

THE EFFECT OF UNEMPLOYMENT ON CONSUMPTION: AN EXPERIMENTAL ANALYSIS*

Enrica Carbone and John D. Hey

This paper reports on an experiment that investigates the apparently robust phenomenon of over-sensitivity of consumption to current income. Using a particularly simple formulation, we also investigate whether individuals correctly respond to their employment status. We find that subjects over-react. Our data enables us to investigate where this over-sensitivity originates; we conclude that economic agents differ in their ability to plan ahead and understand the dynamic process determining their employment status. However, agents seem able to respond appropriately to *changes* in the parameters governing their decision processes, in that the comparative static predictions of the theory are largely confirmed.

A repeatedly observed phenomenon in the empirical literature on the behaviour of consumption is its over-sensitivity (relative to that prescribed by the relevant theory) to current income. We provide here an experimental investigation of this phenomenon, taking advantage of the experimental method to eliminate factors that are not of interest and concentrate attention on those that are. In particular we deliberately adopt a simple model in which income in any period can take just one of two values. We interpret these two values as corresponding to states of employment and unemployment, and we additionally assume that transition between these two states is governed by a first-order Markov process. This formulation represents an empirically relevant advance on previous experimental work and enables us to discover whether individuals react appropriately to their employment status – or whether they over-respond, as the non-experimental empirical literature would suggest. We find they do over-respond. So the behaviour of our subjects is, in that respect, similar to ‘real-life’ economic agents. Using our data, we are then able to look closer at the behaviour of our subjects and begin to infer where this excess sensitivity comes from. Our analysis leads us to the rather obvious conclusion that agents differ in their ability to optimise correctly and that some agents are bad at taking into account the dynamics of the problem – having much too short a horizon. These subjects quite naturally are over-sensitive to changes in current income – leading to an over-sensitivity of the agents taken collectively.

The paper begins, in Section 1, with a brief account of the theory being investigated. We then describe, in Section 2, how this theory was implemented in the laboratory. We also comment briefly on how our model differs from other experimental work on the life cycle consumption model. Section 3 contains a

* We would like to acknowledge financial support from the European Commission under the TMR grant ‘Savings and Pensions’. We would like to thank participants at the TMR conference held in Tilburg in March 1999, the participants to FUR IX held in Marrakech in June 1999, the participants in the IAREP conference in Belgirate in Jun-July 1999 and participants in the meeting of the ‘Gruppo di Studio per le Teorie e Politiche Economiche’ in Benevento in October 1999. We would also like to thank the Editors of this JOURNAL for their contribution to the improvement of this paper.

detailed analysis of the results; we focus attention on how well the solution to the theory explains the behaviour of the subjects (both in absolute terms and in comparative static terms), and we conclude that generally the theory is lacking in its explanatory power, though its power varies from subject to subject. Clearly subjects differ in their ability to solve the dynamic problem posed in our experiment – this differing ability being apparently a function of their differing ability to take into account the full dynamics of the problem. Some subjects are very good at this, some very bad. They seem to differ in their ability to think ahead – some subjects being able to think ahead right to the horizon of the experiment, whilst others are able only to think one or two periods ahead. We thus spend some time estimating and discussing the apparent horizons of the subjects. Once again, we get the conclusion that subjects are different. This helps us to begin to explain the over-sensitivity of consumption to income. We discuss this, and further implications of our findings, in the concluding Section 4. A significant amount of material, including the experimental instructions, is contained in an Appendix that is available to interested readers from the authors on request. Also available is the experimental data and the software that was used in the experiment itself.

1. Theory

We investigate a particular version of the familiar life cycle model of consumption. To be specific, we consider an individual who lives for a finite number T of discrete periods and who has preferences over consumption that satisfy the discounted utility model (with a per-period utility function $U(\cdot)$ and with a discount rate equal to zero).¹ Because of these assumptions, it follows that the objective of the individual at any point of time t is to maximise the expected value of the (not) discounted lifetime utility of consumption, as given by the expected value of

$$U(C_t) + U(C_{t+1}) + \dots + U(C_T) \quad (1)$$

where C_t denotes consumption in period t . We assume that he or she starts life in period 1, in which he or she is employed, and in each of the T discrete periods of his or her lifetime is either employed or unemployed: with income y if employed and z if unemployed. Future states of employment/unemployment follow a first-order Markov process: if employed (unemployed) in one period the probability of remaining (becoming) employed in the subsequent period is $p(q)$. Clearly, income is *ex ante* risky, but the stochastic process defining future incomes is known. There is a certain and known rate of return r (> 1) per period on all money saved.

The strategy that should be followed by the individual from some period t onwards is to choose a sequence of consumption C_t, C_{t+1}, \dots, C_T that maximises the

¹ We put the discount rate equal to zero for two reasons: (1) given that it always enters into the objective function combined in the same way with the rate of return r , it follows that any desired combination of the discount rate and the rate of return can be achieved with a zero discount rate and an appropriately chosen rate of return; (2) while it is possible in principle to induce real discounting in a laboratory experiment through the device of having a random stopping mechanism – a random horizon – previous results (Hey and Dardanoni, 1988) suggests that this does not work, as subjects misunderstand the stationarity property of such a stochastic process.

expected lifetime (non) discounted utility – as defined above. It is well known that the solution to this problem is found using dynamic programming by backward induction, starting in the final period, (period T), in which it is optimal to consume everything, and then working backwards, finding the optimal consumption function in every period. Without going into details, we note that the optimal consumption in any period t is a function of: (i) the wealth W_t in that period; (ii) the value of t itself; (iii) the employment status in that period; and (iv) the various parameters of the model, p , q , r , y , z and $U(\cdot)$. Once the individual has worked backwards from period $t = T$ to period $t = 1$ and has found these optimal consumption functions, he or she can then work forward, substituting in the actual value of wealth, the actual value of t and the actual employment status, into the optimal consumption function to determine the optimal level of consumption in any period. In principle the problem of choosing consumption through the life cycle is solved.

The practical details, however, are rather more complicated, particularly if we assume (as in the experiment itself) that consumption decisions have to be restricted to (non-negative) integers, that wealth itself is to be rounded to the nearest integer and that wealth cannot become negative.² Moreover, given the finite horizon of the problem, the decision maker cannot use stationarity to simplify the solution of the decision problem. These complications imply that the optimal consumption functions cannot be found analytically but instead have to be calculated numerically. We did this using the software Maple. We therefore found the optimal consumption function numerically, as a function of wealth, in each of the periods $t = 1, 2, \dots, T$, for both employment statuses (employed and unemployed) and for each set of parameters p , q , r and y that we used in the experiment.³ Perhaps not surprisingly,⁴ these optimal consumption functions turned out to be approximately⁵ linear functions of wealth in the appropriate range – so we can approximate the optimal strategy in each instance by a relationship of the following form:

$$\begin{aligned}
 C &= 0 && \text{if } a + bW \leq 0 \\
 C &= a + bW && \text{if } 0 < a + bW < W \\
 C &= W && \text{if } W \leq a + bW.
 \end{aligned}$$

The parameters a and b in the above equation depend on: (i) the employment status of the individual; (ii) the time period t ; and (iii) the parameters of the model: p , q , r and y .

² That is, that borrowing is not allowed.

³ We kept z and $U(\cdot)$ invariant in all treatments of our experiment. We therefore take these as given from now on.

⁴ Given the form of the utility function that we used, which was that of constant absolute risk aversion. See Hey (1980).

⁵ It is not exactly linear because of the fact that consumption was forced to be discrete – so the function is a sort of step function – but it can be seen to be approximately linear if a straight line is drawn through the steps.

