Contents lists available at ScienceDirect

# Journal of Economic Psychology



journal homepage: www.elsevier.com/locate/joep

# Strategies in dynamic decision making – An experimental investigation of the rationality of decision behaviour

John D. Hey<sup>a,\*</sup>, Julia A. Knoll<sup>b</sup>

<sup>a</sup> University of York, UK and LUISS, Italy <sup>b</sup> FactWorks GmbH, Berlin, Germany

#### ARTICLE INFO

Article history: Received 27 May 2009 Received in revised form 9 December 2010 Accepted 23 February 2011 Available online 27 February 2011

JEL classification: D9

PsycINFO classification: 2340 2260

Keywords: Dynamic decision making Strategy method Backward induction Heuristics

# ABSTRACT

This paper is concerned with the question of how people tackle dynamic decision problems. It is on the interface between economics and psychology. Economic theory has a well-defined theory of how people should tackle such problems, but experimental evidence suggests that these are not empirically valid, and particularly that people find dynamic decision problems complex and cognitively demanding. Psychologists have long been aware of such issues and have developed a suite of theories to explain behaviour in such contexts, but these have been largely developed in a static context. This paper attempts to build a bridge between the two disciplines by exploring decision processes in a dynamic problem for which economic theory provides clear predictions. To aid us in this quest we use an experimental design which enables us to infer the decision rules that people are using. We identify a number of distinct decision metristics, which could usefully be embodied into economic models of dynamic decision making.

© 2011 Elsevier B.V. All rights reserved.

# 1. Introduction

Though economists and psychologists have long been concerned with the principles and processes of decision making in conditions of risk, little is known about the way people actually tackle stochastic dynamic decision problems. Such problems are especially demanding as their solution requires a sequence of interdependent decisions which takes random changes in the decision environment – so called 'moves of nature' – into account as well as the impact of previous decisions, that is, the earlier moves of the decision maker (Edwards, 1962). This sequence of decisions is also referred to as 'decision strategy' in psychology. According to Beach and Mitchell (1978) a decision strategy is (a) a procedure the decision maker engages in, when attempting to choose among alternative courses of action and (b) a decision rule that dictates how the results of those procedures will be used to make the final choice. Up to now the question how such a decision strategy is chosen cannot be answered unequivocally. The objective of this paper is to provide further insight to dynamic decision behaviour and thereby contribute some clarification.

This paper reports on the decision-making processes that subjects appear to be using in a dynamic decision-problem which has a well-defined optimal solution. We deliberately and crucially adopt an experimental design in which preferences

E-mail address: jdh1@york.ac.uk (J.D. Hey).

0167-4870/\$ - see front matter @ 2011 Elsevier B.V. All rights reserved. doi:10.1016/j.joep.2011.02.011



<sup>\*</sup> Corresponding author. Address: Department of Economics and Related Studies, University of York, Heslington, York YO10 5DD, UK. Tel.: +44 1904 433789.

play no role in defining the optimal solution; the only thing we assume is that preferences satisfy *dominance*; this seems a minimal requirement. The assumption of dominance has various levels: the simplest is that the decision-maker, when faced with a choice between a certain amount of money *x* and a certain amount of money *y*, he or she chooses *x* if *x* is bigger than *y*. We require an extension of this to stochastic choice: let  $X = (x_1, x_2, ..., x_N)$  denote a choice which leads either to a payoff of  $x_1$ , or to a payoff of  $x_2$ , ..., or to a payoff of  $x_N$ , each with probability 1/N. Then we say that the preferences of a decision-maker satisfies dominance if he or she chooses X over Y if  $x_n \ge y_n$  for all n = 1, 2, ..., N with at least one of the inequalities being strict. We assume that our subjects' preferences satisfy dominance in this sense.

We present subjects with a dynamic decision problem in which there are three decision nodes interleaved with three chance nodes. The problem is carefully constructed so that there is a uniquely best optimal strategy for someone whose preferences satisfy dominance.<sup>1</sup> However, because of the way that the problem is presented, the implementation of the optimal strategy is not obvious to someone who is not trained. We are therefore interested in whether subjects can disentangle the complexities of the problem and (perhaps learn to) approach the optimal solution. Our experiment is designed to see what subjects actually do.

