Heuristics in dynamic decision making: Coping with the time constants of a dynamic task by doing something else

Berndt Brehmer

Department of War Studies
Swedish National Defence College
P.O. Box 27805
SE-113 85 Stockholm Sweden
berndt.brehmer@fhs.se

Fredrik Elg
Graduate School of Human-Machine Interaction
Department of Computer Science
Linköping University
SE-581 82 Linköping Sweden
freel@gmx.net

Abstract

Laboratory studies have shown that people cannot handle the time constants in dynamic tasks. Yet they obviously cope with such tasks with some success outside the laboratory. This study is one in a series of studies that examine the hypothesis that people cope by relying on heuristics that allow them to simplify the task. The heuristic studied here was that of relying on frequency differences, i.e., what Reason (1990) calls frequency gambling. It examines the effects of varying the relative frequency of scenarios that require different responding, and where relying on frequency rather than learning the actual time constants will lead to some success. The results show that the participants did not learn the time constants, but the heuristic used did not seem to be frequency as much as a heuristic that could be called "better safe than sorry", i.e., they sent out more than the minimum necessary assets to fight the fire. A variant of this heuristic involving rapid and massive responding has also been identified in earlier studies.

KEYWORDS: DYNAMIC DECISION MAKING, FEEDBACK DELAYS, TIME CONSTANTS, HEURISTICS, MICROWORLDS

The business of system dynamics (SD) is modelling dynamic systems. The purpose is to provide aids that can help people cope better with such systems.

This business could certainly be conducted without any reference to how people actually deal with such systems, except, perhaps, the conviction that they are so bad at doing it that SD is their only hope of salvation. So why should SD practitioners be interested in studies of how people actually handle dynamic systems, such at that presented in the paper now before the reader's eyes?

We think that there are at least two reasons. The first is that SD modelling sometimes involves modelling human behaviour in dynamic systems. Such modelling will be as successful as the understanding of human behaviour upon which it is based. It is all too easy to believe that people's understanding of dynamics is just a corrupt version of the "true" understanding provided by, e,g., a proper STELLA model. As we will show here, there are alternative views that point to substantially different ways of handling dynamic systems.

A second reason is that studies of what people actually do may provide new inputs to SD modelling of dynamic systems. In fact, people actually cope with (at least some) dynamic systems successfully also without the aid of SD and have done so for centuries (see Mayr, 1970, for some interesting examples). The question is how they manage to do this despite the obvious lack of understanding of dynamics as documented in numerous studies (e.g., Boot Sweeney & Sterman, 2000; Jensen & Brehmer, 2003; Moxnes, 2000). Perhaps there is a lesson for SD modellers here, and perhaps there are alternative ways of handling dynamics, alternatives to the formal understanding provided by, for example, SD or control theory. This is not to deny that SD modelling would be useful or needed, for it is obvious that people do not cope successfully with all dynamic systems. Rather, alternative forms of understanding based on what people actually do may provide important new sources for modelling dynamic systems. An understanding of these alternatives may also be important for communicating with users of SD models, as well as in the process of developing such models with cooperation between modellers and users. Specifically, such studies may provide a better understanding of intuitive modes of handling dynamics, which sometimes lead to insurmountable barriers in communication.

The importance of time constants

In the real world, everything takes time, so feedback never follows upon action immediately. While this is a trivial insight, and something that no person would dispute, understanding its consequences for the proper strategy in a dynamic task seems to be far from trivial. There is now quite a number of studies that show that even minimal delays wreak havoc with a person's decision making strategy (see e.g., Brehmer, 1995, Sterman, 2000).

What needs to be done to cope with feedback delays depends of the nature of the delay. It is therefore important to identify, not only that there are delays, but also the nature of the delays. Basically, there are three kinds of delays: *dead time* (the interval between the moment when a decision is made and that when the system that the decision maker wants to affect starts to respond), *time constants* (the time required for a decision to take effect), and *information delays* (the interval between the moment when an action has taken effect and that when the decision maker learns about this effect). All three kinds of delays are revealed to the decision maker in the same way: some time will pass until he/she learns about the result after a decision has been made. To identify the nature of the delay requires additional information about the decision task. It is therefore also important to identify what information about the nature of the delays that may, or may not, be available to the decision maker.