The effect of the employment status is simple: as we will see later, the slope b is not affected by the employment status but the intercept a is – with a higher value of a in the employed state than in the unemployed state. In other words, the optimal consumption function when employed is parallel to, and higher than, the optimal consumption function when unemployed (for the same time period and the same parameters). This is something we can test empirically; anticipating somewhat, our results show that the difference between the estimated intercept when employed and the estimated intercept when unemployed is positive (as indeed it should be) but is generally much bigger than theory predicts. That is, our subjects over responded to moving between employment and unemployment.

The way that the parameters a and b depend on t is not simple: when t is small, a and b are roughly independent of t , but as t approaches 25, a approaches 0 more and more rapidly, and b approaches 1 more and more rapidly. Because there was no obviously apparent functional form (and certainly not one we could justify theoretically) relating a and b to t we decided to treat different periods as different and describe the optimal strategy and estimate the actual strategies *period by period*. We then examined, period by period, the relationship between a and b and the parameters p , q , r and y , for each employment state. To do this, we estimated the relationship between a and b and the parameters p , q , r and y , and the dummy variable e by regression analysis, where $e = 0$ indicates unemployment and $e = 1$ indicates employment. We discovered that the parameter a (the intercept of the optimal consumption strategy) depends on p , q , r , y , e , ep , eq , and ey (these latter three variables denoting the interactive terms $e \times p$, $e \times q$, and $e \times y$). No other combinations of p , q , r , y and e have a significant effect on the intercept a and these variables explain virtually all the variation in a . As far as the slope coefficient b is concerned, we discovered that only r has any effect on it, and that r alone explains virtually all of the variation in b .⁶ So a , in each time period, is a linear function of p , q , r , y , e , ep , eq , and ey ; while b , in each time period, is a linear function of r . If we substitute these back into the linear consumption function $C = a + bW$ above, we conclude that the optimal consumption strategy in each period takes the following form (when it implies a C value between 0 and W):

$$C = (a_0 + a_1p + a_2q + a_3r + a_4y + a_5e + a_6ep + a_7eq + a_8ey) + (b_0 + b_1r)W \quad (2)$$

that is, it is linear in p , q , r , y , e , ep , eq , ey , W and rW (where rW denotes $r \times W$). Table 2 reports the values of the coefficients a_0 , a_1 , a_2 , a_3 , a_4 , a_5 , a_6 , a_7 , a_8 , b_0 and b_1 . The following might usefully be noted:

- (i) the coefficients on p and on ep are both positive and decreasing in t : thus an increase in p – the probability of remaining employed – leads to an upward shift in the consumption function (one that is larger when employed), though one that is decreasing as the horizon approaches;
- (ii) the coefficients on q and on eq are respectively positive and negative and are both decreasing in t : thus an increase in q – the probability of becoming employed – leads to an upward shift in the consumption function (one that

⁶ Some hints as to why this should be so can be found in Hey (1980).

is smaller when employed), though one that is decreasing as the horizon approaches;

- (iii) the coefficients on y and on e_y are both positive and very slightly decreasing in t : thus an increase in y – the income when employed – leads to an upward shift in the consumption function (one that is bigger when employed), though one that is decreasing slightly as the horizon approaches;
- (iv) the coefficient on r by itself is negative whereas the coefficient of rW is positive – so as r rises there is a lowering of the intercept and an increase in the slope of the consumption function – leading to a form of rotation: the increase in the rate of return gives the individual the incentive to save more at low levels of wealth and the opportunity to consume more at high levels of wealth.

We now turn to a discussion of the experimental implementation of the above model.

2. The Experiment Implementation

It goes without saying that we cannot reproduce the life cycle model in real time in the laboratory. Instead we do as other experimentalists do and implement a decision problem with exactly the same structure and the same incentive mechanism. Mathematically, the life cycle model discussed in Section 1 and the experiment that we implemented have the same structure. Whether this means that we have tested the life cycle model in the laboratory is something that can be discussed – and we make some comments on that crucial point in the closing sentences of this Section.

We told subjects that they were taking part in an individual decision problem that would last 25 (we denote this here by T) rounds or periods, after which the experiment would terminate. At the beginning of each and every period, the subject would receive an income, denominated in *tokens*. This token income would, each period, take one of two values, which we denote here by y and z . The actual values of y and z throughout the T periods of the experiment would be determined by a first-order Markov process. Specifically, subjects were told:⁷ if the token income in one period was y , then the probability of it being $y(z)$ next period would be $p(1 - p)$; if the token income in one period was z , then the probability of it being $y(z)$ next period would be $q(1 - q)$. Tokens income could be accumulated to provide wealth, and between periods wealth would earn interest at the rate $r - 1$ (where here r denotes the rate of return on savings). The subjects were told that, after learning their income in tokens in each period, they had to then take a decision – concerning how much of their accumulated token wealth they wished to *convert* into money in that period; we denote here the amount consumed by C . Subjects were informed of the *conversion scale* (from tokens into money); this we denote by $U(\cdot)$. Thus C tokens converted into money yielded $U(C)$ in money. Subjects were told that at the end of the experiment they would be paid in cash the

⁷ Obviously not in these technical terms but in terms accessible to laymen. The full set of instructions is in an Appendix, which is available to interested readers on request.

total amount of money converted from tokens over the T periods of the experiment. However, any unconverted tokens remaining at the end of the experiment would be worthless.

As the first-order Markov process governing the transition between employment and unemployment was a key innovative feature of this experiment, and as misperception of it by the subjects would imply excess consumption volatility, we spent considerable design time ensuring that subjects understood the nature of this process. Appendix Figure 2 shows the visual display that we used in the experiment: after each period, subjects had to nominate 1 card out of 8. The income (and hence the employment status) in the subsequent period was written on the reverse of each card. After nominating a card and having its reverse side revealed, all the reverse sides were shown to the subjects. It was very clear from this that it was more likely that the subject would be employed in the subsequent period if he or she was employed in the present period than if he or she was unemployed. We think that subjects were very clear about the nature of this stochastic process.

Given this payment scheme, it follows that the earnings of a subject over periods t to T of the experiment would be determined by the value of

$$U(C_t) + U(C_{t+1}) + \dots + U(C_T) \quad (3)$$

which implies that their expected earnings would be maximised if they adopted a strategy which maximised the expected value of the expression given in (3) for all t (including $t = 1$). It will be noted that this is exactly the same as that given by (1) above. Thus the objective function of the subjects in our experiment was exactly that of the life cycle model described in Section 1, under the assumption that the subjects were risk neutral (and hence interested in maximising their expected payment for taking part in the experiment). We justify the assumption of risk-neutrality on the grounds that the subjects were paid on all T periods of the experiment and that the amounts of money involved in each period were relatively small.

The experiment itself was computerised⁸ and subjects performed the experiment at individual PCs, proceeding individually at their own pace. Ninety-six undergraduate and postgraduate subjects from the University of York participated in the experiment. On registration, all participants were sent a general set of instructions,⁹ giving them general information about the task involved in the experiment but without giving any particular parameter values. On arrival at their experimental session, participants were then given specific instructions for the parameter set that they would be facing. There was then a general briefing session and the experimenter then showed them a simplified version of the experiment (using computer projection from the experimenter's computer). This simplified version lasted just two periods, and was intended to familiarise the subjects with the screens that they would face, and how the software worked. There was then an opportunity to ask clarificatory questions and then the subjects commenced the experiment and completed it at their own pace.

⁸ The program is available on request from the authors.

⁹ These are in the Appendix available on request from the authors.