We begin our brief discussion of the contributions made by psychologists with the paper by Beach and Mitchell (1978) which suggest classifying decision strategies into *aided-analytical, unaided-analytical* and *non-analytical* strategies. As these names suggest, the three categories differ in their analytical degree, in the amount of required resources, and in the amount of information procurement. Aided-analytical strategies require the application of a prescribed procedure, and usually decision tools (pencil and paper, mathematics, calculators) are used to derive the implied decision. The decision process may be complex and time-consuming since all the relevant information is considered and processed as the decision strategies al-ways require training or invention. In contrast, unaided-analytical decision strategies do not make use of tools. Instead the decision-making processes is entirely carried out in the decision maker's mind. Non-analytical decision strategies comprise simple rules. These are fast decision strategies since little information is processed and the decomposition is omitted. The question that now arises is how the decision maker decides which strategy she will use in any particular decision problem.

Generally it is assumed in psychology that the decision maker has a repertoire of decision strategies from which he or she chooses (Beach & Mitchell, 1978), (Gigerenzer, Todd, & the ABC Research Group, 1999) and (Payne, 1976). This choice is contingent on two factors: the characteristics of the decision problem and the characteristics of the decision maker. The decision problem not only refers to the decision task but also comprises the decision environment (Beach & Mitchell, 1978). Decision tasks can differ in their degree of unfamiliarity, ambiguity of the problem, instability and complexity, while the decision situation can differ in respect of the irreversibility, significance and accountability of the decision as well as of time and/or money constraints (Beach & Mitchell, 1978). The second factor which influences strategy selection is that of individual differences. According to Beach and Mitchell (1978) decision-makers differ in respect of their knowledge about strategies and their chances of success, their ability to implement a particular strategy and their motivation to solve the decision problem. Bettman, Johnson, and Payne (1990) also emphasize the existence of individual differences in respect to cognitive abilities and capacities.

Beach and Mitchell (1978) incorporate both factors in their *Model of Contingency*, which postulates that strategy selection is determined by a cost-benefit analysis, where the expected costs, comprising time, effort or even money, which will emerge by the application of a certain strategy, are counterbalanced by the expected benefit. More formally expressed, the benefits are computed as the product of the probability that the chosen strategy will lead to a correct decision and the utility of making the correct decision. According to Beach and Mitchell the strategy which maximizes net gain (expected benefit minus expected costs) will be chosen. Beach and Mitchell's Model postulates that the choice for an analytical procedure is positively linear related to task and situational demands, as well as to knowledge and ability. However, if the demands become overwhelming, the linear relation breaks down and a non-analytical decision strategy will be chosen. Each of the demand components is weighted in respect to their importance so that the preference of analytical over non-analytical strategies is contingent on the decision situation. Unfortunately a key component of this approach is that it requires the calculation of "the probability that the chosen strategy will lead to a correct decision and the utility of the making the correct decision" both of which are unknown until the optimal strategy has been calculated. If this has been done, there is no need to use a sub-optimal strategy.

It is clear from the psychological literature that, rather than optimizing, decision-makers may adopt *heuristics* which may or may not approximate to the optimal strategy. Kahneman, Tversky, and Slovic (1982) following on the work of Newell and Simon (1976) were among the pioneers at identifying heuristics that people appeared to be using. Gigerenzer et al. (1999) noted that many heuristics (perhaps after a period of learning) approximated quite well the optimal strategy.

Most of the strategies discussed above mainly refer to static decision problems. But what do we know about decision strategies in dynamic decision problems? The answer seems to be 'not much'. The literature appears to be restricted to the work of Busemeyer and his collaborators (for example, Busemeyer, Weg, Barkan, Li, and Ma (2000), Mueller (2001)) and Carbone and Hey (2001). Although in the context of different dynamic decision tasks, the general conclusion of these studies was that subjects try to simplify the decision problem as much as possible by applying heuristics. However, in all these experiments there were subjects who, though not behaving optimally, tried to implement an optimal strategy but

<sup>&</sup>lt;sup>1</sup> This is whether the individuals solve the problem by backward induction or by the strategy method.

made computational mistakes in so doing. Some subjects seem to try out different kinds of decision strategies, and there were some who seemed not to understand the task.

The bottom line of this section is the following: economists have a well-defined story as to how people should tackle the decision problem posed to subjects in our experiment – but this does not seem to be empirically valid. In contrast, psychologists have a battery of stories about how decision problems are tackled – but very few of these have been investigated in a dynamic context. The purpose of this paper is to provide a dynamic setting in which the economic predictions are clear and testable and to investigate them in a laboratory context in which we can observe the decision processes and hence infer whether any of the psychological stories appear to be valid.

