THE UNIVERSITY *of York*

Discussion Papers in Economics

No. 2005/35

Do People Plan?

by

John D Hey

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

DO PEOPLE PLAN? [1]

John D. Hey

Abstract

We report the results of an experiment investigating a key axiom of theories of dynamic decision making – that agents plan ahead. Inferences from previous investigations have been confounded with issues concerning the preference functionals of the agents. Here, we present an innovative experimental design which is driven purely by dominance: if preferences satisfy dominance, we can infer whether subjects are planning ahead. We find that over half the participants in the experiment do not appear to be planning ahead; moreover, their ability to plan ahead does not improve with experience. These findings identify an important lacuna in economic theories.

[1] The author would like to thank the Economic and Social Research Council (ESRC) of the UK for a 28-month Research Fellowship which permitted this research as well as for funds to run the experiments. He would also like to thank Robin Cubitt and John Bone, in particular, for helpful discussions and John Suckling and Jinkwon Lee for help in running the experiments.

I. Introduction

This paper reports the results of an experimental investigation into one of the key axioms of economic theories of dynamic decision making – namely that, one way or another, economic agents plan the future. Previous experimental investigations into whether human beings actually do plan the future have been be-devilled by problems connected with uncertainty about the preference functionals of the individuals concerned: if the individuals do not have Expected Utility preferences, then one cannot infer from the earlier decisions of the individuals either whether they are making plans for the future, or what they are planning to do in the future. In this paper, we present an innovative experimental design which is driven purely by dominance. Accordingly, as long as the preferences of the participants in the experiment satisfy dominance, then we can infer from their behaviour whether they are planning ahead or not. We find, quite remarkably, that over half the participants in the experiment do not appear to be planning ahead. Furthermore, their ability to plan ahead does not appear to improve with experience. These findings have important implications for economic theories of dynamic decision making. We discuss the background to the experiment in section II; in section III we present the design of the experiment; and the results in section IV. Section V concludes.

II. Background and Motivation

Central to all economic theories of dynamic decision making is the concept of a plan. In order to decide what to do today, rational economic agents should first consider what they will do in the future. This is so whether the agents solve the dynamic decision problem by turning it into a strategy problem, or whether they use some form of backward induction. It is also so whether the agents' preferences are Expected Utility or not, though in this latter case the potential problem of dynamic inconsistency may arise. As this idea of planning is central to economic theorising, many economists have sought to test whether it appears to be valid or not. Economists typically have fought shy of simply asking people – on the methodological grounds that it is difficult or impossible to appropriately motivate an honest answer: if one simply asks the agent, but the agent is not forced to implement the stated plan, then there appears to be no motive for answering honestly; if instead, agents are forced to implement the stated plan, then the problem is transformed into one of pre-commitment. Moreover, asking the subjects what they are planning to do suggests to subjects that they might want to plan – if subjects had never thought about doing so then the design of the experiment at least brings the idea to their mind – and hence, perhaps, defeats the very purpose of the experiment. These issues are discussed with great clarity in the papers of Cubitt *et al* [1998 and 2004]. Psychologists, however, have been less reluctant to simply ask subjects. Prominent amongst such psychologists are Jerome R. Busemeyer and his co-workers who have done a series of experiments on dynamic decision making, two good examples[1] being Barkan and Busemeyer [1999] and Busemeyer *et al* [2000]. These experiments suggest that planning by subjects is not as economists imagine it to be. However, economists remain suspicious of these results – for the methodological reasons outlined above. One experiment in economics which attempted to identify the plan (if any) being made by subjects was Hey [2002] but this relied on the assumption that the subjects had Expected Utility preferences[2]. This assumption is too strong for many economists to accept. In retrospect, it seems that the research agenda of Hey [2002] was too ambitious: it attempted not only to identify the plan of the subject (if it existed) but also whether it was implemented. As Hey [2003] showed, this agenda is impossible without some knowledge of the agent's preferences. Accordingly, this present paper has a more modest agenda: rather than try to identify the subject's plan and whether it is implemented, this present paper simply tries to see if the subject does indeed plan[3].

III. The Experimental Design

It will prove simplest to begin with an example. Consider the decision tree in Figure I. This is one of the decision trees used in the experiment. The tree was displayed in colour on the computer screen of the subject. The subjects had previously been given written Instructions and a

PowerPoint presentation (which played at a pre-determined speed on their terminals) of these Instructions. These Instructions tell the subjects that this is a decision tree, that the boxes are nodes – the green ones (the first and the third sets moving horizontally) are *decision* nodes at which they take a decision and the *red* ones (the second and fourth moving horizontally) are *chance* nodes at which Nature takes a decision. They were told that Nature chooses Up or Down with equal probability, independently of past moves by Nature and the decision maker, and what this implies was explained to the subjects at some length. The tree is played out sequentially – with the subject deciding Up or Down at the first decision node (that at the left of the tree) and then Nature moving. The subject then decides again and then Nature makes a final move. This leads to one of the *payoffs* on the right hand side of the tree. The crucial feature of this design is the structure of these payoffs. We now discuss this in some detail - as the structure is essential to the purpose of the experiment.