We ran 16 different treatments, which differed in terms of the parameters used. The relevant parameters are: (i) the two income values y and z ; (ii) the probability p of remaining employed; (iii) the probability q of becoming employed; and (iv) the rate of return r ; (v) the parameter¹⁰ R of the conversion function, which took the particular functional form of the constant absolute risk averse utility function $U(c) = [1 - \exp(-Rc)]$. Throughout all treatments of the experiment we put $R = 0.015$, as our primary interest was not in seeing how changes in risk aversion affected behaviour but rather in how changes in p , q , r and y affect behaviour.¹¹ Accordingly we took two different values¹² for each of p , q , r and y and ran 16 different treatments corresponding to the 16 different combinations possible. Table 1 gives details of these 16 parameter sets. It should be noted, of course, that for each and every subject, his or her parameters were fixed throughout all T periods of the experiment; the variation in the parameter values over treatments was *across* subjects.

It might be useful to note that some experimental work testing the life cycle model has already been carried out. The principal references are Hey and Dardanoni (1988), Kim (1989), and Ballinger *et al.* (2003). Our model differs from the others in that it concentrates attention on the responsiveness of consumption to the employment/unemployment process, modelled as a Markov

Table 1
Parameter Sets

Set	p	q	r	y
1	0.875	0.375	1.4	30
2	0.875	0.375	1.2	30
3	0.875	0.125	1.4	30
4	0.875	0.125	1.2	30
5	0.625	0.375	1.4	30
6	0.625	0.375	1.2	30
7	0.625	0.125	1.4	30
8	0.625	0.125	1.2	30
9	0.875	0.375	1.4	15
10	0.875	0.375	1.2	15
11	0.875	0.125	1.4	15
12	0.875	0.125	1.2	15
13	0.625	0.375	1.4	15
14	0.625	0.375	1.2	15
15	0.625	0.125	1.4	15
16	0.625	0.125	1.2	15

Key: p : probability of remaining employed.

q : probability of becoming employed.

r : rate of return.

y : income when employed.

¹⁰ The parameter α affects the payment to the subjects and not their behaviour (under the assumption of risk neutrality).

¹¹ Note it is the *relative* values of y and z that affect behaviour, so we kept z constant throughout all treatments.

¹² The choice of the particular values was subject to several considerations. In particular, p had to be always greater than q and the probabilities should be ones that would be easily understood by subjects.

process. In contrast the other studies assume an i.i.d. income generating process. We think ours is more realistic. Ballinger *et al*, however, investigate the effect of social learning on behaviour. As such, that paper and ours complement each other. They are both, however, subject to the usual criticism of experimental work, particularly work that tries to reproduce in the laboratory, in a short period of time, a complicated problem that is tackled in the real world over a long period of real time. We appreciate that there are difficulties here but note that there are already good methodological discussions of this and related issues (the interested reader is referred to Kagel and Roth (1995) and to references therein). We also note that, whatever the context, we are testing the theory as mathematically defined: if our experiment is unrealistic, then so is the theory that it is testing.

3. Analyses of the Results

We break down the analysis of the results into 4 sub-sections. In sub-section 3.1, using regression analysis, we estimate the actual consumption strategies apparently employed by the subjects and then provide a comparison of these estimates with the optimal strategies. As will become clear there seem to be strong differences between the actual strategies employed and the optimal. In order to shed some light on the possible causes of these differences, we present, in sub-section 3.2, a direct analysis of the differences between the actual and the optimal. As will be seen, there are two ways we can do this: first, by comparing the actual consumption with what would be optimal given the wealth they actually had in that period; second, by comparing the actual consumption with what would be optimal given the past income stream of the subject and assuming optimal behaviour throughout. We discuss in sub-section 3.2 what exactly we mean by this. Then, in sub-section 3.3, we concentrate particular attention on the important variables and parameters, asking the question whether subjects are responding, at least qualitatively, correctly to changes in the relevant variables and parameters. This is effectively a test of the comparative static properties of the subjects' behaviour. We do this in two different ways, which we describe in detail in that sub-section. Finally, in sub-section 3.4, we try to work out what strategies the subjects are using, given that in general they are not using the optimal strategy. We work on the assumption that subjects generally are not able to solve the problem optimally but instead adopt a simplified heuristic of acting on the basis of a shorter horizon than is actually the case.¹³ That is, we assume that subjects have a subjective horizon, less than the true one, and act on the basis that the actual horizon is this subjective horizon. Obviously this induces dynamically inconsistent behaviour but it allows subjects to simplify the rather complex problem that they face. We explain this in more detail in sub-section 3.4. Anticipating somewhat, our results show that some subjects have very short horizons while others have longer horizons.

¹³ This is what Deaton (1992, pp. 156–7) infers from real consumption data.

3.1. Comparison Between Estimated Actual and Optimal Consumption Strategy

We have discussed in Section 1, and report in (2), the form of the optimal consumption strategy in every period. We used the same functional form to try to explain the actual consumption behaviour of subjects. We took note of the fact that consumption is bounded between 0 and W and therefore carried out Tobit regressions, period by period, of the actual consumption on the variables in (2), namely $p, q, r, y, e, ep, eq, ey, W$ and rW . Obviously here W denotes the actual wealth of the subject in that time period. The results are reported in the Table in the Appendix (which is placed in the Appendix because it is too detailed to print in the text). This Table reports the coefficients in the regressions explaining *optimal* consumption (these are the same as in Table 2 in the text of this paper), the corresponding estimated coefficients (of the above-mentioned 11 variables) in the Tobit regressions of *actual* consumption and the standard deviations of the estimated parameters for the actual consumption regression. This Table also reports whether an estimated coefficient is significantly different from zero (at the 5% level of significance). The results are that the coefficients of W and rW are significantly different from zero while all the other parameters are only occasionally significantly different from zero.

We also report in this Table in the Appendix whether an estimated coefficient is significantly different from zero (at 5%) *and* not significantly different from the optimal coefficient (again at 5%). The results of this analysis is that the coefficient

Table 2
Coefficients of the Optimal Consumption Strategy

t	$const$	p	q	y	r	e	ep	eq	ey	W	rW
2	-59.38	6.37	15.31	0.28	-2.04	-5.48	8.63	-7.66	0.16	-0.53	0.59
3	-55.13	6.19	17.03	0.30	-5.65	-4.47	8.48	-8.91	0.14	-0.53	0.59
4	-51.84	7.06	16.09	0.27	-7.60	-5.68	8.58	-8.98	0.19	-0.53	0.58
5	-46.51	5.82	15.29	0.28	-10.77	-6.15	8.91	-7.44	0.18	-0.52	0.58
6	-41.49	7.22	15.11	0.29	-15.32	-4.14	7.82	-8.06	0.15	-0.52	0.58
7	-37.33	6.88	16.62	0.30	-18.37	-3.65	7.84	-9.14	0.14	-0.52	0.57
8	-27.96	5.65	15.06	0.26	-23.38	-7.34	10.56	-8.32	0.19	-0.51	0.57
9	-24.54	6.43	16.10	0.27	-26.19	-6.34	9.39	-9.32	0.19	-0.50	0.56
10	-17.17	6.48	16.02	0.28	-31.28	-5.55	8.73	-9.50	0.18	-0.49	0.55
11	-8.84	5.85	15.49	0.28	-36.52	-6.31	9.75	-8.52	0.18	-0.48	0.55
12	-0.54	5.44	15.29	0.28	-41.79	-5.65	9.63	-8.79	0.16	-0.46	0.54
13	7.11	5.91	15.58	0.27	-46.87	-5.61	9.41	-9.08	0.16	-0.45	0.53
14	14.99	5.87	15.52	0.27	-51.64	-5.51	9.05	-8.89	0.17	-0.43	0.52
15	23.41	5.49	15.01	0.27	-56.40	-5.78	9.68	-8.26	0.16	-0.41	0.50
16	30.44	5.25	15.59	0.26	-60.06	-6.36	10.36	-9.57	0.18	-0.38	0.49
17	37.15	5.18	15.10	0.26	-63.08	-5.82	9.67	-9.41	0.18	-0.35	0.47
18	42.38	4.92	15.37	0.25	-64.79	-5.66	10.03	-10.37	0.17	-0.32	0.45
19	46.73	3.87	15.19	0.24	-64.83	-6.41	10.87	-10.17	0.19	-0.28	0.42
20	47.86	3.65	14.44	0.23	-62.33	-6.43	10.74	-10.33	0.19	-0.22	0.39
21	46.67	3.36	13.72	0.21	-57.63	-6.61	10.58	-10.35	0.22	-0.15	0.36
22	43.31	1.81	12.45	0.18	-49.97	-7.11	11.52	-10.72	0.22	-0.06	0.32
23	34.35	1.04	10.56	0.14	-37.89	-7.05	11.01	-9.58	0.23	0.07	0.27
24	20.85	0.00	6.79	0.09	-21.44	-5.32	8.02	-6.79	0.21	0.32	0.19
25	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00

