71
Π, group 2	.61	.62	.65	.65	.63	.64	.63	.62	.62	.62	.62	.61
Σ, group 1	.80	.82	.82	.79	.78	.84	.82	.79	.80	.83	.79	.79
Σ, group 2	.70	.71	.73	.71	.76	.71	.70	.70	.69	.69	.70	.74

reactions were at least qualitatively as predicted, i.e. higher initial consumption after "bad news" and lower initial consumption after "good news".

Summarizing these findings we state

Observation 2. On average subjects' reactions to initial chance moves are qualitatively correct.

On the individual level the decision x_2 should be influenced by the information about the first excluded die. If, for instance, the red die is excluded this implies a longer expected life. The average consumption levels $\frac{x_2}{S_2}$ should fulfill the condition

$$\left(\frac{x_2}{S_2}\middle|\neg\text{green}\right) > \left(\frac{x_2}{S_2}\middle|\neg\text{yellow}\right) > \left(\frac{x_2}{S_2}\middle|\neg\text{red}\right). \tag{1}$$

In order to test this we computed individual averages, i.e. we took the means of the relative expenditures $\frac{x_2}{S_2}$ over all 12 rounds for every individual. The binomial test for

H_0: Subjects do not fulfill condition (1)

rejects this hypothesis with $p < 0.005$ for both treatments. It can thus be maintained that subjects tend to react correctly. Accordingly x_3 should depend on the finally remaining die as follows:

$$\left(\frac{x_3}{S_3}\middle|\text{red}\right) > \left(\frac{x_3}{S_3}\middle|\text{yellow}\right) > \left(\frac{x_3}{S_3}\middle|\text{green}\right). \tag{2}$$

The hypothesis

H_0: Subjects do not fulfill condition (2)

was tested in the same way as above and rejected with $p < 0.005$. These results are summarized by

Observation 3. On the individual level, subjects' updating of termination probabilities is qualitatively correct.

When choosing x_1 a participant encounters in all 12 rounds the same decision problem. Nevertheless we observe widely varying x_1-values for the same participant: For U_Π only 11 of 50 participants rely on the same x_1-choice for the first cycle of the 6 possible sequences of initial chance moves; for the second cycle this number increases to 17 (8 participants always chose the same x_1-value). For U_Σ these numbers were 6 for the first cycle, 13 for the second cycle and 5 always. Even after playing 6 rounds of the game, many participants are still experimenting with x_1.[15]

Since the initially available amount S_1 was not prominent, a participant could yield to prominence once by choosing prominent levels of x_t or by inducing prominent values S_t

for $t > 1$. For U_Π the outstanding focal decision was $x_1 = 2.00$ (144 times), for U_Σ it was $x_1 = 3.00$ (111 times). Only for U_Π the values $x_1 = 1.92$ (49 of the 600 observations) and $x_1 = 2.92$ (26 times) were rather prominent (for U_Σ the corresponding frequencies are 11 and 13).

4.3. Behavior in later periods

We now analyse the behavior in the second part of a round consisting of the uncertain periods 4, 5 and 6. When choosing x_t for $t \geq 3$ participants do not know whether "life" ends in period t in which case all saved money would be lost. Optimal as well as boundedly rational behavior requires $x_3 > x_4 > x_5 > x_6$.[16]

Table 3 (for the Π-treatment) and Table 4 (for the Σ-treatment) show the number of plays reaching at least the 4th, 5th or 6th period and the hit rates for the listed criteria. Substituting

Table 3. Facing an uncertain future (Π-treatment).