Like last year's paper (Brehmer & Nählinder, 2004) the present paper is concerned with the second of the kinds of delays mentioned above: the time constants. When a decision maker tries to control a process (such as a fire) using another process (such as a fire fighting process) and when the control process has appreciable time constants (as a fire fighting process has), the process he or she seeks to control will develop before the decision takes effect. In fire fighting, the fire will spread while the fire fighting units get ready to move out, while they travel to the fire and until they have brought the fire under control. This means that when the decision maker makes his decision, he/she has to compensate for what happens after he/she has made that decision. Specifically, the fire chief cannot only send only the number of units that seem to be sufficient at the time of the decision, he/she must send the number that will be required when the units reach the fire.

In the fire fighting task, the time constant is due to the speed with which the FFUs get going, the speed with which they travel to the fire, and the speed with which they extinguish fire. Knowing these, and the fire conditions, and a model of how fires behave, a decision maker can compute the number of FFUs required when these units reach the fire. Even when using a computational approach, coping with the time constants is clearly not a trivial task. In the experiment described here, however, the participants had to use a more intuitive approach, and base their estimates of the time constants on what they could actually *see*, and they could actually see the time constants happen, so to speak, for the movement of the FFUs and their activities were shown directly on the computer screen.

Earlier studies have shown that people seem to employ a very general rule for coping with the time constants in this kind of task: they respond rapidly and massively, i.e., they learn to send as many FFUs to the fire as rapidly as possible once they have learned of a fire's location (e.g., Brehmer, 1989; 1995). Brehmer and Nählinder (2004) wanted to learn whether this expressed a general heuristic or a well calibrated strategy where the number of units was matched to the future size of the fire. They did this by comparing how people responded to fires requiring different numbers of FFUs in an experiment using a microworld called NEWFIRE (Løvborg & Brehmer, 1991) that simulates forest fire fighting. Their results showed that the subjects did not seem to have a well-calibrated strategy for coping with the time constants, but that they used a heuristic. Specifically, they, used a heuristic which involved positioning their FFUs in such a way that they would not have to distinguish between fires requiring encirclement with multiple FFUs and fires requiring direct attack with one FFU; all fires could be handled in approximately the same way. Incidentally, this is a heuristic used in real fire fighting as well. Many U.S. cities require that fire stations be located so that any burning house in the city can be reached within a specified number of minutes. When the participants were prevented from using this heuristic by requiring them to keep their FFUs at their base until a fire had broken out, their performance grew worse, and there was no evidence that they sent the appropriate number of FFUs to a fire. In short, they seemed unable to compensate for the time constants, except by their heuristic.

Now, coping with dynamic tasks is very much a question of being able to handle the feedback delays, and if people cannot handle even the simplest form of delay, time constants (these delays can be considered as simple because they can be *seen* to happen as the FFUs move to the fire and fight it), except by heuristics, and since other kinds of delays that have been studied experimentally have also proved difficult (Brehmer, 1995), this leads to the hypothesis that people may cope with feedback delays generally by the use of heuristics, if they cope at all. This, in turn, raises the question of what other heuristics people are able to use in dynamic tasks as a substitute for handling the delays as such. The heuristic revealed in the Brehmer and Nählinder (2004) study, although successful in that experiment, is of course, quite specific, and if a person is to cope with delays generally he/she will have to rely on more general heuristics.

Reason (1990) has described two “primitives” that the human cognitive system uses as “fallback positions” when it cannot find the actual rule for a task: *similarity matching* and *frequency gambling*. As the name implies, similarity matching means finding a task that is similar to the current one and doing what one usually does in that situation. Frequency gambling involves relying on what one has learned about differences in frequency of success for different behaviours in the past and gambling on that whatever has been successful in the past will prove useful also in the present situation. It is not clear if this means that they would be maximizing, always choosing the alternative with the highest frequency, or if they would exhibit probability matching, i.e., matching the relative frequency of their decisions to the relative frequency of the relevant outcomes in the task. Both kinds of behaviour would be consistent with Reason’s (1990) frequency gambling hypothesis.