of the variable W in the estimated actual regressions has this property just five times. This means that the coefficient of W in the explanation of actual consumption is generally significantly different from the optimal. The variable rW is significantly different from zero and at the same time not significantly different from the optimal just seven times.

A comparison of the estimated coefficients with the optimal ones can be made using this Table¹⁴ in the Appendix. One thing that is clear from the column giving the coefficients on the parameter p that the signs are usually correct, but the magnitudes vary considerably – mainly a consequence of the fact that the regressions are carried out period by period. However, there is a problem of ascertaining whether a particular parameter, or the employment status, has the correct influence on consumption behaviour, as all the parameters and the employment status enter the equation twice – once by themselves and once in interaction with another variable. To get round these problems, we present a more systematic investigation of the ‘comparative static’ effects in sub-section 3.3. But for the record, before we leave this Table in the Appendix, it may be useful to give the following summary:

- (i) the variable p by itself usually has the correct positive sign but it is occasionally (4 times out of 23)¹⁵ negative; moreover the magnitude of the estimated coefficient varies considerably;
- (ii) the variable q by itself has similarly properties, though here the incidence of incorrect negative signs is greater (8 times out of 23), as is the variability of the magnitude of the estimated parameters;
- (iii) the variable y by itself has the same incidence of incorrect negative signs (8 out of 23) as the variable q by itself – but the variability of the estimated coefficients is much lower;
- (iv) the variable r by itself usually has the correct negative sign but there are 6 time periods in which it is positive;
- (v) the variable e by itself has an incorrect positive sign 10 times out of 23 and the variability of the estimated coefficient is extremely high;
- (vi) the interactive term ep has an incorrect negative sign 8 times out of 23, and again the magnitude is highly variable;
- (vii) the interactive term eq has an incorrect positive sign 11 times out of 23, and again the magnitude is highly variable;
- (viii) the interactive term ey has an incorrect negative sign 7 times out of 23, but the variability of the estimated coefficient is not particularly high
- (ix) the variable W by itself has an incorrect positive sign just twice and usually the magnitude of the estimated coefficient is close to that of the optimal;
- (x) the interactive variable rW always has the correct positive sign and the coefficients are not too far (but see later) from the optimal.

¹⁴ This could also be used to check whether behaviour is improving over time, though it is difficult to come to a definite conclusion as there is considerable variability in the estimates.

¹⁵ We restrict attention to periods 2 to 24 – as we have explained already estimation is impossible in period 1; in contrast the estimates on period 25 are wildly distorted by those subjects who did not consume all their wealth in that period.

This discussion highlights the fact that it is difficult from this type of analysis to discover whether subjects are responding correctly to the various variables – because they all appear both by themselves and in some interactive term. We provide in sub-section 3.3 a more satisfactory test of the comparative static predictions. In the meantime, we look at the errors made by the subjects in the experiment.

3.2. *An Analysis of the Difference Between Actual and Optimal Consumption Behaviour*

It is clear from the above that there is a marked divergence between actual behaviour and the optimal. To try to understand what is driving this difference, in this sub-section we present an analysis of the determinants of the difference between actual and optimal consumption. However, we need first to specify what we mean by optimal consumption in the context of a situation where subjects make mistakes in determining their consumption at some stage in their life cycle. There are two possible ways that we could define optimal consumption in some period. The first way is to define optimal consumption as the consumption that would be optimal in that period given the wealth that the individual actually has in a particular period. However, if the individual has not behaved optimally in the past then the wealth that he or she has in that period is *not* the level of wealth that would be optimal given the actual income stream that the individual received in the past. This leads us to the second definition. This is the level of consumption that would be optimal in that period if the wealth of the individual was that which would have been the result of consuming optimally in the past – given the actual income stream received. We have therefore two kinds of ‘error’ – the first being the difference between actual consumption and the optimal consumption defined in the first way (what would be optimal in that period given the wealth of that period) – and the second being the difference between actual consumption and the optimal consumption defined in the second way (what would be optimal in that period given optimal behaviour in the past). Ballinger *et al.* (2003) also consider these two different definitions. They comment that the first type of error ‘...measures only current period deviations from the optimal policy – given that the optimal policy chooses on the basis of subject cash-on-hand¹⁶ – and because of this, we prefer this measure of error.’ (Ballinger *et al.*, 2003, p 17). We agree with them but, like them, use both definitions of error in the analysis that follows.

Let us denote the first kind of error by e_1 and the second by e_2 . If subjects are behaving optimally then these errors will always be zero. If they are behaving almost optimally (that is, they implement the optimal strategy with a ‘tremble’) then these errors should be white noise. That is, no variable should have a significant effect on them. If one does, then this variable is explaining some part of the deviation from optimality. To investigate the possibilities we regress each of these separately against the variables used as independent variables in Table 2 *and* the period number (t). Obviously we can pool all periods together to do this analysis. The results of the regressions are in Table 3.¹⁷ First, we ran the regressions on the full sample. It can

¹⁶ Wealth in our terminology.

¹⁷ These are OLS estimates. We also ran Tobit regressions, and these are almost identical.

Table 3
Regressions Explaining the Error in Consumption

Sample	depvar	const	<i>p</i>	<i>q</i>	<i>y</i>	<i>r</i>	<i>e</i>	<i>ep</i>	<i>eq</i>	<i>ey</i>	<i>W</i>	<i>rW</i>	<i>t</i>
Full	<i>e</i> ₁	-521	94.1	-32	1.70	351	-4.06	14.0	118	-1.08	1.10	-1.08	2.21
		-6.3	1.5	0.6	1.8	7.0	0.1	0.2	1.6	0.9	7.2	9.7	3.4
Restricted	<i>e</i> ₁	51.8	49.7	13.7	0.251	-56.5	-5.2	0.135	-8.5	0.604	-1.09	0.762	0.082
		1.7	2.1	0.7	0.7	2.9	0.2	0	0.3	1.4	15	13	0.3
Full	<i>e</i> ₂	663	-9.13	4.42	-1.35	-445	-8.86	27.5	16.3	-0.89	0.359	-0.156	-12.6
		18	0.3	0.2	3.2	20	0.3	0.8	0.5	1.7	5.4	5.2	-44
Restricted	<i>e</i> ₂	856	-58.8	-0.758	-1.48	-567	-15.1	36.8	-12	-0.438	-0.349	0.438	-13.2
		29	2.7	0.04	4.5	31	0.6	1.3	0.5	1.0	5.0	8.1	59

Key: the first row gives the estimated coefficients.
the second row the *t*-values.
italicised entries indicate estimated coefficients significantly different from zero.