	Cases	%
$T \geq 4$	392	100.0
$x_3 > x_4$	267	68.1
$x_3 \geq x_4$	352	89.8
$T \geq 5$	273	100.0
$x_4 > x_5$	179	65.6
$x_4 \geq x_5$	244	89.4
$T = 6$	211	100.0
$x_5 > x_6$	160	75.8
$x_5 \geq x_6$	192	91.0

	Cases	%
$T \geq 5$	273	100.0
$x_3 > x_4 > x_5$	127	46.5
$x_3 \geq x_4 \geq x_5$	215	78.8
$T = 6$	211	100.0
$x_4 > x_5 > x_6$	98	46.4
$x_4 \geq x_5 \geq x_6$	171	81.1

	Cases	%
$T = 6$	211	100.0
$x_3 > x_4 > x_5 > x_6$	75	35.5
$x_3 \geq x_4 \geq x_5 \geq x_6$	150	71.1

Table 4. Facing an uncertain future (Σ-treatment).

	Cases	%
$T \geq 4$	395	100.0
$x_3 > x_4$	295	74.7
$x_3 \geq x_4$	365	92.4
$T \geq 5$	272	100,0
$x_4 > x_5$	210	77.2
$x_4 \geq x_5$	249	91.5
$T \geq 6$	193	100.0
$x_5 > x_6$	159	82.4
$x_5 \geq x_6$	180	93.3

	Cases	%
$T \geq 5$	272	100.0
$x_3 > x_4 > x_5$	159	58.5
$x_3 \geq x_4 \geq x_5$	234	86.0
$T \geq 6$	193	100.0
$x_4 > x_5 > x_6$	133	68.9
$x_4 \geq x_5 \geq x_6$	170	88.1

	Cases	%
$T \geq 6$	193	100.0
$x_3 > x_4 > x_5 > x_6$	94	48.7
$x_3 \geq x_4 \geq x_5 \geq x_6$	159	82.4

the strict inequality $x_t > x_{t+1}$ by the weak one, $x_t \geq x_{t+1}$, strongly increases the predictive success of the criteria. When, for instance, participants reach the sixth period in the Π-treatment, 35.5% of them obey the criterion $x_3 > x_4 > x_5 > x_6$ and 71.1% the weaker condition $x_3 \geq x_4 \geq x_5 \geq x_6$. The remaining 28.9% failed at least once. Approximately 90% of all cases satisfy $x_t \geq x_{t+1}$ when only two subsequent periods are compared. In general, results are better for the Σ-treatment than for the Π-treatment.

Observation 4. On the individual level many subjects do not satisfy the mild conditions $x_t > x_{t+1}$ for $t \geq 3$ for bounded rationality.

One may argue that optimal consumption behavior cannot be expected when it is practically impossible for the participants to compute it, but that it yields reliable predictions when it is easily derived. Nearly all (S_5, x_5)-observations in the case of the red die lie below the conditionally optimal decision curve which is (piecewise) linear for U_Σ (U_Π). Thus participants have usually left more for period 6 than required by the benchmark solution. Similar, but less clear results apply to the other dice.

4.4. Randomize- versus average-types

Before determining x_1 for the first time each participant is asked whether he prefers the average payoff of all 12 rounds or the payoff of a randomly selected round. We will refer to participants of type A (average) and of type R (random). In case of U_Π 16 of 50 participants are of type R whereas for U_Σ this number is 10. One can view this selection as revealing personal attitudes towards risk. R-types could be seen as more risk-loving since they prefer a payment which is the result of a lottery.[17]

We first compare the behavior of the A- and R-participants when choosing x_1. Table 5 shows the mean of x_1 for each of the twelve rounds and for both types and treatments separately as well as the p-levels resulting from a one-sided Mann–Whitney–U test (H_0 : $\mu_A \leq \mu_R$). For both treatments the means of the R-types are smaller than the means of the A-types, except for the 6th round of U_Σ. But these differences are usually significant only for U_Π (they are not significant in rounds 6 and 9 of the Π-treatment).

The choices x_2 and x_3 do not reveal similarly strong differences between types even when they rely on the same initial chance moves. However, the relative amount $\mu = \frac{x_1+x_2+x_3}{11.92}$

Table 5. Means of x_1 for both types and treatments.

