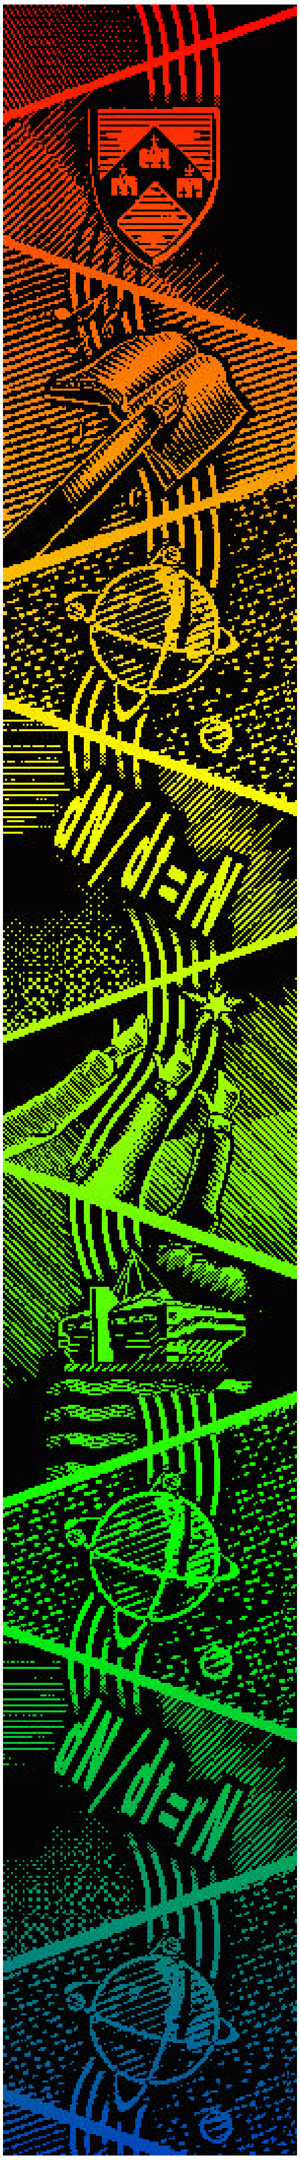

THE UNIVERSITY *of York*

Discussion Papers in Economics

No. 1999/09

A Test of the Principle of Optimality

by

Enrica Carbone and John Hey

Department of Economics and Related Studies
University of York
Heslington
York, YO10 5DD

A Test of the Principle of Optimality*

Enrica Carbone
Universities of Sannio and York

John D. Hey
Universities of York and Bari

August 17, 1998

Abstract

This paper reports on an experimental test of the Principle of Optimality in dynamic decision problems. This Principle, which states that the decision-maker should always choose the optimal decision at each stage of the decision problem, conditional on behaving optimally thereafter, underlies many theories of optimal dynamic decision making, but is normally difficult to test empirically without knowledge of the decision-maker's preference function. In the experiment reported here we use a new experimental procedure to get round this difficulty, which also enables us to shed some light on the decision process that the decision-maker is using if he or she is not using the Principle of Optimality - which appears to be the case in our experiments.

Keywords: dynamic decision making, backward induction, principle of optimality, experiments, software.

JEL classiffications: C91, D81, D90.

1 Introduction

In the economics literature on dynamic decision making there are essentially two main stories about how decision makers tackle dynamic decision problems: (1) that they use backward induction combined with the Principle of Optimality; (2) that they convert the original dynamic decision problem into a static strategy choice problem. Irrespective of which procedure is used the decision-maker needs only to have preferences over static decision problems: with (1) the decision maker works backwards solving each static decision problem in turn (beginning with the final one) and then using these solutions to reduce earlier dynamic problems into static ones - through the elimination of future choices which

*Our thanks to participants in the TMR conference held in Naples in February 1997, in the Italian xperimental conomics conference in Trento in June 1997 and in the conomic Science Association cofererence in Mannheim in June 1998 for helpful comments, and to seminar participants at the Universities of Bonn and Malmo. Our thanks also to the uropean Commission for a grant under the TMR programme "Savings and Pensions" which financed this research.

are sub-optimal; with (2) the choice of a strategy is simply a static decision problem. If the individual's preference function over static decision problems is known then one should be able to infer which of these two procedures is being used from observations on actual decisions - unless, of course, these static preferences are Expected Utility (EU) preferences in which case the two procedures yield exactly the same decision rule in all situations[1]. If, however, these static preferences are not known, then it becomes difficult - from observations on decisions alone - to infer which, if either, of these two procedures an individual is using.

In general if the static preferences are *not* EU preferences then the two procedures will yield different decisions. Procedure (1) is essentially (Belman's) Principle of Optimality; thus if the decision maker uses procedure (2) and does not have EU preferences, then it rather trivially follows that the decision maker is not using the Principle of Optimality - which is not surprising since the decision maker is essentially using a static decision rule. Whether this is rational is an interesting point. However it is not our major concern here. What is our major concern is to try and discover whether or not people use the Principle of Optimality, and, if not, whether they use procedure (2) or some other procedure.