## 2. Experimental design

Our experimental design builds on that used by Carbone and Hey (2001). Dynamic decision problems are represented in the form of a tree. In this experiment, we use what we call 3 + 3 trees. These have three *decision nodes* interleaved with three *chance nodes*. The tree starts with a decision node, and subsequently decision nodes are followed by chance nodes and *vice versa*. After the third chance nodes, there are payoff nodes. In the payoff nodes there are amounts of money which the subject is paid if he or she reaches that node. An example is shown in Fig. 1, in which the subject starts at the bottom and works up through the tree to one of the payoff nodes at the top. At each decision node, there are just two possible decisions – Left or Right. At each chance node, there are just two possible moves by Nature – Left or Right. Subjects are told that Nature moves Left and Right with equal probability and that all moves are independent of each other and of the moves by the subject. Even in this simple 3 + 3 tree branches and payoff nodes proliferate – so that there are 64 of the latter. Each of these contains a payoff denominated in money. The actual payoffs at each payoff node were not shown – and subjects had to click on the node to discover the payoff.

The payoffs in the tree had the property of *dominance*, which is crucial to the design. Payoffs were arranged in such a way that, as long as the subjects' preferences satisfy dominance, there exists a unique optimal solution to the dynamic decision problem. Therefore, if subjects are optimizing, there is a particular solution to the dynamic problem.

We could use the data simply to test economic theory. But we want to do more than that: we want to uncover the decision-making processes used by the subjects. Hence we endowed the experimental software with a special feature: *Notepads*. At each node (whether a decision node or a chance node) subjects could open a Notepad and enter comments and remarks, for example, the computation of the expected value of a move and the resulting decision. Thus, the Notepads were useful to the subjects as a memory and decision aid, and were useful to us in revealing the decision strategy that the subject was using. Additionally, the implementation of these Notepads had the advantage that we did not have to ask subjects directly what they are doing, which reduced the threat that subjects do not change their actual decision behaviour. We should



Fig. 1. 3 + 3 Decision tree.

note that subjects were *not* allowed to use pen or paper or any other recording device – so if they wanted to take and keep notes they had to use the Notepads. In some respects our software is similar to the Mouselab software (see, for example, Johnson, Camerer, Sen, & Rymon, 2002) specifically adopted for this context.

Another crucial feature of the software was the replay function. This enables the experimenters to replay the entire decision process for each subject. Using this feature we could reproduce every move that each subject made – each mouse click, each message left in each Notepad and the actual sequence of mouse and keyboard activity. This feature was invaluable to us in trying to infer the 'strategy' behind the decision process. It was used to produce the results that follow.

The experiment was run at EXEC Laboratory at the University of York. 92 subjects participated. Each subject performed the experiment at his or her own pace on a computer and was screened from the other participants. After a short introduction by the experimenter, subjects then received written instructions, which were repeated in a PowerPoint presentation. Afterwards subjects had the opportunity to ask questions. Subjects had to solve four attempts of the dynamic decision problem, each time with the same set of payoffs. The payment for participation was determined by randomly selecting one of the four payoffs obtained in the course of the experiment.

Before proceeding, we need to make some remarks on the decision task and strategies. We imagined that the majority of subjects were unfamiliar with a decision tree, and hence did not know *ex ante* the appropriate solution to this kind of decision problem. The decision problem remained identical throughout the four attempts, that is, there was no increase or decrease of task complexity between the attempts. However, within the decision problem complexity systematically changed from level to level. Thereby task complexity does not refer to the number of alternatives, which were two (moving left or moving right) on each decision level, but to the number of payoffs that were related to the decision node, and therefore to the amount of information and the number of dimensions of each alternative, respectively. So when working top-down from level three to level one task complexity increases. The optimal strategy is either to use backward induction or to use the strategy method, where subjects examine all possible decision strategies and choose that which is optimal. Either method gives the same solution (because of the dominance property of the tree).

#### 3. Results

The results reported in this section were obtained using the *replay* facility of the software. This facility enables us to see exactly what each subject did during the experiment: each payoff node inspected, each Notepad entry made, and each decision implemented. Moreover it enables us to see the *sequence* in which the subject carried out his or her decision strategy.