Round	1	2	3	4	5	6	7	8	9	10	11	12
Π, type A	2.67	2.81	2.92	2.96	2.76	2.64	2.75	2.67	2.68	3.20	2.75	2.66
Π, type R	2.53	2.28	2.29	2.30	2.38	2.42	2.23	2.24	2.49	2.26	2.31	2.23
p-levels	.073	.031	.025	.060	.071	.251	.018	.015	.210	.008	.071	.025
Σ, type A	3.37	3.61	3.30	3.45	3.76	3.18	3.59	3.48	3.21	3.06	3.16	3.20
Σ, type R	2.82	2.72	2.92	2.62	2.82	3.22	2.52	3.22	2.22	2.47	2.32	2.42
p-levels	.272	.064	.304	.131	.121	.476	.098	.339	.064	.292	.098	.174

Table 6. Means of $\mu = \frac{x_1+x_2+x_3}{11.92}$ for both types and for all sequences.

Sequence	¬gr.,¬yel.	¬yel,¬gr	¬gr.,¬red	¬red,¬gr.	¬yel.,¬red	¬ red,¬yel.
Π, type *A*	0.75	0.74	0.70	0.67	0.64	0.61
Π, type *R*	0.68	0.67	0.66	0.62	0.60	0.57
Σ, type *A*	0.85	0.83	0.79	0.75	0.72	0.70
Σ, type *R*	0.73	0.77	0.72	0.69	0.66	0.66

Table 7. Average results for both types.

	Π-treatment		Σ-treatment	
	A-types	*R*-types	*A*-types	*R*-types
Average payoff	27.14	28.63	6.48	6.56
Std. deviation	14.62	17.52	1.17	1.29
Efficiency rate	0.77	0.81	0.96	0.97

spent in the certain periods $t = 1, 2, 3$ is again type dependent (see Table 6) what, of course, is mainly due to the differences in the choice of x_1. Note, that in our experimental situation there is both the risk of spending too much in the beginning and "living long" and the risk of saving too much in the beginning and "dying early". If R-types were just hoping for a long life this would explain why they consumed less than A-types in the certain periods. This would be consistent with our interpretation of R-types.

Let us also compare the average payoffs, standard deviations and the efficiency rates U_k/U^*, $k \in \{A, R\}$, where U_k denotes the average payoff of $k = A$ and $k = R$ for all twelve rounds and U^* the expected payoff of the benchmark behavior.

R-types are slightly more successful than A-types (Table 7). But these differences are not significant. We summarize this by

Observation 5. Self-selection between random and average payment (R- and A-types) is only weakly correlated with individual differences in intertemporal allocation behavior.

4.5. *Decision times*

The decision times m_t^τ (τ- index of the round, t- index of the period) which were recorded by the computer for every decision x_t^τ provide an interesting experimental observation. In order to always rely on the same number of observations we investigate how the total decision time $m_{1,2,3}^\tau = m_1^\tau + m_2^\tau + m_3^\tau$ for $x_1^\tau, x_2^\tau, x_3^\tau$ changes from "life" to "life". We expected that the time $m_{1,2,3}^\tau$ that a participant needs in round τ depends on τ like $m_{1,2,3}^\tau = \alpha + \frac{\beta}{\tau}$ with $\tau \in \{2, 3, \ldots, 12\}$.[18] Calculating the regression for both treatments we get the following results: For the Π-treatment $\alpha = 52.9$ and $\beta = 543.4$ ($R^2 = 0.976$) and for the Σ-treatment