The present paper is concerned with the frequency gambling alternative. Our question is whether people would use frequency gambling as a heuristic and as an alternative to learning the more complex structure required to cope adequately with the time constants in the fire fighting task. That people would choose the former alternative, given a chance, is likely according to earlier results by Lindahl (1974). Lindahl showed that in a complex problem solving task, the presence of an opportunity for simplification in the form of task dimensions that were correlated with the solution would mask the more complex general rule for the task Thus, people who could learn to respond on the basis of just one dimension of a task would not learn the correct rule for the task as well as people who did not have this opportunity and who could only succeed by learning a more complex rule.

Frequency is a candidate for such a simplification. Hasher and Sachs (1984) have shown that learning frequencies is automatic. In tasks where there are differences in frequency with which different behaviours are required, these differences are likely to be picked up, and could then be used as a basis for frequency gambling in Reason’s (1990) sense. Incidentally, real fire fighting is a task where there are frequency differences. The first author has heard fire chiefs remark that most fires are alike, and that they therefore know what to do about them without thinking much about it. Some have also remarked that the rest (in some cases given as about 5%) are very difficult indeed.

If the decisions required by a dynamic task differ in frequency, so that some decisions are successful more often than others, this may serve as a basis for learning how to handle these tasks. Since frequency coding is automatic and thus learned outside consciousness, mastery of a dynamic task based on such learning might well be the basis of what is usually called intuition (see also Hogarth, 2001, for a discussion of frequency learning as a basis for intuition).

The purposes of the present study was to investigate whether people would learn frequency differences and rely on frequency gambling as a basis for their decisions in a dynamic task rather than learning the actual time constants. Specifically, the task was forest fire fighting as represented in a computer simulation called NEWFIRE (Løvborg & Brehmer, 1991, NEWFIRE is described in detail below). Subjects extinguished a series of fires, some requiring just one fire fighting unit (FFU) and some requiring two FFUs with a marked difference in frequency (80% vs. 20%, or vice versa). They were then tested on a new set of fires of the same kind to assess whether they would discriminate between the two kinds of fires or just respond on the basis of the differences in frequencies in the learning set.

Method

Participants

Forty-six undergraduate students from the University of Uppsala were paid 100 sek (about USD 12) to participate. There were 24 women and 22 men and their average age was 24.3 years.

Microworld

The experiment used NEWFIRE (Løvborg & Brehmer, 1991). The NEWFIRE concept is illustrated in Figure 1. NEWFIRE requires the participant to assume the role of a fire chief charged with the task of fighting forest fires. He/she receives information about a fire from a spotter plane and on the basis of this information, he/she sends out the FFUs. These then report back to him/her about the their location and activities, and he/she then uses this information and further information from the spotter plane to issue new commands to the FFUs and the process goes on until the fire(s) has been extinguished. Figure 1 shows the general concept, and Figure 2 the user interface. As can be seen from this description, the task facing the participant has all the characteristics of a dynamic decision task as defined by Brehmer & Allard (1991):

- It requires a series of decisions
- These decisions are not independent (sending the FFUs to one location precludes or at least delays using them elsewhere)
- The state of the task changes both autonomously (due to the strength and direction of the prevailing wind, the character of the forest, and so on) *and* as a result of the decision maker's actions (i.e., where he/she send the FFUs)
- Decisions must be made in real time, i.e., when the developments in the fire requires action, rather than when the decision maker feels good and ready to make them

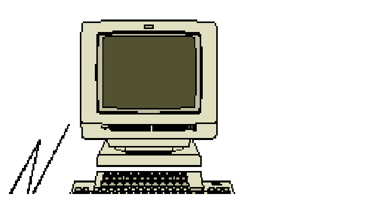

*Figure 1.*The NEWFIRE concept

(For a general discussion of dynamic decision making and the use of microworlds to study it, see Brehmer, 2005; Brehmer & Dörner, 1993).