be seen from Table 3 that the coefficients that are significant are those of the constant, r , W , rW and t . (Obviously here W refers to the actual wealth.) The fact that the coefficient on t is significant and positive means that the magnitude of the error increases through time – effectively that the subjects were relatively under-consuming (relative to the optimal strategy) early on in the experiment and relatively over-consuming (relative to the optimal strategy) later on in the experiment. Care should be taken in interpreting this result – as account should also be taken of the effect of wealth on the error. If the first row of Table 3 is examined we see that the coefficient on W is 1.10 and that on rW is -1.08 . Both are highly significant. Now recall that r is either 1.2 or 1.4, so in either case the coefficient of W in this first error equation is always negative (-0.196 when r is 1.2 and -0.412 when r is 1.4). This means that the error declines as wealth rises – or for low values of wealth the subjects over-consume and for high values of wealth they under-consume. It is clear that many subjects did not save enough (that is they over-consumed) in the early stages of the experiment when their wealth was low, and then they under-consumed later in the experiment as they approached the horizon (as a consequence of over-consuming earlier). When we exclude from the sample those subjects who saved an excessive amount (that is, more than 3,000) then the coefficient on p becomes significant, whilst the coefficients of the constant, r , W , rW and t remain significant. As we have already discussed, it is difficult to use this kind of analysis to help us understand what the subjects were doing – but it does enable us to conclude that the subjects do not correctly respond to their wealth, the rate of return and the time period (and possibly the parameter p – the probability of remaining employed).

When we repeat the analysis of the error using the second definition of the error, we see from Table 3 that the same coefficients are significant but so also is the parameter y (both using the full sample and the restricted sample). The sign is negative – indicating that subjects do not take the effect of an increase in the level of their employment income into account correctly – they do not respond sufficiently to a higher employed income.¹⁸

3.3. Comparative Static Analyses

In this Section we report the results of various comparative static analyses. To be specific, we look at how optimal consumption and actual consumption change when the parameters r , p , q and y change, and we focus particular attention on the effect of the employment status on consumption. We begin however with Table 4, in which we report, period by period, the coefficient on wealth in both the optimal consumption strategy and the estimated consumption strategy, for the two values of the rate of interest used in the experiment. The first column is the period; the second and third columns the coefficient on wealth in the optimal and actual consumption functions for $r = 1.2$; and the fourth and fifth columns the coefficient on wealth in the optimal and actual consumption functions for $r = 1.4$. For example, taking the second column, we see that for the low rate of interest (20%)

¹⁸ We should, of course, note that the comparative static analyses of the parameters p , q , y and r are across subjects – since each subject faced just one set of parameters. But the effect of the employment status e on consumption is both across all subjects and within subjects.

Table 4
Effect of Wealth on Consumption (both estimated and optimal)

Period	Rate of interest = 20%		Rate of interest = 40%	
	Optimal	Actual	Optimal	Actual
2	0.17	-0.70	0.29	-0.47
3	0.17	-0.48	0.29	-0.15
4	0.17	-0.39	0.29	0.04
5	0.17	-0.21	0.29	-0.04
6	0.17	-0.17	0.29	0.02
7	0.17	-0.06	0.29	-0.03
8	0.17	-0.02	0.29	-0.00
9	0.17	0.04	0.29	0.29
10	0.18	0.03	0.29	0.07
11	0.18	0.03	0.29	0.42
12	0.18	0.04	0.29	0.42
13	0.18	0.05	0.29	0.12
14	0.19	0.06	0.29	0.17
15	0.19	0.08	0.29	0.21
16	0.20	0.11	0.30	0.19
17	0.21	0.13	0.30	0.28
18	0.22	0.13	0.31	0.31
19	0.23	0.16	0.32	0.32
20	0.25	0.17	0.33	0.28
21	0.28	0.17	0.35	0.41
22	0.32	0.22	0.39	0.55
23	0.40	0.20	0.45	0.26
24	0.55	0.02	0.58	0.55
25	1.00	-0.10	1.00	-0.90
Averages	0.22	-0.02	0.32	0.18

the optimal marginal propensity to consume out of wealth starts at 0.17 in period 2 (the same figure is valid for period 1), and stays around this level until period 9. It then rises at an increasingly faster rate – until it reaches 1.0 in the final period (when the individual optimally consumes all his or her residual wealth). In contrast, in the third column are the estimated marginal propensities to consume out of wealth at the interest rate of 20%. This Table contains some interesting information: if we start by comparing the second and fourth columns we see that the optimal consumption strategy has the property that the marginal propensity to consume out of wealth is greater in all periods for a higher rate of interest – if the rate of interest rises people should optimally consume more in all periods. If we now compare the third and fifth columns we see that this is generally true of the actual behaviour – the marginal propensity to consume is generally higher at the higher rate of interest. However, in comparing the second column with the third, and in comparing the fourth column with the fifth, the actual marginal propensities to consume are generally lower than the optimal – so behaviour is absolutely wrong but the comparative static predictions are correct. Furthermore the actual marginal propensities to consume show some bizarre effects – actually being negative in the early periods. This is a consequence of the regressions being across all subjects – some subjects with low wealth consumed all their wealth, while others were deliberately building up their wealth stocks in the early periods – thus leading

to a negative marginal propensity to consume in the early periods. But having said this, we note from Table 4 that the actual marginal propensities to consume rise through time – as they should do. There is also some slight evidence that the actual coefficients are getting closer to the optimal as time passes – that is, that the subjects are learning about the problem to be solved. The final row of Table 4 reports the average¹⁹ marginal propensities to consume out of wealth – once again we see that the actual are below the optimal but move in the correct direction when the interest rate rises. In the subsequent analysis we just report and discuss these averages – since there is a high variability in the period-by-period results.²⁰

In Table 5 we present several analyses that shed light on the comparative static effects. Specifically, we look at the effect of e , p , q , y and r on consumption. We begin with the first of these – the effect of e on consumption. As the variable e enters the equation interactively with p , q and y we need to consider the various possible combinations that these may take. In Table 5(a) there are the eight possible combinations. We see that the average effect on optimal consumption of a move from unemployment to employment is always greater when y takes its high value than when it takes its low value. This is also true for the actual effects but the magnitude of the actual effect is everywhere much higher – for example when y is low and both p and q are high, moving from unemployment to employment consumption should increase on average by 1.81 – however, in fact, it actually increases on average by 12.00. Subjects over-react to being employed – they increase their consumption excessively. Thus our subjects are excessively sensitive to their employment status. This reflects the stylised fact that we reported earlier.

In Table 5(b) we look at the effect of p on the consumption level. Because p enters the equations not only by itself but interactively with e , we need to consider the two possible values of e . When e is zero (one) – that is, the subject is unemployed (employed), moving from the low value of p to the high value should on average increase consumption by 5.03 (14.57). In fact the actual average consumption increases by 23.64 (39.89). So subjects are considerably over-reacting to an increase in p – it seems that subjects are not able to take into account correctly the effect of p on their situation – a high value of p induces excess optimism, a low value of p induces excess pessimism. A different story is apparent however in the effect of q on behaviour: as Table 5(c) shows, an increase in q from its low value to its high value should increase consumption on average by 14.73 (5.68) when the subject is unemployed (employed) whereas it actually decreases it by 1.08²¹ (increases it by just 0.15). Here subjects are under-reacting to an increase in q – possibly they are over-pessimistic for high value of q and over-optimistic for low values of q . The asymmetric effects of p and q are particularly interesting.

¹⁹ The average is over periods 2 to 24. We exclude period 25 in all the subsequent analyses since significant distortions are evident in the final period because not all subjects consumed their entire wealth (partly because of a bug in the experimental software which stopped subjects consuming more than 999 in any one period).