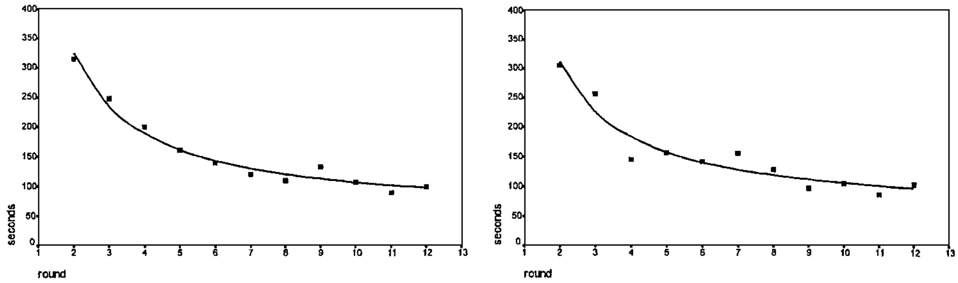

Figure 4. Regressed and actual average for the decision times (Π left and Σ right).

$\alpha = 55.4$ and $\beta = 512.4$, ($R^2 = 0.920$). The mean decision times $m^{\tau}_{1,2,3}$ for x_1, x_2 and x_3 as well as the regression lines are shown in figure 4. These results strongly confirm the "learning by playing"-process.

5. Summary and outlook

Compared to other experimental studies of dynamic decision making our design is more complex since players have to update their termination probabilities which depend on stochastic events during "life". Thus we could investigate whether people react to information (via the choices of x_1, x_2 and x_3 for different sequences of initial chance moves). Intertemporal allocation behavior in the more classical setting of constant termination probabilities has been explored by inspecting the choices of x_3, x_4, and x_5. Using this general frame we implemented different risk structures by implementing different payoff functions. The main findings of our study are:

(i) Average observed behavior displays similar effects as the benchmark solution, based on risk neutral utility maximization.
(ii) In the complex stochastic environment on average subjects react in a qualitatively correct way to "good" or "bad" news, i.e. on average subjects make use of particular information concerning their length of "life".
(iii) With regard to the uncertain horizon we find that subjects are initially too cautious in case of the Σ-treatment, i.e. they usually leave (compared to the benchmark) more for the uncertain periods. In the Π-treatment behavior in early periods depends on which dice are taken out, i.e. how termination probabilities have to be revised.
(iv) Self-selection between random and average payment (R- and A-types) is only weakly correlated with individual differences in intertemporal allocation behavior.
(v) Further data analysis revealed that qualitative learning, e.g. in the sense of directional learning (see Selten and Buchta, 1999, for a more positive result) is only weakly confirmed. When subjects are confronted with a highly stochastic environment it is apparently very difficult to attribute poor results to bad luck or improvable choices.

Findings (i) and (ii) confirm the economic theory of intertemporal choice at least in a qualitative sense and are, thus, in line with the positive results reported in other studies (e.g., Hey, 1982; Gigliotti and Sopher, 1997).

Intertemporal decision making is one prominent topic of the "anomalies and biases"-program (see, e.g., Thaler, 1981; Loewenstein, 1988; and for a survey Loewenstein and Thaler, 1989) questioning the empirical relevance of optimal solutions. Often the experiments did not rely on monetary rewards at all and, if so, their salience is questionable.[19] In our study we distinguish two types of intertemporal preferences of which one is much more salient in the sense that deviating from optimality is much more costly. How can our results be related to this debate? On the one hand it is obvious that the advice to maximize expected payoffs offers no practical help since the problem is much too complex. On the other hand many participants react adequately to changes in life expectation where saliency, as controlled by our two treatments, does not seem to matter much. Thus, for the anomalies and biases-program our results could demonstrate that optimization does not work and that even basic aspects of bounded rationality are neglected. However, since observed average behavior in the complex stochastic environment is much in line with the intuition, provided by the theoretical benchmark solution, one can argue that our results support traditional economic theory.