In NEWFIRE, the participant sees a map depicting a forest (see Figure 2). In this experiment, the forest was homogenous and there were no villages or other objects to protect. The only task was to put out the fire as quickly as possible, using as few FFUs as feasible. Forest was represented by an 18x18 grid with green cells. The start of a fire was signalled by a tone and the cell where the fire started turned red. The fire then spread in the direction of the prevailing wind and with a speed proportional to the strength of the wind according to a general fire model for forest fire propagation. The cells in which FFUs were located were coloured blue. The participant could direct the FFUs to new locations by pointing to a FFU, clicking, pointing to the new location, and then clicking on it. The unit then started to move at the next update of the system. To the right of the map, there was an indicator showing the strength and direction of the prevailing wind, and below this indicator a message panel where messages about the activities and location of each FFU could be seen. Below the message panel was a panel that displayed the last command given to each unit, i.e., the location to which it has been ordered to go. If there was fire in a cell when the FFU arrived, or if fire spread to a cell where a FFU was positioned, it started fighting the fire automatically but only after a delay of one time unit. The time required to extinguish the fire in a cell is a parameter in the program, as is the speed at which the FFUs move. NEWFIRE is a clock driven simulation. The update rate for the screen picture is a parameter in the program.

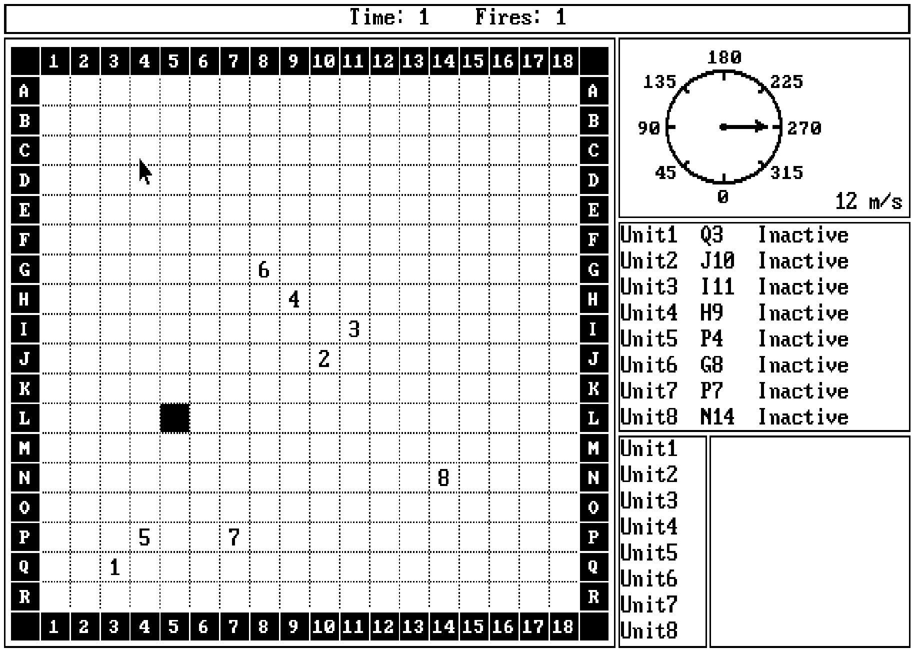

Figure 2. The NEWFIRE interface

In this experiment, NEWFIRE was used in a new way. Rather than having the participants fight every fire until it was extinguished (or until they lost control and the whole forest was destroyed), we studied only the initial commands to each fire. This was because the initial commands very much determine the rest of the fire fighting process. If these first commands are correct, and the appropriate number of FFUs is sent to the appropriate locations, no more needs to be done by the participant; the fire will be extinguished without further work. If they are not correct, there is usually a long process where the participant tries to achieve control. The initial commands are thus the particularly diagnostic of the extent to which a person has learned to compensate for the time constants of the task.

Fire scenarios

Two kinds of scenarios were constructed by positioning the FFUs and choosing a starting location for the fire: scenarios in which the fire could be extinguished by sending just *one* (but the appropriate one, of course) FFU to the fire, and scenarios that required two FFUs to extinguish the fire because it would have spread to two cells as the FFUs moved into position. The number of FFUs required was determined by the time constants: for the 1-FFU scenarios, the FFUs were located in relation to the fire in such a way that it was possible to send one of them directly to the fire and have it extinguish the fire in the burning cell before it would spread to the neighbouring cells. In the 2-FFU scenarios, this was not possible. The participant had to send out two FFUs so as to cover the two cells that would be burning when the FFUs arrived. The participants were informed that there were only these two kinds of fire scenarios, and that their task was to learn to distinguish between them and send out the appropriate number of FFUs to the appropriate location(s) as quickly as possible. As soon as the participants had sent out their FFUs, the scenario was terminated, the program computed what the correct response should be and feedback was given in the form of the word “good” being displayed on the screen if the participant’s decision conformed to the optimal decision as computed by the program. The participant could then choose to see the program actually play the scenario. If their decision was wrong, they were required to do so before they were allowed to go on to the next fire.

