



THE UNIVERSITY *of York*

Discussion Papers in Economics

No. 2006/22

Do People Plan?

by

John Bone, John D Hey and John Suckling

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

Do People Plan?¹

John Bone, John D Hey and John Suckling

Abstract

We report the results of an experimental investigation of a key axiom of economic theories of dynamic decision making – namely, that agents plan. 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 or not. We implement three sets of experiments: the first two (the Individual Treatments) in which the same player takes decisions both in the present and the future; and the third (the Pairs Treatment) in which different players take decisions at different times. The two Individual treatments differed in that, in one, the subjects played sequentially, while, in the other, the subjects had to pre-commit to their future move. In all contexts, according to economic theory, the players in the present should anticipate the decision of the player in the future. We find that over half the participants in all three experimental treatments do not appear to be planning ahead; moreover, their ability to plan ahead does not improve with experience, except possibly when we force subjects to pre-commit to their future decision. These findings identify an important lacuna in economic theories, both for individual behaviour and for behaviour in games.

Address for correspondence: John D Hey, Department of Economics and Related Studies, University of York, Heslington, York YO10 5DD, UK. Tel: 0044 1904 433 786 Fax: 0044 1904 433 759. E-mail: jdh1@york.ac.uk

Bone: University of York, UK, jdb1@york.ac.uk

Hey: LUISS, Italy and University of York, UK, jdh1@york.ac.uk

Suckling: University of York, UK, jrs1@york.ac.uk

¹ The authors would like to thank the ESRC for funds to finance the experiment through a Research Fellowship to Hey (RES-000-27-0077), and the Department of Economics and Related Studies at the University of York for providing research income for Bone and Suckling. The authors would also like to thank Jinkwon Lee for assistance and discussions. Finally the authors would like to thank the Editor, Tim Cason, and two referees for extremely helpful comments which lead to a substantial improvement in the paper.

Do People Plan?

John Bone, John D Hey and John Suckling

1. 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. This should be true both in individual decision problems and in games: individuals should consider their own future behaviour when planning earlier decisions in a dynamic decision problem; players in games should consider the future moves of the other players in the same game.

As far as individual decision problems are concerned, previous experimental investigations into whether human beings actually do plan the future have been confounded by problems connected with uncertainty about the preference functionals of the individuals concerned: if the individuals do not have Expected Utility preferences, then one may not be able to 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 found, in an early pilot experiment, that it seemed to be the case that very few people were acting as if they were planning ahead. We wondered whether this was anything to do with individuals being unsure or unthinking about their own future behaviour, whereas they may well think about the behaviour of others. For this reason, we decided to implement a second experimental treatment: in which the player at the future decision node was replaced by a different player. This latter takes us into the realms of game theory and we discuss the implications in the remainder of this section. However, it should be kept in mind that these games experiments were principally motivated by evidence from individual experiments.

Later, after being prompted by insightful comments from a referee, we implemented a further new individual treatment in which the player was forced to pre-commit to his or her second move when taking the first decision. We shall talk more about this treatment later.

As far as behaviour in games is concerned, it has usually been taken for granted that players think about the future decisions of the other players in the game. In this paper, we can test, using our innovative experimental design, whether this is, in fact, the case. We assume, paralleling our assumptions about individuals, that all players' preferences satisfy dominance and that all players know that that is so.

We ran three treatments – two Individual Treatments (differentiated by whether subjects were asked to pre-commit to their future decision or not) and one Pairs Treatment. We call these INPC (Individual Non-Pre-Commitment), IPC (Individual Pre-Commitment) and P (Pairs) treatments. We find that in all three treatments 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, except marginally in the Individual Pre-Commitment treatment. These findings have important implications for economic theories of dynamic decision making, both applied to individual decision problems and behaviour in games. We discuss the background to the experiment in section 2; in section 3 we present the design of the experiment; and the results in section 4. Section 5 concludes.

2. Background and Motivation

We start with economic theories of *individual* dynamic decision making. Central to all such theories 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 whether they are 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 being Barkan and Busemeyer (1999) and Busemeyer *et al* (2000).² 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.³ 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

² 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.

³ 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.

plan and whether it is implemented, this present paper simply tries to see if the subject does indeed plan.⁴

Turning now to behaviour in games, it could be argued that a test of whether, and what, players think about the other players is unnecessary in that there have been innumerable experimental investigations of behaviour in games. A good example is the early experimental work of Beard and Beil (1994). However, our test is different in that it makes no suppositions about the preferences of the players other than they respect dominance. Moreover, the test does not rely on any assumptions about the presence or absence of altruism or any other confounding issues concerning how the players feel about the others – in our experiment, the payoffs to the two players are identical, and hence there is no conflict of interest. In contrast, the interpretation of previous experimental studies has been confounded by issues concerning altruism, reciprocity and other other-regarding preferences. Our experimental design isolates the key element of this research – whether people plan ahead.