As we have already noted, if we do not know the individual's static preferences we will be unable to infer from observations on decisions alone whether the individual is using procedure (1) or procedure (2) or indeed some other procedure. To do so we need some additional information. Here we report on an experiment with a design which, in principle, gives us that extra information - by directly revealing the decision process that the individual is using. The next section describes that experimental design.

2 Our Experimental Methodologyy

We are interested in studying decision processes in dynamic decision problems under risk[2]. In order to make a decision problem under risk a dynamic one we need at least two decision stages each followed by a move by nature. At each decision stage at least two alternatives

[1]Indeed some economists would argue that this is an important property of U preferences which lend it its normative appeal.

[2]In a *certain* world a dynamic decision problem can be trivially reduced to a static one.

should be available; the simplest case is when there are just two. At each move by nature there should be at least two possibilities[3]; the simplest case is when there are just two, each equally likely. Even under these simplifications, notice how rapidly the complexity of the problem increases with the number of decision stages: if there are just two decision stages each followed by a move by nature, the number of final possibilities is 16 ($= 2^4$); with three decision stages each followed by a move by nature there are 64 ($= 2^8$) final possibilities; with four decision stages 256 ($= 2^{16}$) final possibilities; with five decision stages 1024 ($= 2^{32}$); and so on. To implement an experiment we need to decide the number of stages: in the experiment reported here we chose *three* - implying 64 final possibilities - or *end-nodes*.

decision maker can explore (the residual part of[5]) the tree. In particular the decision maker can explore the payoffs associated with some or all of the end nodes. These payoffs are not shown to the decision maker initially - he or she must discover them by the following procedure. One of the nodes is highlighted white; the highlighted node can be moved around the tree using the cursor keys on the computer. When any node is highlighted, the decision maker can press the 'Enter' key on the computer - thereby opening what we call the *Notebook* associated with that node. Each of these Notebooks has a space in which the decision maker can type comments and leave messages, which stay in that particular node's Notebook until the node is re-visited, the Notebook re-opened and the message amended. In addition, the Notebooks associated with each end node tells the decision maker the payoff associated with that end node. Therefore, to discover the payoff associated with a particular end node, the decision maker must move the highlighted node to that end node, press 'Enter', thereby opening the Notebook associated with that end node, and then read the value of the payoff from that Notebook. Intermediate choice and chance nodes did not have payoffs in the experiment that we ran[6].

The Notebooks were a crucial feature of our experimental design, in conjunction with the fact that we did not allow the subjects in our experiment to use pencil or paper or keep notes of any kind, other than in the Notebooks. Given that the decision problem was quite a complicated one, these features meant that if the decision maker wanted to systematise his or her decision making then he or she would have to keep notes in the Notebooks. Through these notes we could discover the process through which the decision maker was tackling the problem. Indeed, the software recorded every key stroke of the subject, each movement of the cursor, each Notebook opened and every message left in the Notebooks in the appropriate chronological sequence. Moreover, our software enabled us to 'replay' as often as we wanted the actual sequence that each subject had gone through - thereby recreating the actual experimental experience. Through this we get considerable insight into what the subjects were doing, what they were thinking and the motivation behind their decision procedure. In particular, we can check to see whether their procedure

[5]After each move by the decision maker or by Nature, one half of the remaining part of the tree is eliminated by that move.

[6]Though we intend to incorporate such intermediate payoffs in subsequent experiments.

is consistent or not with one or other of the two procedures outlined above and discussed further below.

3 Economic Theory

Let us begin with some notation. Let us call the three decision stages D1, D2 and D3, and the three chance stages C1, C2, C3. There is just one D1 choice node, 2 C1 chance nodes, 4 D2 choice nodes, 8 C2 chance nodes, 16 D3 choice nodes, 32 C3 chance nodes and 64 end nodes.

Using this notation the hypotheses under test can be simply stated. If a subject is following procedure (1) (Backward Induction combined with the Principle of Optimality) then he or she should work as follows: (a) Associated with each of the 16 decision nodes involved at D3 are 4 end nodes - the left hand two being the two equally likely possible end nodes if the decision maker chooses Left at that D3 node, and the right hand two being the two equally lik decision in the Notebook associated with the relevant D2 node - *along with the value of the preference function implied by following the better choice.* (d) He or she should do this for all 4 D2 choice nodes. (e) Then he or she should work back once more - now looking, for the only D1 choice node, at the 4 D2 choice nodes associated with the D1 node - the left hand two being the two equally likely possible D2 nodes if the decision maker chooses Left at the D1 choice node and the right hand two being the two equally likely possible D2 nodes if the decision maker chooses Right at the D1 choice node. Using the fact that the decision maker has already entered into these 4 D2 choice nodes the value of being there (in terms of the decision maker's own monetary evaluation of the benefit of being there) the decision maker should work out, for the D1 node whether Left or Right is better (gives a higher value to the preference function), and then put that decision in the Notebook associated with the D1 node - along with the value of the preference function implied by following the better choice (if he or she wants, though this is not useful information). The decision problem is now solved: in the one D1 choice node Notebook there is a decision, in each of the 4 D2 choice node's Notebook there are decisions and in each of the 16 D2 choice node's Notebook there are decisions. All that remains to do is to implement them, conditional on what ever Nature decides to do at the intervening chance nodes.