In the top half of the tree the payoffs are those in the first row of Panel A of Table I; in the bottom half of the tree the payoffs are those in the second row of Panel A of Table I. All the payoffs are denominated in pounds sterling.

One crucial feature is revealed if we order the entries in these two rows. Doing this, we get Panel B of Table I. It is immediately clear that the entries in the first row *dominate* those in the second row[4] if one assumes for the time being that all 8 numbers in each row are equally likely.

Let us examine first the decision of the subject at the first decision node – that at the left hand side of the tree. Of course, the participants know that they will, after Nature has made her first move, be taking a second decision at the second decision node. However, if the subjects ignore the fact that they will make this second decision, then one obvious procedure is to compare the two rows in Panel A or Panel B of Table I; and to choose Up at the first decision node if the first row is preferred to the second row, and to choose Down at the first decision node if the second row is preferred to the first row. It should be clear that, whatever the preference functional of the subject, the most preferred row is the first row – it *dominates* the second – if one assumes that all 8 outcomes in each half of the tree are equally. So Up would appear to be the best decision at the first decision node[5].

But this is not what the subjects should be doing if they are planning ahead. If they do plan ahead they will plan the following, depending upon which of the four second decision nodes they reach. At the first (the top) of the second decision nodes they will choose Down (because Down leads to either 16 or 8, while Up leads to either 8 or 13 - and because the pair (16,8) *dominates* the pair (8,13)); at the second of the second decision nodes they will choose Up (because Up leads to either 6 or 20, while Down leads to either 6 or 18 - and because the pair (6,20) *dominates* the pair (6,18)); at the third of the second decision nodes they will choose Up (because Up leads to either 15 or 17, while Down leads to either 2 or 4 - and because the pair (15,17) *dominates* the pair (2,4)); at the fourth (the bottom) of the second decision nodes they will choose Up (because Up leads to either 20 or 8, while Down leads to either 8 or 0 - and because the pair (20,8) *dominates* the pair (8,0)).

Anticipating this future behaviour, choosing Up at the first decision node leads to one of the four payoffs in the first row of Panel C of Table I, while choosing Down at the first decision node leads to one of the four payoffs in the second row of Panel C of Table I. If we put these in numerical order, we get Panel D of Table I.

It is immediately clear that the entries in the second row *dominate* those in the first row[6]. Whatever the preference functionals of the subjects - as long as they respect dominance[7] - all subjects will prefer the second row to the first, and hence should choose Down at the first decision node *if, as we have assumed in this discussion, they plan ahead* and anticipate that particular payoffs will be eliminated by their own future decisions. So Down at the first decision node is the actual best decision to take (assuming that the subject's preference functional satisfies dominance). Contrariwise, it follows that Up is the wrong decision, and is so whatever reasons are induced to support it (for example, those above, where implicitly the decision maker is assuming that he or she will choose randomly in the future[8]).

All the trees used in this experiment have what we call this *dominance property*. What we mean by this is that one decision at the first decision node appears to be optimal if the individual does not plan ahead and does not eliminate, whereas the other decision is in fact optimal if the individual does plan ahead and eliminate (assuming throughout that the preference functional satisfies dominance). This simple property enables us to test whether individuals do indeed plan ahead or not.

We should note that this property remains true if individuals do not use backward induction but instead use some alternative method of planning the future. The obvious alternative is that of choosing the optimal *strategy* for the decision problem. A strategy is simply a set of conditional decisions – the conditioning being on the move by Nature at the first chance node. Let us use the notation *(A;B,C)* to indicate a possible strategy: this notation implies that the decision at the first decision node is *A*; and that at the second decision node is *B* if Nature has moved Up at the first chance node, and *C* if Nature has moved Down. There are clearly 8 possible strategies to be considered: *(U;U,U), (U;U,D), (U;D,U), (U;D,D), (D;U,U), (D;U,D), (D;D,U)* and *(D;D,D)*. Each of these leads to 4 possible outcomes. The implied outcomes are listed in Table II. In Panel A we list the outcomes in their original order reading from the top of the tree; in Panel B we list the outcomes in numerical order. We note that, for all strategies, the four respective outcomes are equally likely. It is clear from Table II that the strategy *(D;U,U)* dominates all the other strategies. So playing Down at the first decision node is optimal for a decision maker who plans ahead through the choice of an optimal strategy – again assuming that the decision maker's preferences satisfy dominance.