Design

The experiment was conducted in two stages: *a learning stage* and a *test stage*. In the learning stage, the participants made decisions for 60 fires with the opportunity for feedback, in the test stage, they responded to 30 fires without any opportunity for feedback. There were two experimental conditions, and they differed only with respect to the relative frequency of 1-FFU and 2-FFU scenarios in the learning stage: In the 80-20 condition, there were 80% 1-FFU scenarios and 20% 2-FFU scenarios, in the 20-80 condition, there were 20% 1-FFU scenarios and 80% 2-FFU scenarios. In the test stage, there was an equal number of 1-FFU scenarios and 2-FFU scenarios.

Procedure

Mouse practice session

In this experiment, it is critical that the participants are able to respond quickly and accurately when using the mouse. The experiment therefore started with a mouse practice session. An 18x18 matrix similar to the map in NEWFIRE was presented. Blue fire fighting units would appear in randomly chosen cells of the matrix with a new FFU being presented every 3 seconds until 99 units have been presented. The participant’s task was to point and click on each FFU as it appeared and move it to a designated area with 10x10 cells to the right of the matrix (see Brehmer & Løvborg, 1992 for further description of this facility in NEWFIRE). The 10x10 matrix could be filled in an arbitrary order.

Learning stage

.
As noted above, there were two kinds of scenarios in the learning stage: scenarios where the fire could be extinguished with one FFU and scenarios which required two FFUs. In all scenarios, eight FFUs were located in different randomly chosen cells of the 18x18 map, but in such a way that the fire, when it appeared, could be extinguished by either one or two of the FFUs. The scenarios differed with respect to the strength (between 1 m/sec and 20 m/sec) and direction (North, East, South, West) of the prevailing wind. The participant’s task was to decide which FFU or FFUs to send to the fire, and to click on the respective units and their destinations. The program only allowed them to send out a maximum of two FFUs in each scenario. The NEWFIRE program then calculated the optimal deployment of FFUs on the assumption that each mouse command required 2.5 sec and provided feedback to the participant in the form of the message “good” if he/she had selected the optimal combination of FFUs and destinations. If the message was not “good”, the subject had to click on the word “demonstrate” and the optimal solution, i.e., which one/two FFU(s) should have been positioned where was displayed and the scenario was played out. The learning stage consisted of 60 trials consisting of either 20% 1-FFU scenarios and 80% 2-FFU scenarios, or vice versa, but before starting the 60 trial learning session, the partici-

pants were given two supervised practice trials. The learning stage required about 35 min.

Test stage

The test stage was the same in both conditions. In this stage, participants were given 30 trials, 15 1-FFU scenarios and 15 2-FFU scenarios, but in this stage, the participants were given no information about whether their decisions had been optimal or not, and they did not have the opportunity to view the optimal solution. This stage required about 15 min.

Results

Learning stage

There was no significant blocks effect when response frequencies were analysed in terms of six blocks of 10 trials each. This suggests, that learning was very rapid, or that there was no learning at all and that the participants simply responded in the same manner throughout the learning stage. Since there was no significant blocks effect, data were pooled over blocks for the subsequent analysis

Table 1 shows the conditional probabilities for 1-FFU and 2-FFU decisions for scenarios requiring 1-FFU and 2-FFU decisions respectively for the two learning conditions.

Table 1. Learning stage conditional probabilities for the two learning conditions.