It may be useful to illustrate what we would see with a particular strategy. Suppose, for example, that the subject uses backward induction, one of the specifications of economic theory. He or she would start by considering the decision at each of the 16 third-level decision nodes. At each of these, there are two possible decisions - moving Left or moving Right. For each of these two decisions there are two possible payoffs. So the subject would compare the pair obtained moving Left with the pair obtained moving Right. An implication of the dominance property of the payoffs in the tree is that one of these two pairs dominates the other. This determines the appropriate move at that third-level decision node and also determines the pair of possible payoffs from making that move. So the subject, for each of the 16 third-level decision nodes, first opens the four associated payoff nodes, hence determining the correct move, and then inserts that move and the associated two possible payoffs in the Notepad for that third-level decision node. Then the subject would move back to the 4 s-level decision nodes. At each of these decision nodes there are two possible decisions – moving Left and moving Right. For each of these decisions there are two possible third-level decision nodes – and the subject has already worked out the best move and the two associated payoffs for each of them. It follows that each of the two decisions at each second-level decision nodes leads to four possible payoffs (all equally likely). An implication of the dominance property of the payoffs in the tree is that one of these sets of four possible payoffs dominates the other. This determines the appropriate move at that second-level decision node and also determines the four possible payoffs from making that move. So we should see the subject, for each of the 4 s-level decision nodes, first opening the four associated third-level decision nodes, hence determining the correct move, and then inserting that move and the associated four possible payoffs in the Notepad for that second-level decision node. Finally, the subject would move back to the first-level decision. Again there are two possible decisions – Left and Right – and associated with each of them a pair of second-level decision nodes each with four possible payoffs. Moving Left thus leads to one of eight (equally likely) payoffs and so does moving Right. Once again, the dominance property of the payoffs in the tree means that one of these two sets of eight payoffs dominates the other - thus determining the correct decision at the first-level decision node. This would be entered into the Notepad associated with the first-level decision node.

At this stage the subject has worked out his or her strategy. All that remains to do is to implement it. So we should see the subject moving to the first decision node, implementing the decision written in the Notepad; then waiting to see what Nature does; and subsequently implementing the decision already written in the Notepad for the second-level decision node to which his or her first move and Nature's move has moved him or her; then again waiting for Nature's move and afterwards implementing the decision already written in the Notepad for the third-level decision node to which his or her first move and Nature's move has moved for the third-level decision node to which his or her first move and Nature's move has moved him or her; Nature moves a final time and a payoff is obtained. The subject should repeat this procedure on each of the four attempts – while the strategy may remain unchanged the actual payoff may differ because of Nature's random moves.

We have determined what we should see, when using the replay facility, if the subject was using backward induction correctly. In particular, it should be noted that there are two phases: first the determination of the optimal strategy and then

its implementation. In the first phase the subject works backwards; in the second phase forwards. We note that the subject should leave notes in the Notepads associated with all the decision nodes (except possibly the first as the correct decision there can be easily remembered); these notes should not only detail the correct decision at that node, but also the value (in terms of future prospects) of being at that node. There clearly is no need to use the Notepads associated with the chance nodes.

As we note below, we did find some subjects following this backward induction strategy. For others, they were doing something different, but the replay facility enables us to see exactly what they were doing – both in terms of gathering information (and recording it in the Notepads) and in terms of taking decisions. We therefore get some insight into their decision making *process*.

Using this replay function (and, in particular, reading the Notepad entries) we were able to identify four basic types of subjects. These we define as follows:

- 1. Effort-Minimisers or Ignorers.
- 2. Backward Inductors, which we further divide into the three sub-types:
  - 2.1 Rationalists.
  - 2.2 Quasi-Rationalists.
  - 2.3 Simplifiers.
- 3. Forward Workers.
- 4. Strategy Mixers.

Fig. 2 shows the distribution of the different types of decision makers in our experiment. More detail is provided in Table 1, which also provides information on the average time taken in completing the experiment and the average payments to the subjects, for each type. Standard deviations are also included: these show that there is considerable variation in both these variables within each type. Surprisingly there is a negative (albeit small, that is -0.13) correlation between these time taken and payment earned. It should, of course, be noted that chance plays a role, so that it is not necessarily the case that the correct backward induction strategy leads to the highest actual payment. (Indeed, in the case of the 11 Rationalists, who were generally following the optimal strategy, several were unfortunate to get bad draws; though this was generally less the case with the Quasi-Rationalists.) Also care should be exercised in interpreting the time data as relatively slow subjects might have an incentive to adopt shortcuts. Given that the experimental software allowed subjects to go at their own speed, and hence implicitly invited subjects to choose their own trade-off between time and payoff, it is difficult to know what to infer from the time/payoff data. Perhaps in future experiments we should impose a minimum and maximum time for completing the experiment?

The decision strategies adopted by these different types differ in two dimensions. The first dimension is the *direction* of the solution, which could be either backwards, forwards, or partially backwards and forwards. The second dimension refers to the *type* of decision strategy. Here we distinguish between analytical strategies, non-analytical strategies and a mixture of



Fig. 2. Observed types of decision makers in the sample.

Table 1	
The behaviour of the various	types.