In general, many more research questions can be explored with the help of our experimental data. Moreover, from the very beginning our experimental design has been structured in such a way that it allows for systematic variations such as simplifying or complicating its stochastic nature. Anderhub (1998) analyses a model that is similar to the one developed here but with reduced stochastic complexity (see endnote 20). In this study dice of different colors no longer represent termination probabilities, but rather the actual length of life. This study confirms that subjects' reactions to information are qualitatively correct. Müller (in press) considers alternatives (strategies and heuristics) to the optimal backward induction solution and investigates how the solution is influenced by allowing for risk aversion. Brandstätter and Güth (1998) account for differences in intertemporal allocation behavior by different personality types as elicited by the 16 PA-personality questionnaire. Our experiment has also been performed on the Internet (Anderhub et al., in press).

Appendix—Translated instructions of the Π-treatment

Your task in every round is to distribute an amount of money as well as possible to several periods. The better this is done the higher is your payoff. Altogether you play 12 rounds. In the beginning of the experiment you can choose whether we should draw lots to select one round for which you are paid. Otherwise you will receive the mean of your payoffs of all rounds. In any case you get your payoff in cash after evaluation of the data.

The general aim of each round is to distribute a certain amount of money to several periods. Your payoff of one round is equal to the product of the amounts of money allocated to the single periods. However, there is no certainty about the number of periods to which you have to distribute your money. The game can last for three, four, five, or six periods. Every round will last for at least three periods. Whether you reach the fourth, fifth or sixth period will be determined by throwing a die. There are three different dice colored red, yellow and green. The following table shows in which cases you reach the next period.

Color of die	No further period if die shows	New period if die shows
Red	1, 2, 3	4, 5, 6
Yellow	1, 2	3, 4, 5, 6
Green	1	2, 3, 4, 5, 6

The number of periods of which one round consists can not be higher than six. In the beginning of a round you do not know which die is used. You get this information after you have made some decisions. The general course of the game is as follows:

1st period) You will get an amount of money S, which you can spend in the coming periods. Altogether you can only spend this amount. You choose an amount x_1, which you want to spend in the first period. Think very carefully about how much you want to spend and how much you want to save for the following periods. After your decision one of the three dice is excluded. Now you know that only the two other dice are candidates for the chance moves deciding whether you reach the fourth, fifth or sixth period.

2nd period) You choose an amount x_2, which you want to spend in the second period. You can only spend what you have left from the total amount after the first period. After your decision another die is excluded. Now you know which die remains to be thrown for the fourth, fifth and sixth period.

3rd period) You choose an amount x_3, which you want to spend in the third period. After this decision the computer will throw the remaining die in order to decide whether you reach the fourth period. If you do not reach the fourth period, the round ends here. The amount which has not been spent until now is lost.

4th period) If you reach the fourth period, you will have to choose an amount x_4. Afterwards the die will be thrown again.

5th period) If you reach the fifth period, you will have to choose an amount x_5. Afterwards the die will be thrown again.

6th period) If you reach the sixth period, you do not have to make a decision, because all remaining money is spent automatically.

Your payoff is calculated by the product of all amounts you spent in the periods you reached. For instance if you reached exactly four periods, your payoff is determined by $G = x_1 \cdot x_2 \cdot x_3 \cdot x_4$. When you have reached all six periods, your payoff is determined by $G = x_1 \cdot x_2 \cdot x_3 \cdot x_4 \cdot x_5 \cdot x_6$ where x_6 is the amount you have left after the fifth period. Please think about the following: If you spend an amount of 0 in one period, your payoff will also be 0, because one of the factors is 0. This can happen, for instance, if you spend all money in the fourth period and reach the fifth period. Then you have to spend 0 in the fifth and perhaps also in the sixth period and therefore you get the payoff 0. You have to consider both the risk of spending all your money early and the risk of making your money useless in the case the game ends.

Acknowledgments

We wish to thank Thomas Kötter, two anonymous referees and Arthur Schram for helpful comments and suggestions and Alexandra Schumacher for implementing the software. The financial support of the Deutsche Forschungsgemeinschaft (Gu145/8-1, SFB 373) and the TMR-project "Savings and Pensions" by the European Commission is gratefully acknowledged.