A pilot experiment revealed that more than half the subject pool seemed not to be planning ahead. This could have been for two reasons: first, that their preference functionals violated dominance; second, that the subjects did not fully understand the experiment. Accordingly, in this full-scale experiment reported here, we introduced two changes. First, we got the subjects to play out the whole tree (in the pilot we had just asked them to report their decision at the first node); this gives us information as to what they do at the second node – where the dominance is obvious. Second, we gave them four separate attempts at the tree, and paid them off on a randomly selected one of the four resulting payoffs. As this latter procedure may raise some eyebrows, we should comment further on it. Normally, there are problems in repeating an experiment in which the experiment is played out on each repetition – in that the (expected) wealth of the subject is changing throughout the experiment. This could change the preferences of the subject. However, as this experiment is driven solely by dominance, this should not be a problem – we see no reason why changing wealth should lead to a violation of dominance.

Accordingly, the experiment proceeded as follows. We had a total of 55 participants, organised in 4 separate sessions. They were sat at individual terminals in the EXEC laboratory. They were given written Instructions. Then they were shown a PowerPoint presentation (which played at a predetermined speed) which repeated the Instructions and gave more detail. At this point, the subjects had a chance to ask questions, after which they proceeded to the experiment, with each subject being given four separate attempts/trees and being allowed to work through the trees at their own speed. For each subject and for each attempt, the *set* of payoffs was different, though each set satisfied the *dominance property* defined and discussed above. All payoff sets had two properties: first, that the mean of the payoffs in the actually-dominated half of the tree was at least £2.50 higher than the mean of the payoffs in the actually-dominating half of the tree; second, that the expected payoff making the right decision was always at least £2.50 more than the expected payoff taking the wrong decision. So the wrong decision appeared to quite a lot better than the right one (for those subjects not planning ahead), while the actual payoff from the right decision was quite a lot larger than the actual payoff from the wrong decision. All the payoffs were integers in the range from 0 (pounds sterling) to 20 (pounds sterling). After completing all four attempts, the subject called over an experimenter and the subject was paid in cash a randomly chosen one of the payoffs on the four attempts at the experiment. We now discuss the results.

IV. The Experimental Results[9]

The first crucial point is whether the subjects' preferences respected dominance or not. This is easily answered by looking at the numbers of the decisions at the second decision node which respected dominance. Out of the 55 subjects in the experiment at the second decision node, 52 had decisions on the first attempt which respected dominance, 54 on the second, 54 on the third and 54 on the fourth[10]. (Interestingly, it was almost always a different subject on each attempt whose second decision violated dominance.) One could perhaps argue that, except in one rather blatant case (see the details in the footnotes), these violations of dominance were 'rather minor'. The conclusion, however, is very clear: virtually all the subjects on virtually all the attempts took decisions at the second decision node which respected dominance. Therefore our underlying assumption seems to be valid. In the light of this, we continue to refer to a decision at the first node which respects both dominance and the presumption that subjects plan ahead as the 'correct decision'.

Now let us move back to the first decision node. Out of the 55 subjects, 19 took the correct decision on the first attempt, 21 on the second attempt, 19 on the third attempt and 15 on the fourth attempt. This gives an overall total of 74 correct decisions out of 220 – just 34% correct. The conclusion seems to be clear – the majority of these subjects do not plan ahead. Moreover, the hypothesis that the first decision is taken at random is rejected at the 1% level (t-stat = 2.57), with the deviation going in the wrong direction – more incorrect decisions being made than correct decisions. Moreover, the pattern of correct decisions through the four attempts does not show any improvement through time – experience does not improve the incidence of correct decisions. The fact that the majority of subjects do not appear to plan ahead does not go away with experience.

These aggregate figures do however hide some individual variations – it is clear that some subjects always plan ahead while others never do: of the 55 subjects, 23 never made a correct first decision, 11 made a correct first decision only once, 7 only twice, 7 three times and just 7 got the first decision correct every time. Again one can reject the null hypothesis that all subjects chose at random – in favour of the hypothesis that they are more likely to take the incorrect decision. However, perhaps a better conclusion is that a minority of the subjects did plan ahead while the majority did not. Indeed, suppose we postulate that a proportion p always intend to choose what is the wrong decision, and that a proportion $(1-p)$ always intend to choose what is the right decision, but that both groups implement their intended decision with a tremble t (that is, they choose what they intend to choose with probability $(1-t)$ and choose what they did not intend to choose with probability t). Suppose further that we estimate p and t to minimise the sum of squared differences between the expected frequencies and the actual frequencies, then the estimate of p is 0.718 and that of t is 0.134. On this interpretation, therefore, 71.8% of the subjects intend to choose what is the wrong decision and only 28.2% intend to choose what is the right decision – while both groups implement their intended decisions with probability 86.6%.