Туре	Number of this type	Average time taken (s)		Average payoff (£)	
		Mean	Standard deviation	Mean	Standard deviation
Rationalists	11	1483.27	409.55	11.16	5.47
Quasi-Rationalists	3	1144.67	69.95	13.92	3.58
Simplifiers	15	1016.33	484.14	11.48	4.09
Forward Workers	31	882.19	511.58	12.68	3.74
Effort-Minimisers/Ignorers	22	502.55	281.95	11.38	3.39
Strategy Mixers	10	831.20	393.28	14.93	1.73

the two. There is also another dimension – whether the strategy changed from attempt to attempt. We classify a subject as being in one of the first three strategies listed above if he or she consistently applied that strategy in each of the four attempts. If, instead, a subject implements a different strategy on different attempts then we classify them as type four: Strategy Mixers. It is interesting to note that, among these Strategy Mixers, there was no subject who succeeded in evolving to the correct solution. Therefore we can conclude that, in this experiment, learning, if it was occurring, was incomplete.

We now discuss the different types that we identified.

#### 3.1. The Effort-Minimisers or Ignorers

A particularly extreme method of solving the dynamic decision problem is that of effort minimising or ignoring (information). This method is immediately evidenced by the very small amount of activity of the subjects – their output files (representing their history) are particularly small. This method is characterized either by a denial of the existence of the payoff (that is, subjects did not even click on the payoff nodes to check their values), or an ignorance of the information obtained. In the latter case, subjects seemed to arbitrarily click on payoff nodes, but did not incorporate this crucial information into their decision process. As an extreme example of this ignoring of crucial information we can cite subjects (though admittedly very few) who, on reaching one of the third-level decision nodes, checked the remaining four payoffs and chose the move where the payoffs were dominated, thereby violating our very weak assumption of dominance. By ignoring the payoffs these subjects inevitably worked forward through the decision tree, applying a non-analytical strategy, which due to a lack of Notepad entries in most cases, cannot further be specified. Plausible decision strategies could be 'randomisation of choice', 'choice due to a preference for one side (left or right)' or 'systematic alternation of choice'. The motives for this ignorant behaviour remain somehow clouded and can only be speculated about. In the case of those subjects who violated dominance, a lack of comprehension seems to be the only logical explanation. Others systematically varied their moves apparently in order to determine if moves of Nature were dependent on previous decisions. Either these subjects did not understand the instructions, where it was explicitly mentioned that Nature moves independently of previous moves, or they mistrusted the instructions and expected to be deceived. But the possibility that subjects were either not motivated to devote serious time and effort to the solution, or that subjects simply did not know how to deal with the complexity of the decision problem, should also be taken into consideration. The latter case is predicted by Beach and Mitchell's Model of Contingency (1978). They suppose that excessive computational and cognitive demands implied by the decision problem entail the choice of a non-analytical decision strategy. However, independently of the underlying motivation, subjects in this group preferred to get through the experiment quickly and easily, taking the risk to receive a low payoff. Perhaps they felt intuitively that the costs associated with taking the problem seriously would not be repaid in terms of the benefits from taking the problem seriously. Interestingly, the strategy of Effort Minimising and Ignoring was quite popular among our subjects. In total 24% of our sample were identified as Effort-Minimisers or Ignorers.

#### 3.2. The Backward Workers

The group of Backward Workers constitutes the second largest group (31%) and comprises three subgroups: Rationalists, Quasi-Rationalists and Simplifiers. Backward Workers are characterized by mainly or even completely tackling the decision problem backwards. However, the subgroups differ in their extent of backwardly inducting the optimal solution and therefore in their degree of rationality.

#### 3.2.1. Rationalists

Rationalists behaved according to the predictions of economic theory by backwardly inducting the optimal solution and thus applied an analytical decision strategy. Their decision behaviour can be labelled as completely rational. The optimal solution is derived in a top-down process – as we have described at the beginning of this section. Fig. 3a illustrates the decision process. In total 12% of our subjects showed rational decision behaviour. These results are surprising as optimal behaviour has not been observed before in a similar decision context (Carbone and Hey, 2001). Beach and Mitchell asserted that any aided-analytical decision strategy requires training or invention. For some of the Rationalists, the data indeed suggests that they were familiar with the method of backward induction, maybe due to their studies, since they solved the decision problem straightforwardly. However there were also subjects who did not immediately backward induct the solution but rather worked out this strategy. We can therefore conclude that individuals exist, which when exposed to a new decision problem, think and behave rationally.