If, in contrast, the decision maker is using procedure (2), then a completely different method is required. First the decision maker must list all the possible strategies; then work out the implied probability distribution over final payoffs; and then determine the best of these probability distributions using whatever static preference functional the individual has. The set of possible strategies can be determined as follows: at the D1 stage, there is a single choice, Left or Right; for *each* of these there are two possible responses by Nature, and for each of these the decision maker must decide whether to play Left or Right at the implied D2 node; this gives us $2 * 2 * 2 = 8$ possible strategies up to D2; then, for *each* of these there are two possible responses by Nature, and for each of these the decision maker must decide whether to play Left or Right at the implied D3 node; this gives us a grand total of $8 * 2 * 2 = 32$ possible strategies.

An example is the following, where by L (R) we mean that *Nature* chooses to play L

(R) at C1; and by *LL* (respectively *RL, LR, RR*) we mean that *Nature* chooses to play L at C1 and L at C2 (respectively R at C1 and L at C2, L at C1 and R at C2, R at C1 and R at C2):

D1: L

D2: R if *L*, R if *R*

D3: L if *LL*, R if *RL*, L if *LR*, R if *RR*

In general a strategy is of the form

D1: X_1

D2: X_2 if L, X_3 if R

D3: X_4 if LL, X_5 if RL, X_6 if LR, X_7 if RR

where each of the X_i are either L or R.

For any one of these 32 possible strategies the ultimate end-node will be one of 8 of the 64 possible end-nodes. Given that Nature, in our experiment, is equally likely to play Left or Right at any chance node all 8 of these are equally likely. Consider a particular strategy that the decision maker might employ - for example, consider the strategy of always playing Left. The end nodes that this strategy *could* lead to are (numbering end nodes end nodes from the left) 1 2 5 6 17 18 21 22. (Nodes 33 to 64 are eliminated by the choice of L at D1; then 9 to 16 and 25 to 32 are eliminated by the decision to play L at D2; finally 3 and 4, 7 and 8, 19 and 20, 23 and 24 are eliminated by the decision to play L at D3). Each of these 8 nodes is equally likely. The decision maker can discover the payoffs associated with these 8 end nodes (in the manner previously described) and so the decision maker can evaluate, with his or her static preference functional) the value of following the strategy of playing Left always. This can be repeated for all 32 strategies.

We think it very unlikely that any subject tried to implement this procedure in the context of our experiment. Evidence for this remark is the following. Given that the subjects were not allowed to use paper and pencil the only way that they could keep notes was in the Notebooks provided by our software. Subjects would have to list somewhere all 32 possible strategies and for each of these work out the associated set of 8 end nodes and then evaluate these 8 end nodes (each equally likely) using whatever preference functional they possessed. So we should see somewhere subjects keeping notes of the following form:

L,LL,LLLL.............1 2 5 6 17 18 21 22...............£0 £16 £50 £12 £32 £20 £12 £0

....

R,RR,RRRR..........39 40 47 48 55 56 63 64........£40 £10 £0 £4 £16 £32 £18 £20

where the left hand block indicates the strategy, the middle block the implied end nodes, and the right hand block the associated payoffs (each equally likely). The decision maker would then need to evaluate the relative merits of (£0 £16 £50 £12 £32 £20 £12 £0) against .... against (£40 £10 £0 £4 £16 £32 £18 £20).

We saw no evidence at all of any attempt to implement this procedure. In a sense this is not surprising given the way that the experiment was constructed and the software defined. But more generally we do not find this surprising as the implementation of this procedure is a complex one: even for people skilled in decision making, keeping tracks of the relevant end nodes is difficult and time-consuming (though by no means impossible). Let us now turn to what we did discover, after briefly describing the experimental implementation of the experiment.