A formal probit analysis confirms our findings. We carried out a probit with the dependent variable d talking the values 1 for a correct decision and 0 for an incorrect decision, and with independent variables as follows: a the attempt number (taking values 1, 2, 3 and 4) and n the decision node number (taking values 1 and 2). We found the following, where Z denotes the latent variable such that $d = 1$ if $Z > 0$ and $d = 0$ if $Z <= 0$ (and where the numbers in parentheses are t-ratios):

$$\text{(1)} \qquad \begin{array}{l} Z = -2.72 - 0.021a + 2.34n \quad ll = -180 \quad n = 440 \\ \qquad\ (9.0) \quad\ (0.3) \qquad (12.0) \end{array}$$

The variable a has a negative sign (indicating that subjects are more likely to make an incorrect decision the more experience they have) but is not significant. The variable n is significant – and this reflects the fact that almost all subjects' decisions respected dominance at the second decision node.

The experimental software recorded the time t to take each decision. Interestingly, if we include this variable in the probit, and drop the variable a, we have the following result:

(2)
$$Z = -3.61 + 2.67n + 0.0155t \quad ll = -158 \quad n = 440$$
$$(11.2) \; (12.7) \quad (4.0)$$

So the more time that the subject spends considering the decision, the more likely he or she is to take the correct decision. However, the magnitude of the effect should be taken into account: across all subjects the mean time to take a decision was 22.575 seconds with a standard deviation 21.636. Going from 2 standard deviations below the mean to 2 standard deviations above increases t by 43.272 and thus leads to an increase in Z of 0.67, and hence (starting from $Z = 0$) increases the probability of a correct decision by 0.25.

We also collected some elementary demographic data on the subjects. The subjects were all students at the University of York so we collected data on their degree and their year of study. We also recorded their age and their sex. We included various combinations of these variables into our probit equation. The only one that was anywhere near significant was gender. Denoting by s a dummy variable which takes the value 0 for male and 1 for female, we get the following result.

(3)
$$Z = -3.47 + 2.68n + 0.0148t - 0.305s \quad ll = -157 \quad n = 440$$
$$(10.5) \; (12.6) \quad (3.8) \qquad (1.9)$$

Gender is on the borderline of significance at 5%. The magnitude suggests that, on average, women are 0.12 less likely to take the correct decision as men. This should be set against the fact that at the first decision node, only 34% took the right decision, while at the second decision node 97% did so. This gender effect is confirmed, and perhaps shown better, in Table III, which simply reports the percentages of correct decisions at the first and second decision nodes by gender. The effect is there but we hesitate to speculate how general it is.

V. Conclusions

There are two clear conclusions emerging from this experiment: the majority of subjects do not behave as if they are planning ahead; for the majority of subjects, experience does not improve their ability to behave as if they are planning ahead. As this result conflicts so severely with accepted wisdom amongst economists, it deserves some comment.

First, one could argue that the decision problem, as presented, is either too complicated or unfamiliar. We find it difficult to conceive of a simpler presentation. We could try presenting it in words, but we suspect that that would simply make matters worse. At the very least there would be more noise in the data. We would assert that there is not much noise in this data – confirmed perhaps by the fact that there is not much variation from attempt to attempt. Perhaps we should have given them more attempts, but we suspect that many of the subjects are not learning through experience. We should be careful here in that it is not obvious that experience is synonymous with learning – particularly in the context of this experiment. First, while a subject will obviously observe the outcome of his decision process, it is not immediately apparent to the subject what would have been the outcome if he or she had chosen differently. Indeed, enquiring into what might have been the outcome implies an approach to analysing the decision problem similar to that which the economist uses: a subject who thinks about what might have been may well have taken the correct decision in the first instance. Moreover, if the subject got an 'acceptable' outcome then he or she might not be provoked into thinking whether his or her decision was correct or not; note, in this respect, that the probability of getting an 'acceptable' outcome is quite high. Turn again to Table 1: the subject who does the wrong thing at the first decision node will get at least £16 with probability one-half, while the same is true if he or she does the correct thing. Clearly, in the long run, he or she is going to get more money doing the right thing, but it may take many more than 4 repetitions for this difference to become clear. Following on from this, it could also be argued that the incentives were not strong enough, despite our structuring of the data files to give a reasonably strong incentive to doing the correct thing. This could be tested by repeating the experiment with the payoffs increased by some factor *f*. It would be interesting to see if increasing *f* decreased the number of wrong decisions.

Some people have suggested to us that our findings may be due to some misunderstanding of what Nature does. Nature might be perverse: for example, 'punishing' a subject who was lucky before. There seems some truth in this: if someone has survived on one occasion, he or she may be more risk averse on the next – not wanting to 'push their luck' too much. But, seeing as this whole experiment is driven by dominance, it is not clear in what way Nature might be perverse. One possible argument is the following, perhaps best illustrated by an example, using the notation *(a,b)* to denote a risky prospect that yields £*a* with probability 0.5 and £*b* with probability 0.5. Suppose the decision maker is choosing between (10,20) and (6,14). Obviously the first dominates the second, but the decision maker may argue that, if he or she chose the first, then Nature would 'punish' him or her by giving the outcome £10, whereas if he or she chose the second, the Nature would 'reward' him or her by giving the outcome £14. True, this is one possible story, but, as it does not explain the second decision node behaviour, then it is unlikely to explain the first decision node behaviour. Moreover, we are not aware of any theories in economics that incorporate a malicious or judgemental Nature.