3. The Experimental Design

It will prove simplest to begin with an example. Consider the decision tree in Figure 1. This is one of the decision trees used in the Individual Treatments. 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. The Instructions obviously differed between the Individual Treatments (in which the subject played by him or herself) and the Pairs Treatment (in which there were two subjects, Player 1 and Player 2, playing together but without communication), and we shall distinguish between these in what follows by putting any variations in the Pairs Treatment in square parentheses [].

⁴ 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.

These Instructions tell the subjects that this is a decision tree, that the boxes are nodes – the first and the third sets moving horizontally (the black ones in Figure 1, though they were green in the actual software) are *decision* nodes at which they take a decision and the second and fourth moving horizontally (the white ones in Figure 1, though they were red in the actual software) 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; what this implies was explained to the subjects at some length. The tree is played out sequentially (except in the IPC treatment, where subjects were asked, when taking the decision at the first node, to pre-commit to their decisions at the second node) – with the [Player 1] subject deciding Up or Down at the first decision node (that at the left of the tree) and then Nature moving. Then the [Player 2] subject decides Up or Down at the second decision node and then Nature makes a final move. This leads to one of the *payoffs* on the right hand side of the tree. In the Pairs Treatment *both* players received the same payoff – so there was no conflict of interest between the two players. 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 1; in the bottom half of the tree the payoffs are those in the second row of Panel A of Table 1. 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 1. It is immediately clear that the entries in the first row (first-order stochastically) *dominate* those in the second row 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 [Player 1] at the first decision node – that at the left hand side of the tree. Of course, the subject [Player 1] knows that he or she [Player 2] will, after Nature has made her first move, be taking the decision at the second decision node. However, if the subject [Player 1] ignores the fact that he or she [Player 2] will make this second decision,

then one possible procedure is to compare the two rows in Panel A or Panel B of Table 1; 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 likely. 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 is reached. At the first (the top) of the second decision nodes the subject [Player 2] 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 the subject [Player 2] 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 the subject [Player 2] 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 the subject [Player 2] 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 1, while choosing Down at the first decision node leads to one of the four payoffs in the second row of Panel C of Table 1. If we put these in numerical order, we get Panel D of Table 1.

It is immediately clear that the entries in the second row dominate those in the first row as all the four outcomes are equally likely. Whatever the preference functionals of the subjects - as long as they respect dominance - 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*

⁵ Note, however, that this line of reasoning assumes implicitly that the decision maker chooses randomly in the future.

ahead and anticipate that particular payoffs will be eliminated by their own [Player 2's] future decisions.⁶ 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 incorrect decision, and is so whatever reasons are invoked to support it.

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 people do indeed plan ahead or not.

In the Individual Treatments 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. It can be shown (see the *extra supplementary material* available through the journal website) that, in the tree above, the strategy of choosing Down at the first node and then Up at whatever second node is reached dominates all other strategies. This is true in all trees with this dominance property

Note also the crucial significance of both players receiving the same payoff at whatever final payoff node is reached. This removes any consideration of fairness and other other-regarding motives, and makes the comparison of two final payoffs particularly simple: if at node 1 both players get a whilst at node 2 both players get b , then both players will prefer node 1 to node 2 if $a > b$, irrespective of any feelings of envy, spite, altruism and so on. Moreover, if one choice leads to either a or b (with equal probability) for both players, while another leads to either c or d (again with equal probability) for both players, then both players will prefer the first to the second if the pair (a,b) first-order stochastically dominates the pair (c,d) . Similarly, if one choice leads to one of a, b, c or d (all with equal probability) for both players, while another leads to one of e, f, g or h

⁶ In addition, the subject should be able to attribute dominance to his or her (future) self, or Player 1 should be able to attribute dominance to Player 2.

(again all with equal probability) for both players, then both players will prefer the first to the second if the foursome (a,b,c,d) first-order stochastically dominates the foursome (e,f,g,h) .