4 Experimental Details

Our experimental implementation involved 20 subjects. They were paid the payoff in the end-node that they eventually reached. In the experiment they earned a total of £373 - an average of £18.65 per subject. The time to complete the experiment varied from subject to subject: including a 20 minute briefing, individual subjects took between 45 minutes and 90 minutes. In the invitations, subjects had been told that the experiment might last up to 2 hours. The briefing was extensive, beginning with a verbal (tape-recorded) repetition of the Instructions (which are reproduced here in the Appendix to this paper), and followed by an example partial session with a video display of the computer screen and a hypothetical subject session. Subjects were made aware that they could explore the tree, and leave notes in the Notebooks, as much or as little as they wished at any stage of the decision process. They were also made aware that the payoffs might vary widely from end-node to end-node and that decisions, once taken, were irrevocable. Subjects had ample opportunity to ask questions.

The 20 subjects were divided equally amongst 4 data sets - distinguished by the payoffs in the 64 end nodes. The actual 64 payoffs in the four data sets were the same - though they were distributed differently amongst the 64 end nodes. The details are as follows - where the 64 payoffs are arranged left to right as on the computer screen. The payoffs are all in pounds sterling.

Data set 1: 0 49 20 20 15 15 34 0 10 22 15 15 0 34 22 10 40 0 18 18 18 18 9 29 29 9 0 40 16 26 28 18 40 0 19 19 29 19 17 27 20 20 0 50 16 16 23 11 10 29 19 19 40 0 10 29 16 16 35 0 23 11 0 35

Data set 2: 49 0 20 20 15 15 0 34 35 0 16 16 0 50 20 20 15 15 22 10 0 34 10 22 23 11 35 0 16 16 11 23 18 18 0 40 29 9 18 18 19 19 10 29 40 0 19 19 0 40 29 9 28 18 16 26 17 27 19 29 40 0 29 10

Data set 3: 49 0 20 20 0 50 20 20 15 15 34 0 16 16 0 35 22 10 15 15 16 16 11 23 34 0 10 22 11 23 35 0 19 19 0 40 18 18 40 0 10 29 19 19 29 9 18 18 10 29 40 0 29 9 0 40 29 19 17 27 16 26 18 28

Data set 4: 0 49 50 0 0 34 0 35 10 22 23 11 0 34 35 0 0 40 40 0 9 29 10 29 40 0 0 40 16 26 27 17 20 20 20 20 15 15 16 16 15 15 16 16 10 22 23 11 18 18 19 19 18 18 19 19 9 29 29 10 18 28 29 19

To a certain extent the actual payoffs should not influence the way that 'rational' subjects tackle the decision problem[8]. However, if subjects are not 'rational' (in the sense prescribed by the procedures outlined above), and instead adopt some kind of heuristic, then the payoffs might influence the *procedure* - in which case, varying the payoff structure might give valuable insight. We intend to explore the implications of this in future experiments.

[8]Unless, of course, for example, the payoffs are arranged in such a way that the problem becomes trivial and the subject notices this: for example if end-nodes 1 to 32 have a payoff of £0 while end-nodes 33 to 64 have a payoff of £50.

5 Experimental Results

We can make available on request the files containing the subjects' data and the program for 'replaying' the experimental sessions. These make very clear what the subjects were doing (though not always why they were doing it). One feature that is immediately apparent is that no two subjects were doing (exactly) the same thing. Moreover, it is clear that no subject was using the Principle of Optimality as described above: no subjects left decisions *and* values in all decision nodes - and no subjects seemed to be taking into account that they would be optimising - or at least taking decisions - in the future. So the immediate answer to the question in the title of this paper is "no" - none of the 20 seemed to be using the Principle of Optimality. Nor were they using procedure (2) - as we have already remarked.

What then were they doing? In answering this question we adopt the economist's usual strategy - of trying to state broad generalisations (of the observed behaviour) that have empirically distinguishable (and hence testable) implications. We do not think it of value to record every nuance of behaviour, nor to simply conclude that 'different people do different things'.

What broad generalisations emerge?

First, we notice *a very common tendency to avoid bad outcomes*. This could be characterised as very risk averse behaviour or even maximin behaviour (of two risky options choosing the one where the worst outcome is least bad). As it happens, in the context of our experiment, where all the risky prospects involved just two possibilities, and in which we had chosen the prospects in such a way that, whenever the subject was confronting two risky options, it would always be the case that the risky option had the higher mean, these two forms of behaviour are impossible to distinguish empirically. So we cannot be sure whether our subjects were risk averse or maximiners. But, of course, this in itself is irrelevant to the hypothesis under test.