There are clearly alternative ways of interpreting our experimental evidence. One is through observing that the decision problem is complex, particularly at the first decision node. At the second decision node, the decision is straightforward and the dominance there should be obvious. However, it could be argued that the dominance at the first decision node is not so obvious, even for someone who plans ahead. This could be for two reasons, the first of which is that the subjects have to realise that the four implied outcomes are all equally likely. In principle this requires the compounding of probabilities, but in practice we could argue that this experiment is particularly simple in its characterisation of the probabilities – we deliberately specified that Nature was equally likely to move Up or Down. Nevertheless, we accept that some subjects may have been doing the wrong thing (despite planning) because of their inability to compound probabilities. It is difficult to see how one might design an experiment which eliminates this possibility – given that the chance moves are necessary to get trees with the dominance property. As this point is important, we should state it more explicitly – without the chance moves it is impossible to get trees with this dominance property – and hence it is impossible to have an experimental test solely of planning without any implicit assumption concerning the subjects' preferences. We should emphasise that our experimental design is unique in this respect and therefore that our findings do add substantially to previous experimental results, all of which have relied on making some assumption (either explicitly or implicitly) about subjects' preferences.

A second possible interpretation of our results is that some subjects may have been doing the wrong thing (despite planning) because the dominance at the first decision node is not as obvious as at the second node. Consider the example given in Table I – and consider the payoffs in their 'natural order' from the tree. To see the dominance at the second decision nodes one has to see that (16,8) dominates (8,13), that (6,20) dominates (6,18), that (15,17) dominates (2,4), and that (20,8) dominates (8,0). All of these may be 'obvious'. In contrast, for the backward inductor, the dominance of (15,17,20,8) over (16,8,20,6) has to be seen. This is less obvious. For a strategy player, the dominance of *(D;U,U)* over all the other strategies has to be seen. From Panel A of Table II it is clear that this is not immediately obvious – though it becomes more obvious from Panel B of Table II. So one failure of 'planning' may simply be the inability to see dominance when it is not obvious. This suggests that planning issues may be compounded with complexity – what we interpret here as the failure of planning may simply the increasing disability of subjects to solve increasingly complex decision problems. It is not clear how one might disentangle these two effects, though in a sense it may not be necessary to do so if these things always go hand in hand – whether it is the failure of planning or the inability to solve complex problems, the outcome is still *as if* the subjects were not planning. That surely is an important message.

It could also be remarked that we have 'tricked' the subjects into doing the wrong thing – by using trees with the dominance property. This makes the wrong choice look best for those subjects who do not plan ahead and hence who do not eliminate rejected outcomes from the set under consideration. Such a criticism misses the point. For someone who does plan ahead, the apparent dominance of the worst decision is irrelevant – and hence anyone doing the wrong thing cannot be planning ahead. This is our bottom line. But there is a second point, by using trees with the dominance property we have gained some insight into the decision process of those people who do not plan ahead. As well over half did the wrong thing – it seems to be the case that the non-planners where looking at the complete set of payoffs when taking their decision. So, in addition to the conclusion that the behaviour of the majority of subjects was not consistent with planning ahead, we emerge with a tentative hypothesis for future study: while thinking ahead, subjects do not plan ahead. They think of all possibilities but do not eliminate those that should/will be rejected by their future choice. This tentative conclusion is reinforced by the fact that 23 out of the 55 subjects *always* took the wrong decision at the first decision node.

In conclusion, we think that the bottom line is clear: the majority of subjects do not plan ahead in the way that economists say they should. While they may well *look* ahead, the act of planning ahead in the way that economists conceive it is alien to many human beings. Our next plan is to explore the implications of this finding for economic theories that involve dynamic decision making.

REFERENCES

Barkan, Rachel and Jerome R. Busemeyer "Changing Plans: Dynamic Inconsistency and the Effect of Experience on the Reference Point", *Psychonomic Bulletin and Review,* VI (1999) 547-554.

Bone, John D, John D. Hey and John R. Suckling "Do People Plan Ahead?", *Applied Economics Letters*, X (2003) 277-280.

Busemeyer, Jerome R, Ethan Weg, Rachel Barkan, X. Li and Zhengping Ma. "Dynamic Consequential Consistency of Choices Between Paths of Decision Trees", *Journal of Experimental Psychology*: General, CXXIX (2000) 530-45.

Carbone Enrica and John D. Hey, "A Test of the Principle of Optimality." *Theory and Decision,* L (2001) 263-281.