A pilot experiment of the Individual Treatment 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 the three treatments of the 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 for the INPC and P treatments proceeded as follows. In the INPC Treatment 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 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. (The payoffs in the trees used in the experiment are available via the journal website). 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 correct decision was always at least £2.50 more than the expected payoff taking the incorrect decision. So the incorrect decision appeared to quite a lot better than the correct one (for those subjects not planning ahead), while the actual payoff from the correct decision was quite a lot larger than the actual payoff from the incorrect decision. All the payoffs were integers between (and including) £0 to £20. 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.

In the P Treatment, a total of 52 subjects participated in the experiment – divided into 26 pairs. In each pair, one player was chosen to be Player 1 and the other Player 2. Neither player knew the identity of the other player; there was no way that they could communicate with each other. In all other respects the design of the experiment was exactly the same as in the Pairs Treatment, except for minor variations in the paying of the subjects. After completing all four attempts, the two subjects in a pair independently called over an experimenter and each subject was paid in cash a randomly chosen one of the payoffs on the four attempts at the experiment. As before, we note that the use of this *random lottery incentive mechanism* does not create problems in this particular experiment as the whole experiment is driven by dominance – so changing expected wealth through the experiment should not affect behaviour. The subjects completed a very brief questionnaire and signed a receipt.

After running these two treatments and obtaining results which seemed to indicate that people were not planning ahead, we realised that it would be useful to run a third treatment – the IPC treatment – in which subjects were asked to pre-commit to their decision at the second decision node when taking their decision at the first.

4. The Experimental Results⁷

⁷ The data is available in a variety of formats via the journal website. There is a file of input data and three raw files of output, one for each treatment. There is also a document “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 (though the use of this software requires the approval of the authors).

We begin with a discussion of the data from the INPC treatment. The data are summarised in Table 2. 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.⁸ (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.⁹ 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. A one-tailed test of the hypothesis that the first decision is taken at random is rejected at the 1.1%, 4.0%, 1.1% and 0.4% level on attempts 1, 2, 3 and 4 respectively, with the deviation going in the wrong direction – more incorrect decisions being made than correct decisions. 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.

⁸ 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 chose $(19,17)$ in favour of $(20,17)$.

⁹ Though we should note that non-violation of dominance at the second decision node does not necessarily imply non-violation at the first node – since the first node decision is more complex and the dominance more difficult to 'see'.

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.

Let us do a more formal analysis. Suppose we postulate that a proportion p of the population from which our subjects were drawn always intend to choose what is the incorrect decision, and that a proportion $(1-p)$ always intend to choose what is the correct decision. Let us suppose further that all subjects 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 incorrect decision and only 28.2% intend to choose what is the correct decision – while all subjects implement their intended decisions with probability 86.6%. Again a large majority seem to be doing the incorrect thing.

We now turn to an analysis of the data from the Pairs Treatment. Again, these are summarised in Table 2. As before, we start with the second decision nodes and the decisions of the subjects who were Player 2. In the 26 pairs, 25 took the correct decision (the dominating one) on the first attempt, 25 on the second, 24 on the third and 24 on the fourth. It is interesting to note that in 23 of the pairs, Player 2 always took the correct decision, while there was one Player 2 who took the incorrect decision on one attempt, another Player 2 who took the incorrect decision on two attempts, and one Player 2 who managed to take the incorrect decision on three attempts. The

conclusion is that the vast majority of the Player 2's took the correct decision, though one is left wondering about the three Player 2's who managed to take incorrect decisions.¹⁰

Let us now turn to the crucial decisions from the point of view of the hypothesis we want to test: the decisions of the Player 1's. Of the 26 pairs, only 5 took the correct decision on the first attempt, 9 on the second, 9 on the third and 9 on the fourth. While it could be argued that there was some improvement between the first and the second attempts, there is no evidence of any further improvement with experience. We must admit that we had expected to see such an improvement – because in the vast majority of the pairs Player 2 took the correct decision – and so Player 1 (if he or she had wanted to be) could have been in no doubt about the rationality of Player 2. If we combine together the data from all 4 attempts, we see that just 32 out of 104 first decisions were correct. This is 30.8%. A one-tailed test of the null hypothesis that the first decision was taken at random is rejected at the 0.1%, 5.9%, 5.9% and 5.9% level on attempts 1, 2, 3 and 4 respectively, with the departure being in the wrong direction – implying that the subjects were more likely to take the incorrect decision rather than the correct decision. The conclusion seems to be clear – the Player 1's do not think ahead and anticipate the Player 2's decision as Game Theory would have us believe.