Crucially relevant is our second broad generalisation: *many subjects were using backward induction* - as the Principle of Optimality would have them do - *but were ignoring the fact that they would be behaving optimally thereafter*. So, for example, in deciding whether to go Left or Right at D1, such subjects were looking at *all* 32 end nodes associated with choosing Left at D1 and comparing these 32 end nodes with *all* 32 end nodes associated with choosing Right at D2 - *ignoring the fact that three-quarters of such end nodes would be excluded by future decisions*. A prime example of such a subject is Subject 1: this was a subject who appeared to be basically risk-neutral (but who seemed prepared to trade-off the mean return against the numbers of zeros involved with any prospect) and who was therefore calculating the mean (or total) payoff associated with any decision - but ignoring the fact that he or she would be taking decisions later on. So, for example, in deciding whether to go Left or Right at D1, this subject calculated the mean payoff across end nodes 1 to 32 and compared this with the mean payoff across end nodes 33 to 64. Having taken the decision at D1 on the basis of this criterion, the subject repeated the procedure at D2. Here there were just 16 end nodes remaining (half of the original 64 having been eliminated by the decision at D1 and half the remaining 32 having been eliminated by Nature's move at C1) - half of these associated with Left at D2 and half with R at D2. The subject again calculated the mean payoff attached to the Left 8 and compared this with the mean payoff attached to the Right 8 - and again took the decision (now at D2) on the basis of which of these two were highest. Further, at D3, the procedure was repeated - though, of course, at this decision node, with no further decisions to come, this procedure was justifiable.

We feel that this broad generalisation (possibly characterising all or part of the behaviour of over half our subjects) is the most important hypothesis emerging from our experiment. If empirically verified in future work, it has significant implications. Moreover, we suspect, though we have not yet explored this fully yet, that the predictions of important economic applications (such as savings and investment behaviour) using this 'principle' rather than the Principle of Optimality, are likely to be significantly different - leading possibly to more plausible explanations of economic phenomena than those provided hitherto by conventional theory.

Consider, for example, the implications for a risk-neutral subject who uses this 'principle'. By ignoring the fact that he or she will choose ('somehow' - thought not optimally[9]) in

[9] Remember that the decision maker is not optimising since he or she will commit the same mistake throughout the decision process.

the future, and by calculating the mean payoff over *all* subsequent end nodes, the decision maker is implicitly acting on the presumption that he or she is *equally likely* to take either of any pair of choices in the future. This is an interesting story - and one that seems to tie in neatly with a third observation: that subjects often, even though they had earlier entered decisions into their Notebooks at the choice nodes at D2 and D3, *checked them* using their decision procedure, when they arrived at that choice node. This suggests that these earlier decisions were provisional in some sense - either because they did not really know what their preferences would be once they arrived at that node, or perhaps because they did not trust their earlier calculations[10].

Consider also the implications for a maximiner. Such a person, by ignoring the fact that he or she will choose in the future, while choosing between prospects on the basis of their worst outcomes is implicitly acting on the presumption that he or she *will take the worst decision*[11] *in the future*! This emphasises the fact that a maximiner can be viewed as a deeply pessimistic person - but now not only about what Nature is likely to do but also about what she herself (or he himself) is going to do.

Some justification of this procedure might, as we have already argued, come from considerations of computational cost. Considerations of bounded rationality might also lead to a similar conclusion. Other evidence that such considerations might be important comes from a third broad generalisation: *that some subjects try and simplify the problem in some way*. One obvious way is to take the decision at D1 randomly - this reduces the problem from that of considering 64 end nodes to that of considering just 16, a significant simplification, and one partially 'justified' by subjects scanning across all 64 end nodes trying to get a general impression as to whether there is any significant difference between nodes 1 to 32 and nodes 33 to 64[12]. Several subjects used these procedures, either at D1 when it is particularly cost-saving, or later.

Other generalisations could be made[13] but the three above seem the important ones to

[10]This latter is an interesting thought and one that makes sense in a world where computation and thought is costly: in such a world it may well make sense to approximate in the first instance and then go back and calculate more carefully if necessary.

[11]From the point of view of his or her preference function

[12]Of course, if subjects are using the Principle of Optimality, this is uninformative.

[13]For example, that subjects start out the experiment with a clear strategy but that they change it or

us. We intend to work on them further in the future - both to test empirically their validity and to explore their theoretical implications. In the meantime, we invite others to examine our data - and do not rule out the possibility of other broad generalisations emerging.

6 Conclusions

The most immediate conclusion is that the subjects in our experiment neither used the Principle of Optimality (our procedure (1)) nor the alternative procedure (2). This could have been for any number of reasons, particularly the experimental design. Some of these we have already discussed. Of course, the subjects in the particular experimental implementation reported here were not given the chance to learn through experience, nor did we try and educate them by telling them about the Principle of Optimality. Such avenues remain to be explored in future experiments, though whether they are relevant depends upon the intended application of the theory.

Although negative in terms of the hypotheses under test, our experiment was positive in terms of alternative hypotheses: whilst our observations that subjects are very risk averse (possibly even maximiners) and that they simplify, approximate and check are interesting, we feel that the observation that subjects ignore the fact they *they* will subsequently take decisions (and hence eliminate future possible end nodes) is the most profound, and the one with the most far-reaching implications. We look forward to their exploration.