#### 3.2.2. Quasi-Rationalists

The decision behaviour of Quasi-Rationalists resembles the decision behaviour of Rationalists except for the inference of the decision on the first level. In contrast to Rationalists, subjects labelled as Quasi-Rationalists spuriously take irrelevant payoff information into account when deriving the decision for decision level one conditional on their decisions on decision level two. This means they consider payoffs that will be eliminated by the move on the second level. Hence, the decision on the first level is sub-optimal, as it is based on some irrelevant information. Still, Quasi-Rationalists apply an aided-analytical though sub-optimal strategy as they are apparently unaware of their mistake. In total only three subjects were identified as Quasi-Rationalists. It seems plausible to assume, that these subjects are Rationalists in reality, whose decision behaviour was



Fig. 3a. The backward workers' strategy.

impaired by a lack of attention and/or cognitive exhaustion. The decision process of the Quasi-Rationalists is shown in Fig. 3b.

# 3.2.3. Simplifiers

In contrast to Rationalists and Quasi-Rationalists, Simplifiers (16%) try to simplify the decision problem by reducing the amount of information. Two types of Simplifiers were identified: the Desperates and the Time and Effort Savers. Their decision process is visualized in Figs. 3c and 3d. In principle Desperates work backward and infer the optimal solution for decision levels three and two by backward induction. But reaching decision level it seems that that they do not know how to further reduce the information. Their only way out of this misery – and that is why this group of subjects is called Desperates – is to move to decision level one and take a bottom-up decision. This could be a random decision or a decision based on intuition or the impression of the payoffs. Afterwards they solve the decision problem according to their notepad entries on decision level two and three. Thus, Desperates first apply an analytical decision strategy, but, as a consequence of the



Fig. 3b. The quasi-rationalists' strategy.



Fig. 3c. Simplifiers - the desperates' strategy.



Fig. 3d. The time and effort savers' strategy.

increase in task complexity on decision level one, they switch to a non-analytical decision strategy. Though in a different context, this decision behaviour was observed by Payne (1976) and is predicted by Beach and Mitchell's Model of Contingency.

Time and Effort Savers at first work bottom-up and apply a non-analytical decision strategy. This reduces the number of decisions to be taken to only six and hence the amount of payoff information that needs to be considered for strategy inference. For the remaining two decision levels, Time and Effort Savers now backwardly-induct the optimal decision – conditionally on their decision for the first decision node. This can be referred to as local optimization,<sup>2</sup> as the optimal decision is only derived for a part of the decision problem and dependent on the previous simplification. This switch from non-analytical to analytical decision strategies was also observed by Payne (1976), who supposed that the simplification in the first place serves to reduce task complexity and therefore cognitive strain.

<sup>&</sup>lt;sup>2</sup> On the contrary, Rationalists optimise globally, as they infer the overall optimal solution to the decision problem.

# 3.3. Forward Workers

Forward Workers are subjects whose decision behaviour is characterized exclusively by working from the bottom upwards. That is, Forward Workers start at decision level one, checking (most of) the 64 payoffs and subsequently make a move. This first decision reduces the number of payoffs to 16, which are considered for deriving a decision for the move on level two. This move then eliminates 12 payoffs so that only four payoffs remain on decision level three. Fig. 4 summarizes the forward-working decision process.

The special characteristic of this decision strategy is that subjects do not make any decisions in advance. That is, Forward Workers do not work out a decision for every single decision node as Backward Inductors do, but solely consider the payoffs related to the decision node that has been reached by their previous decision.

Moreover, Forward Workers generally base their decisions on rather simple decision criteria. It could be inferred from the notepad entries that 'avoiding zero payoffs' is a very popular decision criterion. In using this criterion subjects eliminate branches which contain a zero payoff, or, in the case of several zero payoffs, the branch which contains most zero payoffs. If neither alternative contains any zero payoffs, subjects remove the option which contains the lowest payoff. The popularity of this decision strategy might originate from its simplicity. Overall 'avoiding zero' is easy to derive and apply, as it does not require much computational and cognitive effort. Additionally, it is a rather fast strategy. A slightly more complex, though still rather simple and fast decision strategy that has been applied by Forward Workers is the computation of cut-off points. Here the subject fixes a certain cut-off value, for example, a payoff of £10, and subsequently counts the number of payoffs that exceed this cut-off point. The decision is hence made in favour of the branch which contains more payoffs above the cut-off criterion.

Irrespective of the specific decision criterion the data analysis demonstrated that all Forward Workers base their decisions on non-analytical decision strategies. Though working bottom-up through the decision tree reduces complexity with every decision, a shift from non-analytical to analytical decision strategies was not observed for Forward Workers. In our experiment Forward Working was the most popular decision strategy (33%). This result is consistent with Mueller's (2001) observations made in an experiment on the intertemporal allocation of money. He stated that subjects solve the consumption/saving problems in a forward working manner.