80% 2-FFU, 20% 1-FFU condition

FFU used/FFU required	**p (FFU used/FFU required) Actual results**	**p (FFU used/FFU required) if partipants had learned time constants**	**p (FFU used/FFU required) if participants rely on frequency gambling by matching**	**p (FFU used/FFU required) if participants rely on frequency gambling by maximizing**
2-2	0.92	1.00	0.80	1.00
2-1	0.55	0.00	0,80	0.00
1-2	0.08	0.00	0.20	0.00
1-1	0.43	1.00	0.20	0.00

80% 1-FFU, 20% 2-FFU condition

FFU used/FFU required	p (FFU used/FFU required) Actual results	p (FFU used/FFU required) if partipants had learned time constants	p (FFU used/FFU required) if participants rely on frequency gambling by matching	p (FFU used/FFU required) if participants rely on frequency gambling by maximizing
1-1	0.83	1.00	0.80	1.00
1-2	0.42	0.00	0.80	1.00
2-1	0.17	0.00	0.20	0.00
2-2	0.58	1.00	0.20	0.00

If the participants had learned to compensate perfectly for the time constants, they should have responded with the high frequency decision when it was required and the low frequency decision when it was required, and there should have been no high frequency decisions when low frequency decisions were required and no low frequency decisions when a high frequency decision were required, as shown in the third column in Table 1. The actual distribution of decisions clearly does not conform to this pattern in either condition. Nor does it conform to the pattern expected if the participants had frequency gambled by maximizing, i.e., always selecting the high frequency decision. In this case, they would always have responded with the high frequency decision, so that the probability of that decision would have been 1.00 for both scenarios requiring the high frequency decision and the low frequency decision, and 0.00 for the low frequency decision for both those scenarios that required that decision and those requiring a high frequency decision. Instead, there is a marked difference between the high and low frequency scenarios. For the former, the pattern of decisions is close to what would have been expected if the participants had made their decisions on the basis of frequency gambling by means of frequency matching. For the low frequency scenarios, on the other hand, the pattern of decisions resembles what would be expected on the basis of random responding, the conditional probabilities for both decisions are close to 0.50. However, no strong conclusions can be drawn from the results from the learning stage since the decision pattern is still being learned from feedback in this stage. For more definitive conclusions, we must turn to the test stage, where the participants were tested without feedback and where we have less reason to expect that they are changing their decision rules.

Test stage

Table 2 shows the overall decision probabilities for 1-FFU and 2-FFU decisions in the two learning conditions.

Table 2. Overall probabilities of 1-FFU and 2-FFU decisions in the two learning conditions

Condition	p(2-FFU decisions)	p(1-FFU decisions)
80% 2-FFU, 20% 1-FFU	0.73	0.27
80% 1-FFU, 20% 2-FFU	0.44	0.56

As the reader will recall, the percentage of scenarios requiring one and two FFUs was the same in both conditions, i.e., 50%. Therefore, if the participants had learned the time constants, we would have expected 50% 1-FFU decisions and 50% 2-FFU decisions in both conditions. As the table shows, the results do not support this interpretation, and they are different for the two conditions. For the 80% 1.FFU, 20% 2-FFU condition, the results come close to the 50% 1-FFU, 50% 2-FFU decisions, but for the 80% 2-FFU, 20% 1-FFU condition, the results agree with what would have been expected if the participants had used a frequency matching strategy. One possible explanation for this would be that the participants learned different things in the two conditions. There is, however, an alternative explanation. Looking at the table, we note that both conditions are alike in that there are too many 2-FFU conditions, compared to what would have been expected, albeit on different grounds in the two conditions. In the 80% 2-FFU, 20% 1-FFU condition there are too many 2-FFU decisions compared to what would have been expected if the participants had learned the time constants, suggesting frequency matching. In the 80% 1-FFU, 20% 2-FFU condition, there are too many 2-FFU decisions compared to what would be expected if the participants were frequency matching, suggesting that they learned the time constants. A possible explanation for this apparent difference is that the participants in both conditions failed to learn, but used a different heuristic, which we may call "better safe than sorry". There is an important difference between making an error in for the two kinds of fire scenarios. i.e., making a 1-FFU decision when a 2-FFU decision is required, and making a 2-FFU decision when a 1-FFU decision is required. In the former case, the participant will loose control over the fire. In the latter case, the fire will be extinguished, albeit at a higher cost, i.e., by using more FFUs than required.

Table 3 shows conditional probability of 1-FFU and 2-FFU decisions for scenarios requiring 1-FFU and 2-FFUs for the two learning conditions in the test stage.

Table 3. Test stage conditional probabilities for 1-FFU and 2-FFU decisions for scenarios requiring 1-FFU and 2-FFUs for the two experimental conditions together with the probabilities that would have been expected if the participants would have learned the time constants, if they had frequency gambled by matching and maximizing.