7 Bibliographical Background

A brief note on the bibliographical background to this experiment might be useful, though we omit all of the associated experiments involved with testing the *implications* of various principles of optimality[14]. Our prime concern here is that of discovering the decision *processes* underlying the actual decisions. In the economics experimental literature, precedents are less numerous, though a clear intellectual antecedent of our experiment is the work of

get tired half way through - which interestingly suggests that they had not correctly anticipated the costs associated with their initial strategy.

[14]Such experiments, in the context of dynamic decision making, are numerous: for example (Hey, 1993), (Hey and Dardanoni, 1988) and (Strobel, 1996).

Colin Camerer and various associates, (Camerer et al., 1993) and (Camerer et al., 1998), on behaviour in games, exploiting the Mouselab software (developed by psychologists) which facilitates the discovery of the thought processes of subjects in experiments. More recently there has been a number of papers with a similar motivation coming out of the experimental laboratory of the University of Trento, under the supervision of Massimo Egidi. Some other similarly motivated work has also recently been done at the University of Bonn, under the supervision of Reinhard Selten.

References

Camerer, C. F., Johnston, E., Rymon, T., and Sen, S. (1993). Cognition and framing in sequential bargaining for gains and losses. In Kirman, A. and Tani, P., editors, *Frontiers of Game Theory*, pages 27–48. MIT Press.

Camerer, C. F., Johnston, E., Rymon, T., and Sen, S. (1998). Detecting failures of backward induction: Evidence from sequential bargaining experiments. Unpublished working paper.

Hey, J. D. (1993). Dynamic decision making under uncertainty: An experimental study of the dynamic competitive firm. *Oxford conomic Papers*, pages 58–82.

Hey, J. D. and Dardanoni, V. (1988). A large scale experimental investigation into optimal consumption under uncertainty. *conomic Journal*, 98:105–116.

Strobel, M. (1996). Some fundamental aspects of saving, updating and dynamic programming - an experimental analysis of dynamic saving behaviour in a stochastic environment. Unpublished working paper of Humboldt University Berlin.