The predominant application of forward-looking decision strategies could originate from a preference for fast and cognitively less demanding decision procedures. Another explanation, as Carbone and Hey (2001) already suggested due to their results, might be a lack of ability to plan ahead, that is, subjects did not realize that they will make decisions in the future which will eliminate certain outcomes.

### 3.4. Strategy Mixers

Strategy Mixers are those subjects that use different strategies in different attempts and therefore could not be assigned unequivocally to one of the former groups. Their decision behaviour resembles Mueller's trial policy, where subjects simply try out a number of different decision strategies.



Fig. 4. The forward workers' strategy.

The inconsistency of decision behaviour could be due to an inability to solve the decision problem. In this case subjects would guess. Another explanation might be that subjects are not satisfied with their decision strategy, somehow realize that this was not the optimal solution and hence are still in search of an improvement of the decision. There are also some subjects who seem to be overwhelmed by the huge amount of payoff information. These subjects only check a few payoff nodes on the first attempt, but then get 'braver' from attempt to attempt and consider more and more information. This intake of new information leads to a change of their decision behaviour.

# 4. Discussion

In our experiment we could identify four different types of decision makers: first, the group of Effort-Minimisers/Ignorers, who working forward through the decision tree ignore most of the information provided and seem to have a clear preference for solving the decision problem with minimal effort. Second, Backward Workers who primarily tackled the decision problem backwards but differed in their degree of rationality, that is, their ability to completely backwardly-induct the optimal strategy. Third, Forward Workers who exclusively solve the decision problem using simple and fast decision strategies in working their way bottom-up through the decision tree and finally the group of Strategy Mixers, who switch strategies between the attempts in search of the appropriate one.

Overall only the Rationalists, a subgroup of Backward Workers, succeeded in deriving the optimal strategy by backward induction, thus behaving according to economic theory. Although their share in our sample is only 12%, the number of rational decision makers in our experiment is much larger than in previous ones of comparable complexity (for example, Carbone and Hey (2001)). Moreover, another 19% of the sample, the Quasi-Rationalists and Simplifiers, also subgroups of the Backward Workers, attempted to solve the decision problem rationally but did not get the optimal strategy 100% correct. As a result they simplified their strategy at some point in the decision process. In contrast to previous research on dynamic decision making (for example, Busemeyer et al. (2000), Mueller (2001) and Carbone and Hey (2001)), we thus observed that a relatively large number of people – the group of Backward Workers, constituting 31% of the sample – at least try to derive the optimal solution and maximize outcome, with more than one third of them even succeeding in doing so. For this reason we cannot conclude that subjects generally use forward-looking decision strategies and try to simplify the decision problem as much as possible by applying heuristics – although forward-working, as executed by the group of Forward Workers (34%) and Effort-Minimisers/Ignorers (23%), remains the dominant way of solving the decision problem. Instead our results show that, in a dynamic decision context of moderate complexity, economic theory does apply at least to a certain degree, even though it cannot generally account for the observed decision behaviour.

But our results not only partially support economic theory. Our experiment also yields empirical evidence for psychological theory, which postulates that strategy choice depends on task complexity and cognitive demand and assumes a strategy switch in case of an increase or decrease in the latter parameters (Payne (1976) and Beach and Mitchell (1978)). This strategy switch could be observed among Backward Workers. Quasi-Rationalists and Desperates, a subgroup of the Simplifiers, switched during their decision making process from a complex to a simple strategy: They started tackling the decision problem by backward induction, but in doing so complexity increases with each decision level as the number of payoff nodes related to a decision node increases when working backwards. Thus, at a certain point in the decision process, namely at level one, these subjects switched to a simpler strategy in order to derive a solution. Time and Effort Savers, also a subgroup of Simplifiers, in contrast switched their strategy from simple to complex: they started at decision level one with a simple decision strategy, thus reducing complexity, that is, the amount of payoff information that needed to be considered, by eliminating one branch of the tree with a forward-looking strategy and only in the second move started backward induction. These results show that psychological theory not only holds for static decisions but can also explain decision behaviour in a dynamic context.

Concerning the Effort-Minimisers/Ignorers and Forward Workers, neither economic nor psychological theory can explain their decision making behaviour. It thus remains unclear whether these subjects were unable to optimize or simply unwilling to do so, preferring a fast and simple decision strategy that reduces cognitive workload over a maximized outcome. Notepad entries of Forward Workers suggest that the majority of these subjects might not have realized that backward induction is the appropriate solution to the decision problem. Due to a lack of notepad entries in the group of Effort-Minimisers/Ignorers their underlying motives can only be speculated about. The decision behaviour of the Strategy Mixers in contrast reveals that these subjects were still in search of the appropriate strategy hence in the first place being unable to solve the decision problem optimally. Still, some of them seem to learn over the four attempts in the sense that they successively considered more information in their decision making process.