Cubitt, Robin P, Chris S. Starmer and Robert Sugden, "Dynamic Choice and the Common Ratio Effect: An Experimental Investigation." *Economic Journal*, 1998, CXIIX (1998) 1362-80.

___________, _________ and ___________ "Dynamic Decisions Under Uncertainty: Some Recent Evidence from Economics and Psychology," in *The Psychology of Economic Decisions, Volume 2: Reasons and Choices* edited by Brocas, Isabelle and Carrillo, Juan D. Oxford University Press, (2004).

Hey, John D. "Experimental economics and the Theory of Decision Making Under Risk and Uncertainty." *Geneva Papers on Risk and Insurance Theory*, XXIIX (2002) 5-21.

Hey, John D. "Are Revealed Intentions Possible?" unpublished paper (2003).

Figure I

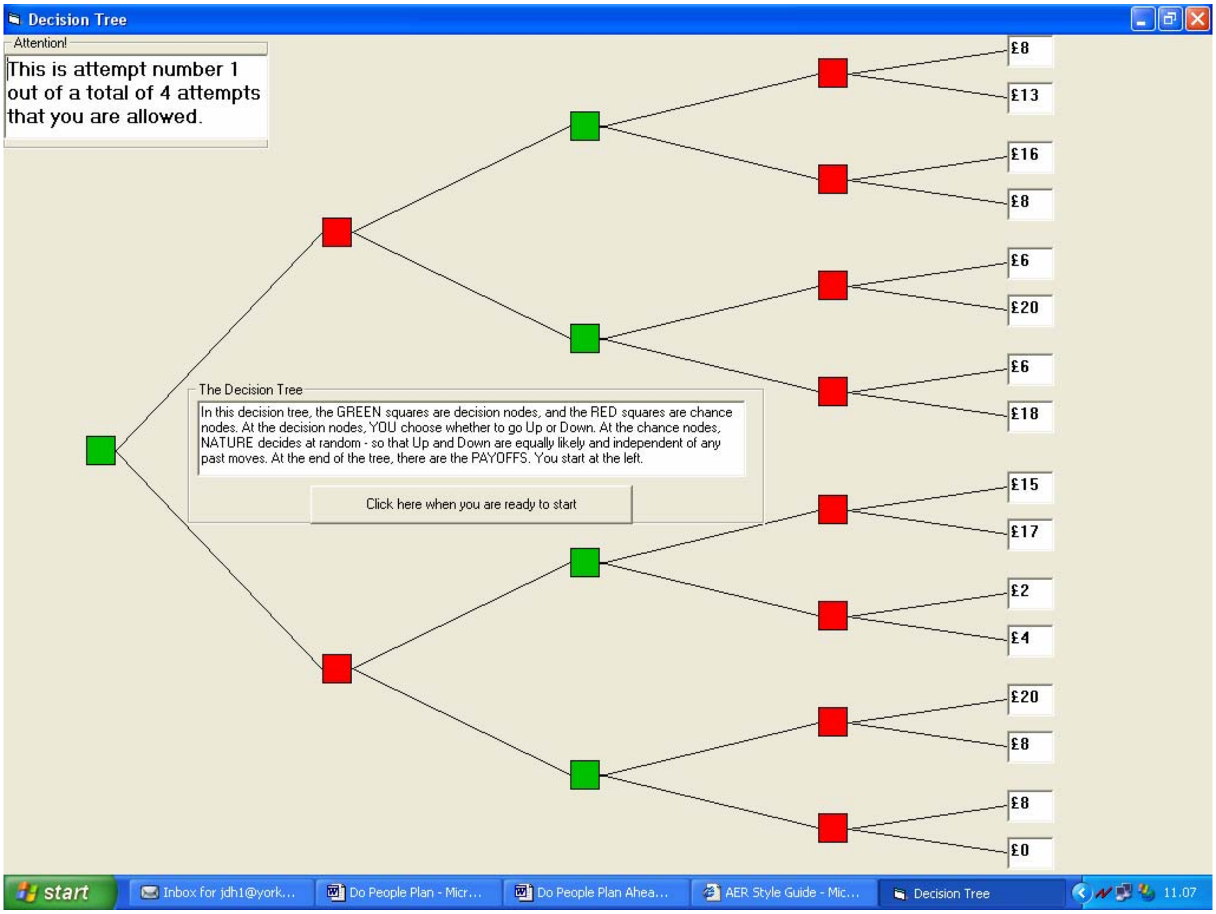

Decision Tree
Attention!
This is attempt number 1 out of a total of 4 attempts that you are allowed.
The Decision Tree
In this decision tree, the GREEN squares are decision nodes, and the RED squares are chance nodes. At the decision nodes, YOU choose whether to go Up or Down. At the chance nodes, NATURE decides at random - so that Up and Down are equally likely and independent of any past moves. At the end of the tree, there are the PAYOFFS. You start at the left.
Click here when you are ready to start
£8
£13
£16
£8
£6
£20
£6
£18
£15
£17
£2
£4
£20
£8
£8
£0
start
Inbox for jdh1@york...
Do People Plan - Micr...
Do People Plan Ahea...
AER Style Guide - Mic...
Decision Tree
11.07