²⁰ If we had estimated the regressions over all periods, taking into account the way that coefficients should change, then we could have reduced this variability, but it would have introduced further complications – caused by the necessary inclusion of many interact terms in the regression equations. This would have made the subsequent analysis of the results particularly difficult.

²¹ Though this is not significantly different from zero.

Table 5
Average Effect of Other Parameters and Employment Status on Consumption
 (both estimated and optimal)

<i>(a) average effect of employment status on consumption</i>				
Probabilities	Employed income low (15)		Employed income high (30)	
	Optimal	Actual	Optimal	Actual
p high and q high	1.81	12.00	4.52	19.85
p high and q low	4.07	11.69	6.78	19.54
p low and q high	-0.57	7.94	2.13	15.79
p low and q low	1.69	7.63	4.40	15.48
<i>(b) average effect of probability of remaining employed on consumption</i>				
Unemployed: Optimal	Actual	Employed: Optimal	Actual	
5.03	23.64	14.57	39.89	
<i>(c) average effect of probability of becoming employed on consumption</i>				
Unemployed: Optimal	Actual	Employed: Optimal	Actual	
14.73	-1.08	5.68	0.15	
<i>(d) average effect of employment income on consumption</i>				
Unemployed: Optimal	Actual	Employed: Optimal	Actual	
0.25	0.24	0.43	0.76	
<i>(e) average effect of rate of return on consumption</i>				
Wealth = 100: Optimal	Actual	Wealth = 200: Optimal	Actual	
59.71	139.82	156.62	340.51	

Key: table entries show average effect (over periods 2 to 24) of

- (a) moving from unemployment to employment.
- (b) changing the value of p from its low value to its high value (0.625 and 0.875).
- (c) changing the value of q from its low value to its high value (0.125 to 0.375).
- (d) changing the value of y from its low value to its high value (15 and 30).
- (e) changing the value of r from its low value to its high value (1.2 and 1.4).

In contrast, if we look at Table 5(d) we see that subjects respond almost correctly to a change in the employed income y . An increase of y from its low value to its high value should increase consumption on average by 0.25 (0.43) when unemployed (employed) – the actual average increase is 0.24 (0.76). One might conclude from this that subjects find y (a fixed number) easier to think about than p and q (probabilities). However this conjecture seems to be refuted by Table 5(e) which shows that subjects over-react significantly to changes in the rate of interest. Perhaps this reflects the fact that people under-estimate the effects of compound interest?

An alternative way of thinking about these comparative static effects is presented in Tables 6, 7, 8 and 9. Whereas in Table 5 we used the results of the regression analyses, Table 6 to 9 present a descriptive exercise, in which we present the average (both actual and optimal) consumption over all periods for each parameter individually: in Tables 6, 7, 8 and 9 for the parameters r , y , p and q

Table 6
Period-by-period Effect of the Rate of Return on Consumption

Period	Low					High				
	(1)	(2)	(3)	1-2	1-3	(1)	(2)	(3)	1-2	1-3
1	8	0	0	8	8	10	0	0	10	10
2	10	0	0	10	10	12	0	0	12	12
3	12	0	0	12	12	16	0	0	16	16
4	14	0	0	14	14	19	1	6	18	13
5	12	0	0	12	12	21	2	16	19	5
6	12	1	3	11	9	21	5	33	16	-12
7	14	2	6	12	8	22	10	55	12	-33
8	17	3	12	13	5	29	18	78	11	-49
9	21	6	20	16	1	56	28	100	28	-44
10	22	8	31	14	-9	47	32	123	15	-76
11	23	11	43	12	-20	68	41	145	26	-78
12	27	16	55	12	-28	64	46	168	17	-105
13	30	20	68	10	-37	59	55	191	4	-132
14	33	26	80	8	-47	82	71	214	10	-133
15	41	32	93	9	-52	88	84	237	4	-148
16	43	38	105	6	-61	100	99	260	1	-160
17	52	45	117	7	-65	124	117	283	7	-160
18	56	52	129	4	-73	130	131	306	-2	-176
19	63	60	141	3	-78	139	144	328	-5	-189
20	67	69	153	-1	-85	151	160	351	-9	-200
21	78	80	165	-2	-88	174	179	373	-5	-199
22	94	92	179	2	-85	180	196	396	-16	-216
23	95	103	191	-8	-97	166	199	418	-33	-252
24	65	116	203	-50	-137	198	221	440	-23	-243
25	118	138	217	-19	-99	193	218	463	-25	-269

(1) – average actual consumption.

(2) – average optimal consumption given the wealth of the period.

(3) – average optimal consumption given the past income stream.

1-2 – difference between (1) and (2).

1-3 – difference between (1) and (3).

respectively. In each Table we give averages for the low value of the parameter and for the high value – and for each we give the average actual and the average optimal, calculated in the two different ways that we have already discussed: (2) is the optimal consumption given the wealth they actually had that period; (3) is the optimal consumption if they had behaved optimally throughout.

The tables are structured as follows: the first column indicates the period, then we have two sets of five columns – the first five for the low value of the relevant variable and the second five for the high value of the variable. The five columns contain: (1), the average (across subjects) of actual consumption; (2), the average (across subjects) of the optimal consumption using the first definition of optimal consumption (that based on the wealth they actually had in that period); (3), the average (across subjects) of the optimal consumption using the second definition of optimal consumption (that based on the income stream that they had received to date); 1-2, the difference between actual and optimal according to the first definition; 1-3, the difference between actual and optimal according to the second definition.

Table 7
Period-by-period Effect of the Employed Income on Consumption

Period	Low					High				
	(1)	(2)	(3)	1-2	1-3	(1)	(2)	(3)	1-2	1-3
1	5	0	0	5	5	12	0	0	12	12
2	6	0	0	6	6	16	0	0	16	16
3	9	0	0	9	9	19	0	0	19	19
4	13	0	0	13	13	20	1	6	19	14
5	12	0	0	12	12	21	2	16	20	5
6	12	0	7	12	6	20	5	30	15	-9
7	13	2	17	11	-4	23	10	44	13	-21
8	20	5	28	14	-9	26	16	61	10	-35
9	39	10	42	29	-3	38	23	77	14	-40
10	24	9	58	15	-34	45	31	95	14	-51
11	24	14	75	10	-51	67	39	113	28	-47
12	25	21	92	4	-67	66	41	131	25	-65
13	36	30	110	5	-74	54	45	149	8	-95
14	53	40	127	13	-74	62	56	167	5	-106
15	62	47	144	15	-82	67	69	186	-2	-119
16	69	54	162	15	-93	75	83	203	-9	-128
17	73	60	179	13	-106	102	101	221	1	-119
18	84	66	196	18	-112	101	117	238	-16	-137
19	73	71	214	2	-141	130	133	255	-4	-125
20	88	82	231	6	-143	130	146	273	-16	-142
21	98	92	248	6	-150	154	167	290	-13	-136
22	106	99	266	7	-160	168	190	309	-22	-141
23	99	104	283	-5	-184	162	198	326	-36	-165
24	90	115	300	-26	-210	174	221	343	-48	-170
25	130	130	317	-0	-187	181	225	363	-44	-181

Notes: see Table 6.

Table 6 reports the results for the interest rate. It is clear from this Table that individuals over-consume for the first 9 periods and under-consume thereafter. An increase in the interest rate seems to increase the difference between the actual consumption and the optimal consumption (calculated with both definitions) so to begin with they consume more than they should and afterwards in order to compensate they have to consume increasingly less. However, the comparative static prediction is verified – when the rate of interest increases, individuals should and do consume more everywhere.

In Table 7 we repeat the analysis for the parameter γ . An increase in γ should increase consumption everywhere – and it does. However, the absolute magnitudes of consumption are incorrect.