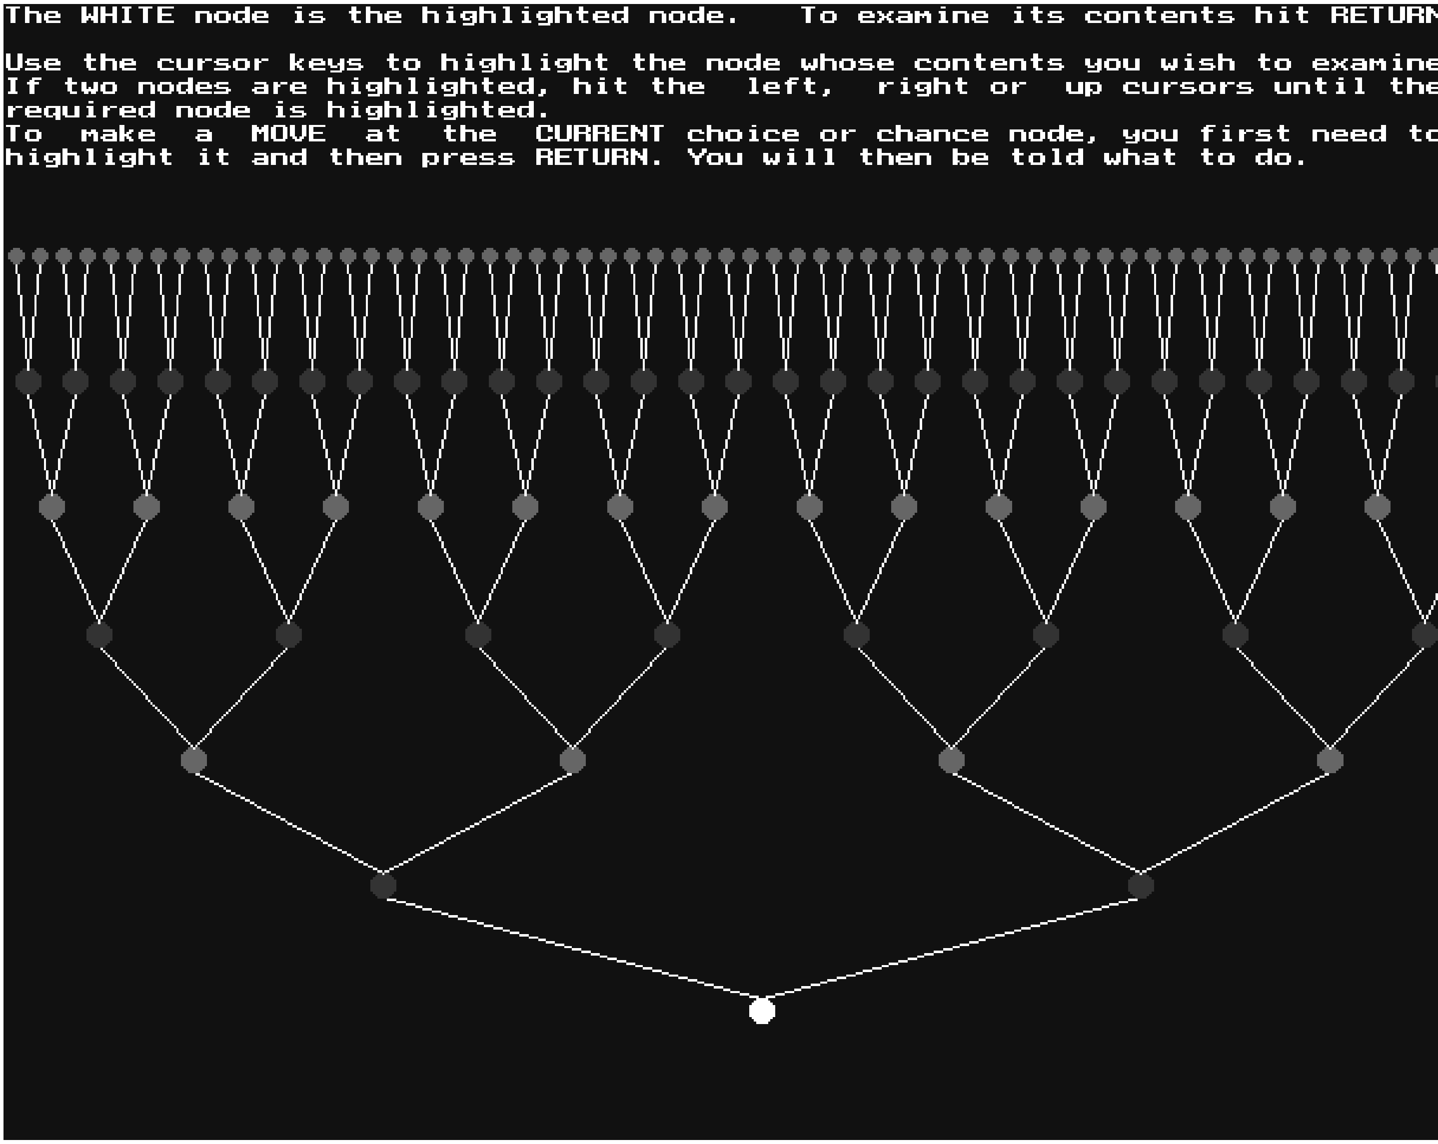

Figure 1: The Basic Screen for the Experiment

INSTRUCTIONS

Welcome to this experiment. It is an experiment on the economics of dynamic decision making under risk. Various research funding bodies have provided the funds to finance this research. The instructions are straightforward, and if you follow them carefully you may earn a considerable amount of money which will be paid to you in cash immediately after the end of the experiment. Please read the instructions carefully and take as much time as you need. There are no right or wrong ways to complete the experiment, but what you do will have implications for what you are paid at the end of the experiment.

Basic Structure of the Experiment

The experiment proceeds in a sequence of *choice* and *chance* nodes. At each node there two subsequent paths to follow: Left and Right. At each *choice node* you will have to take a *decision* - in each case whether to go Left or Right. At each *chance node* a chance device (which you will operate) will determine whether Left or Right is chosen. In total there are *three* choice nodes and *three* chance nodes, starting with a *choice* node and then alternating the two types until the final *chance* node. So the entire sequence is: choice, chance, choice, chance, choice, chance. After the third and final chance node is played out you will arrive at an *end* ***node***. Each *end node* has associated with it a *payoff* - a certain amount of money which is yours to keep if you arrive at that end node. To discover the payoff associated with any particular *end node* you will have to *highlight* that node (I shall explain shortly how that is done) and 'open' it by hitting the 'Return' key on your keyboard. This will reveal the payoff associated with that node and also the *Notebook* associated with that node.

Notebooks

All nodes have associated *Notebooks* which can be opened in the same fashion: by highlighting that node and pressing the 'Return' key on your keyboard. Once open, you can leave notes in the *Notebook* for future reference. These Notebooks may prove useful to you in deciding whether to move Left or Right at each choice node. Indeed, since we will not allow you to use paper and pencil nor to take any other notes during the experiment, the use of these Notebooks is the *only* way you can take and keep notes to help you in your decision making. You can use these Notebooks as much or as little as you wish.

More Details

You should read these instructions carefully, and then turn to the computer. It begins by asking you some details to help us identify you; these details will be kept confidential. The program will then briefly overview these instructions before proceeding to the main screen. This contains the decision tree: starting at the bottom with a single highlighted (white) choice node, and then proceeding upwards through a level containing 2 chance nodes (red circles); then to a level containing 4 choice nodes (green circles); then to a level containing 8 chance nodes (red circles); then to a level containing 16 choice nodes (green circles); then to a level containing 32 chance nodes (red circles); and finally to the top level containing 64 end nodes (brown circles). See the figure attached to these instructions for a preview. Ultimately you will move up through the tree until you reach one of the end nodes; which end node you end up at will determine your payment. At each choice node you come to, you will eventually have to decide whether to go Left or Right; at each chance node you come to, a chance device (which you will operate) will decide whether you go Left or Right. So the final end node you end up at depends partly on your choices and partly on chance.