Table I: the payoffs in the tree in Figure I

Panel A: In their original order

Top half	8	13	16	8	6	20	6	18
Bottom half	15	17	2	4	20	8	8	0

Panel B: The payoffs in Panel A in numerical order

Top half	20	18	16	13	8	8	6	6
Bottom half	20	17	15	8	8	4	2	0

Panel C: The payoffs after eliminating those that will be rejected by future choice

Top half	16	8	20	6
Bottom half	15	17	20	8

Panel D: The payoffs in Panel C in numerical order

Top half	20	16	8	6
Bottom half	20	17	15	8

Table II: Possible strategies and their implied possible outcomes

Panel A – Outcomes in their original order

Strategy	Outcome 1	Outcome 2	Outcome 3	Outcome 4
(U;U,U)	8	13	6	20
(U;U,D)	8	13	6	18
(U;D,U)	16	8	6	20
(U;D,D)	16	8	6	18
(D;U,U)	15	17	20	8
(D;U,D)	15	17	8	0
(D;D,U)	2	4	20	8
(D;D,D)	2	4	8	0

Panel B – Outcomes in numerical order

Strategy	Highest outcome	Second highest outcome	Second lowest outcome	Lowest outcome
(U;U,U)	20	13	8	6
(U;U,D)	18	13	8	6
(U;D,U)	20	16	8	6
(U;D,D)	18	16	8	6
(D;U,U)	20	17	15	8
(D;U,D)	17	15	8	0
(D;D,U)	20	8	4	2
(D;D,D)	8	4	2	0

Table III: Gender Effects

	percent correct at first decision node	percent correct at second decision node	numbers
male	39.5	98.4	31
female	26.0	95.8	24
all	33.6	97.3	55

Written Experimental Instructions (not intended for publication)

EXEC

INSTRUCTIONS

Welcome to this experiment. It is an experiment on the economics of dynamic decision making under risk. The Economic and Social Research Council of the UK (ESRC) has 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. There is no participation fee for this experiment – what you are paid at the end depends partly on the decisions that you take during the experiment and partly on chance. At the end of the experiment you will be asked to sign a receipt for any payment that you received, and to acknowledge that you participated voluntarily in the experiment. The results of the experiment will be used for the purpose of academic research and will be published in such a way that your anonymity will be preserved.

The Experiment

The experiment concerns a *Decision Tree*. This decision tree is simply a short sequence of decisions to be taken by you, interlaced with moves taken by Nature. Nature is a random device, representing risk, whose behaviour will be explained to you shortly. Each sequence of decisions by you and moves by Nature leads to a payoff, which is an amount of money. You will be allowed several attempts at the decision tree. On each attempt there will be a payoff, denominated in money. Your payment for participating in this experiment will be one of these payoffs – chosen at random from the set of payoffs on the various attempts that you will have completed.

The Decision Tree

The *Decision Tree* is characterised by a sequence of *decision* and *chance* **nodes**. At each node there two subsequent paths to follow: Up and Down. At each *decision node* **you** will have to take a *decision* - in each case whether to go Up or Down. At each *chance node* a chance device - which we call **Nature** - will determine whether Up or Down is chosen. Nature operates in a totally random way – so that Up and Down are equally likely and independent of any past moves either by you or Nature. In total there are *two* decision nodes and *two* chance nodes, starting with a *decision* node and then alternating the two types until the final *chance* node. So the entire sequence is: decision, chance, decision, chance. After the second and final chance node is played out you will arrive at an *end node*. Each *end node* has associated with it a *payoff* - an amount of money. The payoffs associated with each end node are written in the end nodes.

Nature

‘Nature’ is our way of describing a totally random device. It is important that you understand what this means. At any chance node, when Nature moves, it moves in such a way that Up or Down are equally likely and independent of any moves made by you or by Nature at any time. This means that it is impossible to predict what Nature is going to do and the only information on which you can work is simply that Up and Down are equally likely. It may be useful to you to note that the way that Nature is implemented on the computer is through using the random number generating mechanism of the computer software. Even with this knowledge you are unable to predict any move of Nature.