Notes

1. Early field evidence is reviewed by Hall (1978).
2. This survey is confined to studies in experimental economics. Intertemporal allocation behavior, especially saving behavior is, of course, also a topic in economic psychology (see, for instance, Wärneryd, 1999).
3. Positive termination probabilities imply less concern for the future like positive discount rates, i.e. discount factors smaller than 1 (see Anderhub and Güth, 1999, who present various formulations and their experimental results).
4. ECU—Experimental Currency Unit.
5. We did not record how often and to what purpose the calculator was used.
6. We wanted to control the exclusion of the dice in order to get more comparable results and to test for experience effects. We did not find any evidence that participants noticed this minor regularity.
7. For given $S_5(>0)$, termination probability w with $0 < w < 1$ and $C = x_1 \cdot x_2 \cdot x_3 \cdot x_4(>0)$ the total payoff is $U_\Pi = C \cdot x_5[w + (1 - w)(S_5 - x_5)]$. From $\frac{\partial U_\Pi}{\partial x_5} = 0$ and $\frac{\partial^2 U_\Pi}{\partial x_5 \partial x_5} = -2C(1 - w) < 0$ one obtains $x_5^*(S_5, w) = \min\{\frac{w}{2(1-w)} + \frac{S_5}{2},\ S_5\}$ for all $x_1, x_2, x_3, x_4, S_5 > 0$. Inserting x_5^* into U_Π yields already two different optimisation tasks for $t = 4$.
8. As demonstrated in Müller (in press), in the case of U_Π all consumption levels are positive if one allows for a sufficiently high degree of risk aversion.
9. For a control group of 14 participants we used the binary lottery-technique (see Roth and Malouf, 1979) for the Π-treatment. The results did not reveal obvious deviations from the results discussed here. Moreover, for evidence questioning the validity of the binary lottery-technique see Güth et al. (1993) and Selten et al. (1999).
10. The complete set of data is availiable from the authors upon request.
11. For $U = U_\Sigma$ the relative deviation measure $rd(x_t) = \frac{|x_t^*(S_t) - x_t|}{S_t}$ is usually larger than for $U = U_\Pi$ where $x_t^*(S_t)$ denotes the conditionally optimal choice, given the availiable fund S_t, and x_t the actual choice in period t.
12. The choice of $x_1 = S_1$, for instance, implies $U_\Pi = 0$ and $U_\Sigma = \sqrt{S_1}$.
13. The measures $\mu(\Pi)$ and $\mu(\Sigma)$ did not change very much from the first to the second cycle.
14. The Mann-Whitney U-test usually yields $p < .01$. Exceptions are rounds 7 and 11 of the Π-treatment and round 11 of the Σ-treatment where $p < .05$. In round 6 of the Π-treatment and rounds 4 and 12 of the Σ-treatment we get $p < .1$.
15. The data reveal that more variations in x_1 resulted in lower payoffs. This can be seen by looking at the correlation coefficient between the standard deviation of x_1 and the mean payoffs over all 12 rounds, which is $-.63$ for the Π-treatment and $-.37$ for the Σ-treatment.
16. Even extreme risk aversion does not lead to $x_t = x_{t+1}$ for $t \geq 3$ (see Müller, in press).
17. Note, however, that the lottery in which the R-types finally wish to participate is not exogenously given but is—to a large extent—determined by these subjects' performance during the experiment.
18. When the first decision screen appeared many participants started to read the instructions again. We therefore exclude round $\tau = 1$.
19. Harrison (1994) argues that "several of the most widely cited pieces of experimental evidence contrary to EUT [Expected Utility Theory] and Bayes Rule [. . .] do not satisfy the accepted precepts of experimental economics" and shows that "modifications to the experiments to remedy these design weaknesses result in observed choice behavior consistent with the predictions of economic theory" (p. 251).

20. Anderhub et al. (2000) add a further (self-selected) stochastic aspect by allowing to invest S_2 in a profitable but risky prospect.