Highlighting Specific Nodes

At any stage in the whole process, either before taking any choices or playing out any chance nodes, or after taking one or more choices or after playing out one or more chance nodes, you can explore the tree and leave notes in the Notebooks associated with any node. This is simply done by highlighting the node and then hitting the 'Return' key. To highlight a node - turning it white - you use the cursor keys on the keyboard: the $\rightarrow, \leftarrow, \uparrow$ and $\downarrow$ keys. To move the highlighted node up a level you hit the $\uparrow$ key; to move it down a level you hit the $\downarrow$ key; to move it left you hit the $\leftarrow$ key; and to move it right you hit the $\rightarrow$ key. After moving up a level you will find that two nodes are highlighted; in order to highlight just one of these nodes you will have to use either the $\leftarrow$ key or the $\rightarrow$ key.

Examining the Notebooks

Once a node is highlighted - it turns white - you can examine the contents of its associated Notebook (and the associate payoff if it is an end node) by hitting the 'Return' key. You can then add to or amend the contents of the Notebook simply by typing in the usual fashion (including using delete keys etc). To then close the Notebook you press the 'Return' key.

Making Moves

To make a move at the current node you should press the 'Escape' key (at the top lefthand part of your keyboard). If the current node is a choice node you will be asked whether you wish to move Left or Right; you will then be asked to confirm your choice by hitting 'Y' for Yes. If you want to change your choice - or indeed not make a choice after all - you should press any other key. If the current node is a chance node, you will be asked to call the experimenter. He or she will then get you to pick one of two cards - one labelled Left and one labelled Right - at random. The one you pick will determine your move at that chance node. It is important that you do call the experimenter at this stage: if you do not, your experiment will be declared null and void and you will not get paid for taking part in the experiment. Please note that you can make moves only at the current node.

Final Payoff

The payoff for the tree is the payoff in the final end node. This will be printed out on the screen when you reach that final end node. At this stage you should call the Experimenter.

Other

If there is any aspect of these instructions about which you are not clear, please ask the Experimenter. It is clearly in your interests to understand these instructions as fully as possible. Please also feel free to call the Experimenter at any time.

SUMMARY OF THE SUBJECTS' BEHAVIOUR

Subject 1 Averages over all subsequent end-nodes; trades off number of zeros against average.

Subject 2 Overview for D1; checks for dominance across all remaining end-nodes at D2; chooses safer option at D3

Subject 3 Scans for zeros in choosing R at D1 and D2, I think; goes for safer at D3.

Subject 4 Starts out very systematically and then chooses R at D1 with no obvious reason; scans across 16 remaining to choose at D2; same at D3.

Subject 5 Partially systematic - but only very partial; guesses at D1?; more systematic afterwards - but no values brought back.

Subject 6 Starts out very systematically, then takes D1 apparently at random; effectively ignores previous painstaking Notebooks - re-scans, finds dominance; goes for risky at D3.

Subject 7 Uses Backward Induction on minima.

Subject 8 Uses 'scan technique'; keeps nothing in Notebooks; but respects dominance at D3.

Subject 9 Uses 'scan technique' for D1; then compares 8 end-nodes of Left with 8 end-nodes of Right for D2; then chooses riskier at D3 (same as at D2).

Subject 10 Completely ignores Right hand half; chooses safer option at D3; not clear how D2 taken.

Subject 11 Very systematically works back with sets of 4 numbers; takes decisions at D2 and D3 then scans; just decisions, but no values, in D2 and D3; how working back?

Subject 12 Backward Induction of totals - having taken a scan decision at D1; effectively the same at D2 and D3; ignores subsequent optimising.

Subject 13 Guy in a hurry; chooses at D1 immediately; scans D2; correctly chooses dominating choice at D3.

Subject 14 Groups 4 end-node payoffs in the D3 Notebooks; scans for D1; then groups 4 sets of 4 to take D2; then safer at D3.

Subject 15 Works out decisions at D3 (safer); tries backward induction on minima - incompletely.

Subject 16 Scans for D1; then very systematically brings back all numbers - but no use of subsequent optimising.

Subject 17 Works on minima; brings back to D2 - then scans, I think.

Subject 18 Totally confused - decisions at random?

Subject 19 Dislikes zero; finding safe route seem to be strategy - does not take account of future optimising.

Subject 20 Avoids zeros; looks for minima; seems to evolve his strategy as he goes; spots that Right hand side has no zeros; goes for that.