Though there is some evidence that irrational decision behaviour can be mainly attributed to a certain inability to derive the optimal solution, other explanations should not be denied and shortly discussed at this point. The subjects' motivation to put effort in the derivation of the optimal strategy as well as a "cost-benefit analysis" in which the subject weighs the benefits at stake (the payoffs) against the costs (the computational effort as well as the time needed to derive and execute the optimal decision strategy) could also be plausible reasons for the observed sub-optimal decision behaviour. Maybe the maximum payoff of £20 in our experiment was just not high enough to outweigh the costs of maximization behaviour. But more experiments will be needed for clarification.

#### 5. Conclusions

The objective of our experiment and this paper was to contribute new insights to the possible answers to the question of how people tackle dynamic decision problems. We have identified three broad types of decision strategies that seem to be applied in a dynamic context (though within each type there are sub-types): backward-working, forward-working, and subjects who seem to operate with no strategy at all or randomly. Only the Backward Workers are those who are consistent with economic theory – though they differ in their ability to completely backwardly-induct the optimal strategy. Forward Workers in contrast crucially seem to ignore the fact that they are going to take decisions in the future; while the others either seem to have a random strategy or try to work out the appropriate one.

In economics it is typically assumed that all agents are backward workers. In psychology instead it is assumed that agents do not maximize but adapt their decision strategy to the complexity of the decision problem and the required cognitive demand. In our experiment we have found evidence for both theories. The conclusion from our results seems to be that neither economic nor psychological theory alone explain dynamic decision making. Rather the truth seems to lie somewhere in between: It seems that economic theory can account for dynamic decision making as long as the decision problem is not too complex and the agent is motivated to derive and apply an analytical strategy, has the time, the required cognitive abilities and decision aids (for example, notepads or pencil and paper) to do so and moreover weighs the benefits of optimization higher than its costs. In every other case decision behaviour can be better explained by psychological theory.

Consequently individual differences and the decision context play a crucial role in dynamic decision making. Still, we feel that we have identified important new decision strategies that are general enough to be observed in different decision contexts with different samples. But this is only a first step and future work in economics as well as in psychology will need to identify general rules of dynamic decision making, taking individual differences and the decision context into account in order to be able to predict and not only describe decision behaviour.

We should note in conclusion that our experiment broadly represents a whole class of economic problems of crucial relevance to economic policy, in particular savings and investment problems. Although a simplification of such problems, the structure of this decision problem is exactly of the same form as these crucial real-world decision problems. As such it helps to shed light on real-world saving and investment decisions.

#### References

Beach, L. R., & Mitchell, T. R. (1978). A Contingency model for the selection of decision strategies. Academy of Management Review, 3, 439–449.
Bettman, J. R., Johnson, E. J., & Payne, J. W. (1990). A componential analysis of cognitive effort in choice. Organization Behavior and Human Decision Processes, 45, 111–139.

Busemeyer, J. R., Weg, E., Barkan, R., Li, X., & Ma, Z. (2000). Dynamic and consequential consistency of choices between paths of decision trees. Journal of Experimental Psychology: General, 129, 530–545.

Carbone, E., & Hey, J. D. (2001). A test of the principal of optimality. Theory and Decision, 50, 263-281.

Edwards, W. (1962). Dynamic decision theory and probabilistic information processing. Human Factors, 4, 59-73.

Gigerenzer, G., Todd, P. M., & the ABC Research Group (1999). Simple heuristics that make us smart. Oxford: Oxford University Press.

Johnson, E. J., Camerer, C., Sen, S., & Rymon, T. (2002). Detecting failures of backward induction: Monitoring information search in sequential bargaining. Journal of Economic Theory, 104, 16-47.

Kahneman, D., Tversky, A., & Slovic, P. (Eds.). (1982). Judgment under uncertainty: Heuristics & Biases. Cambridge, UK: Cambridge University Press.

Mueller, W. (2001). Strategies, heuristics, and the relevance of risk-aversion in a dynamic decision problem. *Journal of Economic Psychology*, 22, 493–522. Newell, A., & Simon, H. A. (1976). Computer science as empirical enquiry. *Communications of the ACM*, 19, 113–126.

Payne, J. W. (1976). Task complexity and contingent processing in decision making: An information search and protocol analysis. Organizational Behavior and Human Performance, 16, 366–387.