We should note that these aggregate figures hide some individual variations – it is clear that some players always think about the other player while others never do: of the 26 Player 1's, 10 never made a correct decision, 7 made a correct decision only once, 4 only twice, 3 three times and just 2 got the decision correct every time. Again one can reject the null hypothesis that all Player 1's chose at random – in favour of the hypothesis that they are more likely to take the incorrect

¹⁰ Here we give some detail on the violation of dominance at the second node. The one Player 2 who took the incorrect decision on one attempt did so by going for (0,4) instead of (10,12), where we list the equally likely outcomes for each of the two choices, putting first the one near the top of the tree. This seems to be quite a major mistake. The Player 2 who took 2 incorrect decisions went for (20,13) instead of (14,20) and went for (7,19) instead of (9,20). These errors look relatively minor. The Player 2 who took the incorrect decision on 3 attempts went for (19,3) instead of (6,20), for (20,6) instead of (20,10) and for (5,1) instead of (20,2). While at least one of these could be classified as a minor error, at least one is major. It is difficult to find explanations for some of these decisions.

decision. However, perhaps a better conclusion is that a minority of the Player 1's did think about the other player while the majority did not.

At this stage, it was pointed out to us that we should carry out a further individual treatment in order to shed some light on what is going on. In this further treatment, when taking the decision at the first node, subjects were asked what they wanted to do at the second node – after Nature has moved. This required them to give two conditional decisions – one if Nature moved Up and one if Nature moved Down. These conditional decisions were then implemented. Obviously this makes sense only with individuals – and so we call this third treatment the Individual Pre-Commitment, or IPC, treatment. Here we have rather more data, and can analyse the correctness of both conditional decisions. We had 45 subjects on this treatment: on the first attempt, 85 of the 90 conditional second-stage decisions were correct; on the second attempt, 83 were correct; on the third, 86 were correct; and on the fourth, 81 were correct. Once again the vast majority of the second stage decisions respected dominance. Table 2 gives the details, in the same form as the INPC and P treatments, of the actual decisions. It will be seen that, at the second node, of the 45 subjects in the experiment, 43 had decisions on the first attempt which respected dominance, 41 on the second, 44 on the third and 40 on the fourth. Once again, dominance seems to be largely respected at the second node. At the first decision node, of the 45 subjects, 13 took the correct decision on the first attempt, 14 on the second attempt, 15 on the third attempt and 18 on the fourth attempt. This gives an overall total of 60 correct decisions out of 180 – just $33\frac{1}{3}\%$ correct – remarkably close to the 34% in the INPC treatment. The conclusion seems to be clear – the majority of these subjects do not plan ahead. Once again, a one-tailed test of the null hypothesis that the first decision was taken at random is rejected at the 0.2%, 0.5%, 1.3% and 9.0% level on attempts 1, 2, 3 and 4 respectively, with the departure being in the wrong direction – implying that the subjects were more likely to take the incorrect decision rather than the correct decision. However, in this treatment the pattern of correct decisions through the four attempts does show some modest improvement through time –

though even by the fourth attempt only 40% are taking the correct decision. One can only speculate as to whether this would improve further with time.

Finally, let us turn to a comparison of the data from the three Treatments, using the data summarised in Table 2. At the second decision nodes, the overall percentage taking the correct decision was 97% in the INPC treatment, 93% in the IPC treatment and 93% in the P treatment. At the first nodes, the respective overall percentages were 34%, 33% and 31%. These are not statistically significantly different one from the other. It seems in all cases that the majority of the subjects were not planning ahead.

V. Conclusions

We set up an experiment with two decision nodes, for which one move at the first decision node is correct if the subjects are planning ahead (either about their own future behaviour or about the behaviour of the other player) while the other move is incorrect. All we need to support this prediction is that subjects' preferences do not violate dominance (and, in the pairs treatment, that this is common knowledge). We carried out three treatments of this experiment: two individual treatments (one without pre-commitment and one with) and one pairs treatment. In all three treatments we find that the majority of the subjects are taking the incorrect decision at the first node. Moreover, behaviour does not improve with experience, except marginally in the Individual Pre-Commitment treatment. Our conclusion is stark: subjects do not plan ahead. What is going on?

We find the results of the IPC treatment particularly illuminating in shedding some light on the possible answers to this question. We must admit that we were reluctant to implement this IPC treatment as we thought that the results would be clear: forcing the subjects to pre-commit to their second move would force them to think more carefully about their first move. We expected many more correct first-stage decisions. However, this does not appear to have been the case – unless one can argue that the slight improvement through the different attempts in the IPC treatment is

evidence of learning and understanding. This may well be true, but we think that there is another factor at work.