80% 2-FFU, 20% 1-FFU condition

FFU used/FFU required	**p (FFU used/FFU required) Actual results**	**p (FFU used/FFU required) if partipants had learned time constants**	**p (FFU used/FFU required) if participants rely on frequency gambling by matching**	**p (FFU used/FFU required) if participants rely on frequency gambling by maximizing**
2-2	0.92	1.00	0.80	1.00
2-1	0.49	0.00	0,80	1.00

1-2	0.08	0.00	0.20	0.00
1-1	0.51	1.00	0.20	0.00

80% 1-FFU, 20% 2-FFU condition

FFU used/FFU required	**p (FFU used/FFU required) Actual results**	**p (FFU used/FFU required) if partipants had learned time constants**	**p (FFU used/FFU required) if participants rely on frequency gambling by matching**	**p (FFU used/FFU required) if participants rely on frequency gambling by maximizing**
1-1	0.84	1.00	0.80	1.00
1-2	0.27	0.00	0.80	1.00
2-1	0.16	0.00	0.20	0.00
2-2	0.73	1.00	0.20	0.00

The results for the test stage are similar to those for the learning stage. Thus, there is a marked difference between the high and low frequency scenarios. For the high frequency scenarios, the response pattern is close to what would be expected on the basis of frequency gambling by means of frequency matching, and for the low frequency scenarios, the pattern looks more like random responding. Further analysis shows, however, that for the latter scenarios, there is a tendency to make more 2-FFU decisions than 1-FFU decisions in both conditions. The mean conditional probability of 2-FFU decisions for the low frequency scenarios is 0.49 for the 80% 2-FFU, 20% 1-FFU while for the 80% 2.FFU, 20% 1-FFU condition, it is 0,74. This suggests that there is more to the decision making than just relying on frequency for the low frequency scenarios. Specifically, the results suggest that when in doubt, the participants send two rather than one FFU, as discussed above.

Discussion

The results of the present experiment agree with those of earlier experiments in showing that participants do not compensate very well for the time constants, even after considerable practice (compared to the amounts of practice we can expect in the real world). If the participants had learned this, we would not have observed the marked difference between the high and the low frequency scenarios: the same time constants apply to both.

However, the participants obviously discriminate between the high and low frequency scenarios in both learning conditions, and this implies some learning of the time constants, even though it is far from perfect. Indeed, the participants learned to cope reasonably well with the high frequency scenarios. For these scenarios, they seem to match the relative frequency of their 1-FFU and 2-FFU decisions to the relative frequency of 1-FFU and 2-FFU scenarios. For the low frequency scenarios, there is no evidence of frequency gambling, even though these scenarios are the best candidates for this, since there was less practice for these kinds of scenarios and thus less opportunity to learn. Instead, the decision making appears to be random at a first look. Closer study suggests, however, that the decisions for these scenarios rely on a principle of caution, in that the participants respond with a 2-FFU decision rather than 1-FFU decision. This is true, to some extent also for the high frequency scenarios. Even for these scenarios, there is a high frequency of 2-FFU decisions (0.92 in the condition where there had been 80% scenarios requiring 2-FFUs in the learning stage, and 0.16 in the condition where there had been 20% 1-FFU in the learning stage. For the latter condition, there is also a much higher probability of the low frequency decision (0.74 2-FFU decisions for the scenarios requiring 2 FFUs than for the condition where there had been 80% 2-FFU scenarios, where the probability of the low frequency 1-FFU decision for the low frequency 1-FFU scenarios was 0.51). As noted above, a 2-FFU decision is, of course, safer, because whereas the fire in a 1-FFU scenario can always be extinguished with two FFUs, that in a 2-FFU scenario cannot be extinguished with one FFU. This may be seen as example of the heuristic we had observed in earlier studies (Brehmer, 1989; 1995), i.e., that of rapid or massive responding, i.e., sending out as many FFUs as possible to a fire as rapidly as possible, which is, of course, exactly what responding on the basis of a "better safe than sorry" heuristic would lead to.