References

Anderhub, V. (1998). “Savings Decisions When Uncertainty is Reduced—An Experimental Study.” Discussion paper No. 73, Sonderforschungsbereich 373, Humboldt University Berlin.

Anderhub, V. and Güth, W. (1999). “On Intertemporal Allocation Behavior—A Selective Survey of Saving Experiments.” *Ifo Studien.* 45, 303–333.

Anderhub, V., Güth, W., and Knust, F. (2000). “On Saving and Investing—An Experimental Study of Intertemporal Decision Making in a Complex Stochastic Environment.” Discussion paper No. 2, Sonderforschungsbereich 373, Humboldt University Berlin.

Anderhub, V., Müller, R., and Schmidt, C. (in press). “Design and Evaluation of an Economic Experiment via the Internet.” *Journal of Economic Behavior and Organization.* Forthcoming.

Brandstätter, H. (1988). “Sechzehn Persönlichkeits–Adjektivskalen (16 PA) als Forschungsinstrument anstelle des 16 PF.” *Zeitschrift für experimentelle und angewandte Psychologie.* XXXV, Heft 3, 370–391.

Brandstätter, H. and Güth, W. (1998). “A Psychological Approach to Individual Differences in Intertemporal Consumption Patterns.” Discussion Paper, Humboldt University Berlin.

Fehr, E. and Zych, P. (1995). “Die Macht der Versuchung: Irrationaler Überkonsum in einem Suchtexperiment.” *Zeitschrift für Wirtschafts–und Sozialwissenschaften* 114, Heft 4, 569–604.

Gigliotti, G. and Sopher, B. (1997). “Violations of Present–Value Maximization in Income Choice.” *Theory and Decision.* 43, 45–69.

Güth, W., van Damme, E., and Weber, M. (1993). “The Normative and Behavioral Concept of Risk Aversion—An Experimental Study.” Working Paper, Tilburg University.

Hall, R.E. (1978). “Stochastic Implications of the Life Cycle—Permanent Income Hypothesis: Theory and Evidence.” *Journal of Political Economy.* 86, 971–987.

Harrison, G.W. (1994). “Expected Utility Theory and the Experimentalists.” *Empirical Economics.* 19, 223–253.

Hey, J.D. (1982). “Search for Rules of Search.” *Journal of Economic Behavior and Organization.* 3, 65–81.

Johnson, S., Kotlikoff, L., and Samuelson, W. (1987). “Can People Compute? An Experimental Test of the Life Cycle Consumption Model.” Working Paper, Harvard University.

Loewenstein, G.F. (1988). “Frames of Mind in Intertemporal Choice.” *Management Science.* 34, 200–214.

Loewenstein, G.F. and Thaler, R.H. (1989). “Anomalies: Intertemporal Choice.” *Journal of Economic Perspectives.* 3, 181–193.

Müller, W. (in press). “Strategies, Heuristics and the Relevance of Risk Aversion in a Dynamic Decision Problem.” *Journal of Economic Psychology.* Forthcoming.

Roth, A. and Malouf, M. (1979). “Game Theoretic Models and the Role of Information in Bargaining: An Experimental Study.” *Pschological Review.* 86, 474–94.

Selten, R. and Buchta, J. (1999). “Experimental Sealed Bid First Price Auction with Directly Observed Bid Functions.” In D. Budescu, I. Erev, and R. Zwick (eds.), *Games and Human Behavior: Essays in the Honor of Amnon Rapoport.* Mahwah, NJ: Lawrenz Associates.

Selten, R., Sadrieh, A., and Abbink, K. (1999). “Money Does not Induce Risk Neutral Behavior, But Binary Lotteries Do Even Worse.” *Theory and Decision.* 46, 211–249.

Thaler, R.H. (1981). “Some Empirical Evidence on Dynamic Inconsistency.” *Economic Letters.* 8, 201–207.

Wärneryd, K.-E. (1999). *The Psychology of Saving.* Edward Elgar Publishing.