Table 8 reports on the effect on consumption of a change in the parameter p . From the results it appears that the increase in p does have a noticeable effect on consumption (in the direction indicated by the theory) and increases the magnitude of the errors made by the subjects.

Table 9 reports on the effect on consumption of an increase in the parameter q . We see that when q increases the actual consumption for the first 7 periods of the experiment decreases and the difference between actual consumption and optimal consumption (according to both definitions) decreases. However, after the

Table 8
Period-by-period Effect of the Probability of Remaining Employed on Consumption

Period	Low					High				
	(1)	(2)	(3)	1-2	1-3	(1)	(2)	(3)	1-2	1-3
1	9	0	0	9	9	9	0	0	9	9
2	8	0	0	8	8	14	0	0	14	14
3	10	0	0	10	10	18	0	0	18	18
4	13	0	1	13	12	20	1	5	19	15
5	12	1	5	11	7	22	1	11	20	11
6	13	2	13	10	-0	20	3	24	17	-3
7	14	5	24	9	-9	22	6	38	16	-16
8	20	10	37	11	-16	25	11	53	14	-28
9	25	15	50	10	-25	52	18	70	33	-18
10	30	20	67	10	-37	39	20	87	19	-48
11	50	27	83	23	-33	40	26	105	15	-64
12	50	28	101	21	-51	41	33	123	8	-82
13	39	31	118	8	-79	50	45	140	5	-90
14	46	40	136	6	-91	69	57	158	12	-89
15	53	50	154	3	-101	76	66	176	10	-100
16	60	61	171	-1	-111	83	76	193	7	-110
17	72	75	189	-3	-118	104	87	211	17	-107
18	78	88	206	-10	-128	107	95	228	12	-121
19	93	99	223	-6	-130	109	105	246	4	-136
20	86	111	241	-26	-155	133	117	263	16	-130
21	110	135	258	-25	-148	142	124	280	18	-138
22	127	159	276	-33	-149	147	129	299	18	-152
23	120	171	292	-51	-173	141	131	317	10	-176
24	139	202	309	-62	-170	124	135	334	-11	-210
25	157	201	327	-44	-169	154	154	353	0	-199

Notes: see Table 6.

seventh period, the situation reverses and the actual consumption increases, as it should.

3.4. *The Planning Horizons of the Subjects*

It is very clear from the results that we have already discussed that not all subjects solve the dynamic optimisation problem optimally. To do so requires a process of backward induction extending to the 25-period horizon of the problem. We could say that a fully-optimising subject has in this experiment a 25-period horizon – in the sense that he or she has the ability to plan ahead 25 periods in the first period, 24 periods ahead in the second and so on. This is a computationally complex problem so it is not surprising that some subjects appear not to have 25 period horizons. It could be argued that they adopt some procedure to simplify the problem.²² One such procedure is that elaborated on by Ballinger *et al.* (2003), in which subjects do not look as far as the correct horizon but instead act as if there were a shorter horizon, which they ‘roll forward’ as time passes. For example a subject with a planning horizon of two periods will always act as if the next period is to be the last (except of course in the 25th period which they know *is* the last); a

²² See also the remarks by Deaton (1992, pp. 156-7) in his analysis of non-experimental data.

Table 9
Period-by-period Effect of the Probability of Becoming Employed on Consumption

Period	Low					High				
	(1)	(2)	(3)	1-2	1-3	(1)	(2)	(3)	1-2	1-3
1	9	0	0	9	9	8	0	0	8	8
2	13	0	0	13	13	10	0	0	10	10
3	15	0	0	15	15	13	0	0	13	13
4	13	0	3	13	10	20	1	3	19	17
5	16	1	7	15	8	18	1	9	17	9
6	17	2	17	15	-1	16	3	19	13	-3
7	18	4	29	14	-11	18	7	32	11	-14
8	20	8	43	12	-23	26	13	47	12	-21
9	41	13	57	27	-17	36	20	62	16	-26
10	26	15	75	11	-49	43	25	79	18	-36
11	47	23	92	24	-45	43	30	96	13	-53
12	50	25	110	24	-60	41	37	113	4	-72
13	39	28	127	11	-88	50	47	131	3	-81
14	47	38	145	10	-98	67	59	149	9	-82
15	54	48	163	6	-109	75	68	167	7	-92
16	61	58	181	3	-119	82	79	184	3	-102
17	72	70	199	2	-127	104	92	201	12	-97
18	82	83	216	-0	-134	103	101	218	2	-116
19	93	90	233	3	-139	109	114	236	-5	-127
20	94	100	250	-6	-155	124	128	254	-4	-129
21	117	115	267	2	-151	135	144	271	-8	-136
22	125	128	286	-3	-161	148	160	289	-12	-140
23	122	134	303	-12	-181	138	168	306	-30	-167
24	141	154	322	-14	-182	122	182	321	-60	-198
25	126	150	339	-24	-213	185	206	341	-20	-156

Notes: see Table 6.

subject with a planning horizon of three periods will always act as if the next-but-one-period is to be the last (except of course in the 24th and 25th period when they know that they are the next-to-last and the last respectively). A completely myopic subject has a one-period horizon in that he or she always acts as if the present period is to be the last (and therefore consumes all their income every period). More generally a subject with an h period horizon acts, in period t , as if period $t + h - 1$ is to be last (except, of course, when $t + h - 1$ is greater than 25, in which case they correctly perceive how far away is the horizon).

We follow the procedure suggested by Ballinger *et al.* in trying to estimate the apparent planning horizons of the subjects. We do this as follows. For any planning horizon h we can work out the optimal consumption of the subjects using the optimal strategy (that we have already calculated) for the fully optimal subject. In this, optimal consumption in period t is a function of t , of wealth at the beginning of t and of the employment status in that period. The fully optimising subject – one with a 25 period horizon – always uses the function relevant for t in period t . In contrast a subject with planning horizon h uses the function relevant for period $26 - h$ in period t (instead of the one relevant for period t) whenever t is less than or equal to $26 - h$, and then uses the correct one (that is the one relevant for period t) when t is greater than or equal to $26 - h$. As before there are two

definitions of the 'optimal' consumption in any period: (1) the consumption that would be optimal given the wealth that the subject actually has in that period; (2) the consumption that would be optimal in that period given the income stream that the subject actually had and given that he or she had optimised in the past. We repeat that the second definition is more strict in that errors are compounded – there could be departures from this definition of optimality both through current period non-optimising and through having the wrong wealth at the start of the period through non-optimising in the past. Ballinger *et al.* prefer the first definition. It has certain advantages – particularly in that it lets us see more clearly whether subjects' behaviour is improving through time.

We have estimated the apparent planning horizon of the subjects using both definitions of optimal consumption. Moreover, we follow Ballinger *et al.* in estimating the apparent horizon as that which minimises the mean *squared* difference between actual consumption and optimal consumption.²³ The results are presented in Table 10, which gives information on the 'best' apparent horizon for each subject for each error definition. It will be seen that the error definition has only a minor effect on the results.

It is clear from this Table that there are substantive differences between subjects. There are some subjects with extremely short planning horizons. For example, subject 6 on parameter set 9 appears completely myopic – having an apparent planning horizon of just one period on both error definitions. Effectively this subject is simply consuming his or her income every period. In contrast there are subjects with very long planning horizons – many having apparent horizons of 20 periods.²⁴ If we look at these latter subjects we see that they are indeed the ones who saved up reasonable amounts of wealth during the experiment – realising that they could benefit from the high rates of interest. In contrast, those subjects with relatively small apparent horizons built up relatively small stocks of wealth during the experiment – not realising the returns that were possible from the high rates of interest.