The Various Attempts

You will be allowed several attempts at the tree. The several attempts are all independent of each other. In particular, the moves by Nature in one attempt are independent of the moves by Nature in other attempts. Moreover, there is no reason why your moves should not be independent – but those decisions are entirely up to you: your decisions on any one attempt are not in any way constrained by what you decided on other attempts. Your decisions are entirely up to you – though obviously your payment will depend on what you decide. The basic *structure* of the tree will remain the same from attempt to attempt, in the sense that there will always be a decision node, then a chance node, then a decision node, then a chance node and then a payoff node, in each attempt. Moreover, Nature will always behave completely randomly. **The one thing that will differ from attempt to attempt is the set of payoffs. You should therefore carefully check the set of payoffs on each attempt.**

Your Payment for Participating in the Experiment

As we have already remarked, you will be allowed several attempts at the tree. The precise number of attempts will be told to you when you start the experiment, and you will be reminded throughout of how many attempts you have done and how many remain to be done. On each attempt there will be a *payoff*, denominated in money. Your payment for the experiment will be a randomly chosen one of these payoffs. For example, suppose you are allowed 4 attempts at the tree. There will be 4 payoffs – one for each attempt. At the end of the experiment, you will be invited to call over one of the experimenters. He or she will have 4 cards, numbered from 1 to 4. These cards will be shuffled and you will be invited to pick one of the cards (obviously without seeing the number written on it). The number on the card that you pick will be noted and you will be paid the payoff on that numbered attempt.

How the Experiment will Proceed

The experiment will begin with a PowerPoint presentation of these Instructions. Then you will turn to the experiment itself. The opening screen displays the **EXEC** logo. When everyone is ready to start, the **EXEC** logo will disappear. You will then be told how many attempts at the tree you will be allowed. The decision tree will then be displayed. You should study this carefully, and particularly the various possible end (payoff) nodes. You will end up at one of these payoff nodes in any one attempt. You will then be invited to work through the tree, starting at the left-hand node, which is a decision node. At each decision node, you will be asked to indicate whether you want to move Up or Down and then you will be asked to confirm your decision by clicking on the button “Click here to confirm”; your decision will then be implemented, with the part of the tree that your decision has excluded turning grey to indicate that that part is no longer available. At each chance node, you will be asked to get Nature to move by clicking on the button “Click here to get Nature to move”; you will then be told the move by Nature, and it will be implemented, with the part of the tree that Nature’s move has excluded turning grey to indicate that it is no longer available. After the second and final move by Nature, you will see that only one end (payoff) node remains available. This is the payoff for that attempt.

The end of the experiment

After you have completed all the attempts, the **EXEC** logo will once again be displayed, along with a message informing you that the experiment is over. The message will also list the payoffs on the various attempts At this point, you should call over one of the experimenters. He or she will then carry out the procedure described above for determining your payment for the experiment. He or she will ask you to complete a brief questionnaire and will pay you your payment. You will be asked to sign a receipt for the payment.

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.

THANK YOU FOR YOUR PARTICIPATION

Author

John D. Hey

Department of Economics, University of York, Heslington, York, YO10 5DD, United Kingdom

and

Dipartimento di Scienze Economiche e Aziendale, Libera Universita Internazionale degli Studi Sociali (LUISS), Via O. Tommasini 1, 00162 - Roma ROMA, ITALY.

e-mail: jdh1@york.ac.uk

[1] We should note that the work of Busemeyer and his colleagues is far-reaching and goes far beyond the concerns of this present paper. Here we refer only to issues concerned specifically with planning.

[2] If the agents preferences are non-Expected Utility, then different sequences of decisions may be implied depending upon whether the agents use the strategy method or some method of backward induction. There is also the thorny problem of how such agents may resolve any potential problem of dynamic inconsistency.

[3] As this distinction is important, we should elaborate on it here. In the Hey [2002] experiment, the idea was to separate the subjects into two groups – (1) those moderately risk-averse and (2) those less risk-averse and risk-loving. The idea was that different routes through the tree would be planned by the two different groups. Moreover, this would enable us to predict the second choice from the first – if the subjects plan. In this present experiment, however, there is only one route through the tree for those subjects who plan.

[4] By which we mean first order stochastic dominance.

[5] Note, however, that this line of reasoning assumes implicitly that the decision maker chooses randomly in the future. On this, see later.

[6] Note that all four outcomes are all equally likely.

[7] In addition, the subject should be able to attribute dominance to his or her (future) self.

[8] Indirect evidence that this is in fact what subjects implicitly assume can be found in Carbone and Hey [2001].

[9] The data is available in a variety of formats. There is a file of input data, called *tinp.01* and a raw file of output data *all.dat*. There is also a document "A Guide to the Data Files" which describes these data files. The experimental software can also be made available in the form of a compiled Visual Basic 6 program (the use of this software requires the approval of the author).

[10] Using *(a,b)* to denote a 50-50 gamble between *a* and *b*, we note that, on the first repetition, of the three whose decisions violated dominance, one chose *(19,13*) in favour of *(20,14)*; one chose *(12,3)* in favour of *(16,5)* and one chose *(17,3)* in favour of *(18,8)*; on the second repetition, the one who violated dominance chose *(2,1)* in favour of *(18,16)*; on the third repetition, the one who violated dominance chose *(19,6)* in favour of *(20,6)*; and on the fourth repetition, the one who violated dominance choose *(19,17)* in favour of *(20,17)*.