In the IPC treatment we asked subjects to pre-commit to the possible second moves – conditional on what Nature did. However, we did not ask them to pre-commit to the counter-factual conditional moves if they had chosen differently at the first stage. As this point seems to us to be important, we should be more specific. Consider Figure 1. Suppose they chose Up at the first node. We then asked them to pre-commit to a move at the top second-stage decision node and a move at the second-to-the-top second-stage decision node. However, we did not ask them what they would do if they were at the second-to-the-bottom second-stage decision node nor at the bottom second-stage decision node. So while they might understand the possible payoffs from choosing Up at the first node we did not force them to think about the possible payoffs from choosing Down at the first node. *Asking them to pre-commit to the second-stage decision does not get them to start thinking like economists.* We are possibly forcing them to realise (in the context of the tree in Figure 1): “if I choose Up then my payoffs will be one of £16, £8, £20 and £6”. But we are not forcing them to realise the counter-factual – “if I chose Down then my payoffs will be one of £16, £17, £20 and £8”. They are not being forced to face the real choice.

Of course this is speculation at this stage. To test this idea, we would need to carry out a further experiment in which they were asked to take conditional decisions *at all possible future decision nodes*. But there are problems in so doing. What rationale do we give the subjects for doing this? What incentive do they have for reporting honestly what they would do at nodes that they have no intention of reaching? And, are we not almost forcing them to think like economists? Does this not destroy the point of the whole exercise? In any case, we have discovered a deeply disturbing anomaly.

VI. References

Barkan R, and Busemeyer J R (1999), “Changing Plans: Dynamic Inconsistency and the Effect of Experience on the Reference Point”, *Psychonomic Bulletin and Review*, 6, 547-554.

Beard T R and Beil R O, (1994), “Do People Rely on the Self-Interested Maximization of Others? An Experimental Test”, *Management Science*, 40, 252-262.

Bone J D, Hey J D and Suckling J R (2003), “Do People Plan Ahead?”, *Applied Economics Letters*, 10, 277-280.

Busemeyer J R, Weg E, Barkan R, Li X and Ma Z (2000), “Dynamic Consequential Consistency of Choices Between Paths of Decision Trees”, *Journal of Experimental Psychology: General*, 79, 530-45.

Cubitt R P, Starmer C S and Sugden R (1998), “Dynamic Choice and the Common Ratio Effect: An Experimental Investigation” *Economic Journal*, 108, 1362-80.

Cubitt R P, Starmer C S and Sugden R (2004), “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.

Fehr E (2003), Theories of Fairness and Reciprocity – Evidence and Economic Applications. Invited Lecture at the 8th World Congress of the Econometric Society. In: M Dewatripont, L Hansen and S Turnovsky (Eds.), *Advances in Economics and Econometrics – 8th World Congress*, Econometric Society Monographs, Cambridge, Cambridge University Press, (with K Schmidt)

Fehr E (forthcoming), The Economics of Fairness, Reciprocity and Altruism, *Handbook of Altruism, Gift Giving and Reciprocity*, Amsterdam, North Holland (with K Schmidt).

Hey J D (2002), “Experimental economics and the Theory of Decision Making Under Risk and Uncertainty” *Geneva Papers on Risk and Insurance Theory*, 28, 5-21.

Hey J D (2003), “Are Revealed Intentions Possible?”, Discussion Paper, University of York.

Figure 1: Individual treatment: an example of the tree as viewed by the player at the first node




Table 1: the payoffs in the tree in Figure 1

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 2: Decisions

<i>Treatment</i>	first attempt		second attempt		third attempt		fourth attempt		all attempts	
	correct	incorrect	correct	incorrect	correct	incorrect	correct	incorrect	correct	incorrect
Decisions at Second Decision Node										
INPC	52 95%	3 5%	54 98%	1 2%	54 98%	1 2%	54 98%	1 2%	214 97%	6 3%
IPC	43 96%	2 4%	41 91%	4 9%	44 98%	1 2%	40 89%	5 11%	168 93%	12 7%
P	25 96%	1 4%	24 92%	2 8%	24 92%	2 8%	24 92%	2 8%	97 93%	7 7%
Decisions at First Decision Node										
INPC	19 35%	36 65%	21 38%	34 62%	19 35%	36 65%	15 27%	40 73%	74 34%	146 66%
IPC	13 29%	32 71%	14 31%	31 69%	15 33%	30 67%	18 40%	27 60%	60 33%	120 67%
P	5 19%	21 81%	9 35%	17 65%	9 35%	17 65%	9 35%	17 65%	32 31%	72 69%