It would be of interest to see if there is any connection between the apparent horizons of the subjects and the parameters of the model but, as can be seen from Table 10, such an analysis is somewhat confused by the high variability between subjects with the same parameter set. For example, with parameter set 1, the apparent horizon varies from 2 to 20. However, we can get rid of some of this variability by averaging over subjects with the same value of some particular parameter. For example, half the subjects had a high rate of interest (40%), half a low rate of interest (20%). If we use the mean squared difference measure based on the second error definition, then for the first half of the subjects (those with a high rate of interest) the average apparent horizon is 5.65 periods and for the second half of the subjects (those with a low rate of interest) an average apparent

²³ We have also done the analysis using the mean *absolute* difference – but the results are broadly similar.

²⁴ We should note that the optimal strategy is effectively the same for at least the first 10 periods of the experiment – so a subject loses very little by having a horizon of 15 instead of 25. Indeed for the subject with a reported apparent planning period of 20 in these Tables, the mean differences between actual and optimal is the same for apparent planning horizons of 20 to 25.

Table 10
Horizon Determined by Minimisation of Mean Squared Difference Between Actual and 'Optimal' Consumption

First error definition – table shows apparent horizons of subjects						
Subject parameter set	1	2	3	4	5	6
1	2	5	2	20	2	2
2	9	3	19	3	5	3
3	3	2	2	5	3	11
4	2	3	2	6	2	21
5	11	5	2	20	2	4
6	5	20	4	2	15	20
7	5	4	13	2	14	17
8	4	3	11	7	13	3
9	2	14	18	2	3	1
10	14	2	17	2	17	2
11	7	2	6	5	6	3
12	1	3	2	7	8	5
13	4	2	17	5	2	2
14	4	13	2	2	5	3
15	3	2	2	3	1	2
16	2	2	12	8	1	2

Second error definition – table shows apparent horizons of subjects						
Subject parameter set	1	2	3	4	5	6
1	2	5	2	20	2	2
2	8	3	21	3	4	3
3	3	2	2	5	3	6
4	2	3	2	6	2	21
5	4	2	2	19	2	4
6	4	19	4	2	13	20
7	6	4	15	2	19	20
8	4	2	9	7	7	3
9	2	19	19	2	3	1
10	11	2	8	2	17	2
11	5	3	5	5	5	3
12	1	3	2	7	8	5
13	4	2	19	5	2	2
14	3	9	2	2	5	3
15	3	2	2	3	1	1
16	2	2	16	6	2	2

horizon of 6.13 periods. There seems to be nothing of economic significance here and certainly they are not statistically significantly different. Similarly for the subjects with a high probability of *remaining* employed, their average apparent horizon is 5.67 and the average for those with a low probability of remaining employed 6.10. For the subjects with a high probability of *becoming* employed their average apparent horizon is 6.58 and those with a low probability of becoming employed it is 5.19. Although these are not statistically significantly different, the difference has a modest economic significance – suggesting that those who are more likely to leave unemployment think more carefully about the future. Finally for those with high incomes, when employed, the average apparent horizon is 6.77, while for those with a low income when employed, the average apparent horizon is

5.00. This is interesting – the greater payoff to being employed seems to induce subjects to think more carefully about their future. But none of these differences are statistically significant. The significant differences seem to be between subjects and not between parameter sets. Put simply – some subjects are better than others in that they have longer planning horizons.

4. Conclusions

One of the important conclusions of this study is contained in the sentence above: subjects differ in their ability to solve the task. In particular, subjects differ in their ability to think ahead – some subjects seem to be able to think a long way ahead, others only a little way. We could classify subjects according to their ‘apparent planning horizon’ – as discussed in sub-section 3.4. Alternatively we could classify subjects according to the way they tackled the problem: on this criterion there seem to be four basic types:

- (i) those who understand the basic nature of the problem – including the returns from saving and the diminishing returns from consuming – and who approach the optimal strategy in varying degrees,
- (ii) those who are pre-occupied with the present and who seem to think little about the future;
- (iii) those who simply seem to like to have wealth, who build up excessive amounts of wealth during the experiment; and – those who seem rather confused
- (iv) building up stocks of wealth over cycles of around 4 or 5 periods and then consuming almost all of these built-up stocks of wealth.

If we exclude from the discussion those subjects who seemed to get pleasure from building up enormous stocks of wealth, we could conclude that virtually all the others were trying, with varying degrees of success, to solve the optimisation problem – though most with a too-short planning horizon or with a variable planning horizon. As a consequence of this apparent myopia, it follows that the behaviour of the majority of the subjects was such that they consumed too much in the early stages of the experiment (when their wealth was low) and as a consequence had too low levels of wealth in the later stages of the experiment and thus consumed too little in these later stages. This behaviour could result *either* from a too-short planning horizon *or* from an underestimation of the effect of interest on savings. However, it is clear that subjects take into account the role of the rate of interest – indeed increasing their consumption excessively when the rate of interest rises.

We also observe significant over-reaction to the current employment status: subjects tend to consume too much in periods of employment and too little in periods of unemployment. This is perhaps a manifestation of a more general phenomenon: subjects do not seem to be able to *smooth* their consumption stream sufficiently – with current consumption too closely tracking current income. This was also observed by Ballinger *et al* (2003) and is frequently observed in analyses based on questionnaire data. In this context it implies that subjects are worse off in

periods of unemployment than they need to be and better off in periods of employment than they should be. It prompts the question: how should governments take into account such myopia when planning the state unemployment insurance scheme?

It seems clear that subjects have difficulty in taking into account the probabilistic structure of the income process correctly. Although a particularly simple process (a first-order Markov process) from the point of view of a statistician, it is difficult for a non-statistician to assimilate. Perhaps this is the reason for the inability to smooth the consumption stream sufficiently and, perhaps, also the reason why the key parameters p and q have effects different from those predicted by the theory. In particular, subjects over-respond to an increase in p , while *generally*²⁵ under-responding to an increase in q . Perhaps they regard the high value of p (0.875) as effectively a case of permanent employment (while the low value of p is clearly a risky case), and the low value of q (0.125) as effectively a case of permanent unemployment (while the high value of q is clearly a risky case)?

Whilst subjects seem to have difficulty with the stochastic nature of the income process, they seem to have much less difficulty in understanding the actual values of income – as evidenced by the fact that they respond almost exactly correctly to an increase in the level of employed income.

So individuals have trouble optimising and have trouble understanding the stochastic structure of the problem. As a consequence they smooth their consumption stream insufficiently and over-respond to the current situation. In the light of this under-optimisation, it is an interesting open question as to how governments should respond.

University of York and Università di Bari

Date of receipt of first submission: May 2001

Date of receipt of final typescript: October 2003

Dataset is available for this paper: www.res.org.uk

References

- Ballinger T.P., Palumbo M.G. and Wilcox N.T. (2003). 'Precautionary saving and social learning across generations: an experiment', *ECONOMIC JOURNAL*, vol. 113(490), pp. 920–47.
- Deaton A.S. (1992). *Understanding Consumption*, Oxford: Clarendon Press.
- Hey J.D. (1980). 'Optimal consumption under income uncertainty: an example and a conjecture', *Economics Letters*, vol. 5(2), pp. 129–33.
- Hey J.D. and Dardanoni V. (1988). 'Optimal consumption under uncertainty: an experimental investigation'. *ECONOMIC JOURNAL*, vol. 98(2), pp. 105–16.
- Kagel J.H. and Roth A.E. (1995). *Handbook of Experimental Economics*, Princeton University Press.
- Kim H. (1989). 'An experimental study of consumption: test of the permanent income and life cycle hypotheses', Indiana University unpublished dissertation.

²⁵ Though there is an interesting twist to this: in the early stages of the experiment subjects actually respond in the *wrong direction* to an increase in q . Possible reasons for this are discussed in sub-section 3.3.