This study failed to support the hypothesis that the participants would use frequency gambling, but this does not mean that they would never use this heuristic. One reason may be that the participants actually did learn to distinguish between the two kinds of scenarios, i.e., they learn to compensate for the time constants to some extent, especially for the high frequency scenarios. In addition, there was an obvious (to the participants) alternative to the frequency heuristic: the "better safe than sorry" heuristic. Together, the albeit imperfect learning of the time constants and the "better safe than sorry" heuristic, the participants actually succeeded reasonably well and actually managed to extinguish most fires, albeit at a higher cost than necessary. When there is no such alternative, or when the time constants are more difficult to learn than they appear to have been in the present scenarios, frequency may well be an alternative. It is an important task for future studies to find the conditions under which a frequency heuristic may operate, as well as the conditions under which participants resort to the other "cognitive primitive" described by Reason (1990), i.e., "similarity matching".

Although we did not find the frequency gambling heuristic we had expected, the results of this study, as well as those of Brehmer and Nählinder (2004) nevertheless suggest that the hypothesis that people cope with time constants by means of heuristics may have considerable explanatory value. However, the heuristic in operation here is a "better safe than sorry" heuristic". As the heuristic found by Brehmer and Nählinder (2004), it works quite well in that gets the job done, although not at minimum cost. But then optimal responding may be of greater concern to decision theorists than it is to people in general, and may provide a useful guide to understanding what people actually do in dynamic tasks.

References

Booth Sweeney, L. & Sterman, J.D. (2000). Bathtub dynamics: Initial results of a systems thinking inventory. *System Dynamics Review*, *16*, 249-286.

Brehmer, B. (1989). Feedback delays and control in complex dynamic systems. In P. M. Milling & E. O. K. Zahn (Eds.), *Computer-based management of complex systems*. (pp. 189-196). Berlin: Springer-Verlag.

Brehmer, B. (1995). Feedback delays in complex dynamic decision tasks. In P. Frensch & J. Funke (Eds.), *Complex problem solving: The European perspective.* (pp.103-130). Hillsdale, NJ: Lawrence Earlbaum.

Brehmer, B. (2005). Micro-worlds and the circular relation between people and their environment. *Theoretical Issues in Ergonomics Science, 6*, 73-94.

Brehmer, B. & Allard, R. (1991). Dynamic decision making: effects of complexity and feedback delays. In J. Rasmussen, B. Brehmer & J. Leplat (Eds.), *Distributed decision making*. Chichester: Wiley.

Brehmer, B.& Dörner, D. (1993).Experiments with computer-simulated microworlds: Escaping both the narrow straits of the laboratory and the deep blue sea of the field study. Computers in Human Behavior, 9, 171-184.

Brehmer, B. & Nählinder, S.(2004). Achieving what cannot be done: Coping with the time constants in dynamic decision tasks. Proceedings of the 22nd International Conference of the System Dynamics Society, Oxford,U.K., July 25-29.

Hasher, I. & Zacks, R. T. (1984). Automatic processing of fundamental information: The case of frequency of occurrence. *American Psychologist, 39*, 1372-1388.

Hogarth, R. M. (2001). *Educating intuition*. Chicago: University of Chicago Press.

Jensen, E. & Brehmer, B. (2003). Understanding and control of a simple dynamic system. *System Dynamics Review*, *19*, 119-137.

Lindahl, M.-B. (1974). *Conceptual learning effects of the intradimensional variation of the instances*. Uppsala: Uppsala University: Ph.D. Dissertation.

Løvborg, L. & Brehmer, B. (1991). *NEWFIRE – A flexible system for running simulated fire fighting experiments*. Risø National Laboratory, Denmark, Risø-M-2953.

Løvborg, L. & Brehmer, B. (1992). *New facilities in NEWFIRE*. Working papers in cognitive science and HCI (Ed. O. Ravnholt). Roskilde, Denmark: Centre for Cognitive Science, University of Roskilde WPCS-92-8.

Mayr, O. (1970). *The origins of feedback control*. Cambridge, Mass.: MIT Press.

Moxnes, E. (2000). Not only the tragedy of the commons: misperceptions of feedback and policies for sustainable development. *System Dynamics Review, 16,* 325-348.

Reason, J. (1990). *Human error*. Cambridge, Mass: Cambridge University Press.

Sterman, J. D. (2000). *Business dynamics. Systems thinking and modelling for a complex world*. Boston, Mass.: Irwin McGraw-Hill.