Table of Contents

Go Back

Achieving what cannot be done: Coping with the time constants in a dynamic decision task

Berndt Brehmer
Department of War Studies
Swedish National Defence College
P.O. Box 27805
SE-115 93 Stockholm, Sweden
Tel. +46 8 788 8495
berndt.brehmer@fhs.mil.se

Staffan Nählinder
Command and Control Systems
Swedish Defence Research Agency
P.O. Box 1165
SE-581 11 Linköping, Sweden
Tel. +46 13 37 8163
staffan.nahlinder@foi.se

Abstract

This study examines how people handle time constants in dynamic tasks, using a microworld called NEWFIRE, which simulates forest fire fighting. The participants did not discriminate between fires requiring different number of units if they were not allowed to move any units before the fire started. If they were allowed to do so, performance improved. This suggests that the participants do not learn the time constants, and that they used the heuristic to preposition the units to avoid having to do so. Using such heuristics may well be how people handle dynamic tasks also in other circumstances. More effort should therefore be put into studying what people actually do in dynamic tasks, not only into whether or not they perform optimally.

KEY WORDS: DYNAMIC DECISION MAKING, TIME CONSTANTS, HEURISTICS, MICROWORLDS

There is mounting evidence that people encounter problems when trying to control dynamic systems, even when these systems are simple (e.g., Booth-Sweeney & Sterman, 2000; Jensen & Brehmer, 2003). Yet, the world is full of such systems so people must have found ways of coping with them, even if what they do may deviate from what would be required on the basis of normative theory, such as control theory or system dynamics. This would hardly be surprising. Control theory and system dynamics are recent inventions, and hardly part of people's natural intuitions. In another class of tasks that also requires the use of recent normative theory for optimal performance, i.e. decision theory, people do not make decisions according to normative theory either, but use different forms of heuristics (see Hastie & Dawes, 2002, for a recent overview of this research). Perhaps we would find people using various heuristics also in dynamic tasks, if we looked for them?

This paper presents two experiments concerned with this problem. Specifically, the experiments concern how people cope with one of the three forms of feedback delays in a dynamic task: the *time constants*.

The task studied is that of forest fire fighting, using NEWFIRE, a microworld designed by Løvborg and Brehmer (1991), which has found widespread use in studies of dynamic decision making. In experiments with this microworld, the participants are asked to assume the role of a fire chief located in his command post. He or she receives information about fires from a spotter plane, and on the basis of this information, he or she sends out orders where to go to the fire fighting units (FFU). These units then report back about their activities and locations, and so it goes until the fire or fires have been extinguished. The concept is shown in Figure 1.

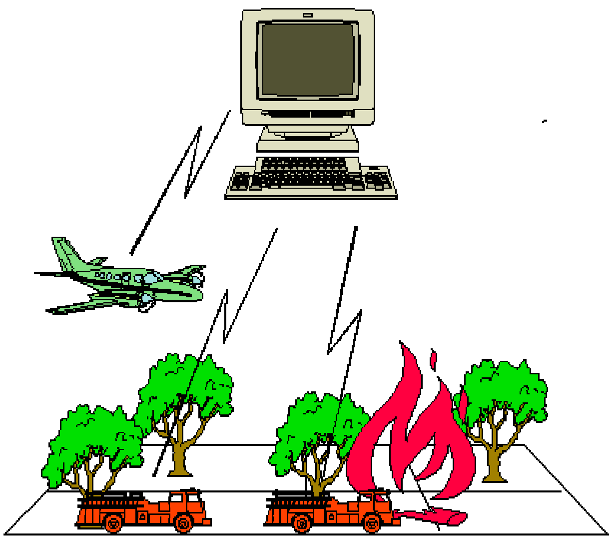

Figure 1. The NEWFIRE concept.

This task has all the characteristics of a dynamic decision task as defined by Brehmer and Allard (1991):

- It requires a series of decisions
- These decisions are not independent
- The state of the task changes, both autonomously, e.g., due to changes in the direction and strength of the prevailing wind, and as a consequence of the decision maker's actions, i.e., how he or she uses the FFUs
- Decisions must be made in real time

To cope with this task, the participant must be able to understand how fire spreads and how the FFUs can be used to fight the fire. Specifically, the task requires the participant to place the FFUs on the fire, or in the fire's way, depending on the extent to which the fire has spread and the distance that the FFUs would have to travel to reach the fire. When the fire and the FFU meet, the extinction process starts automatically.

As in all dynamic tasks, the participant must understand and cope with the various forms of feedback delays in the task. One form of delay that can never be avoided is the time constant, i.e., the fact that the FFUs require time to reach the fire, and that the fire extinction process requires time. The problem that the participant faces here is that while the FFUs are on their way, the fire spreads even further. As a consequence, the participant must send the appropriate number of fire fighting units to the fire. If there is a significant distance to that fire, he or she cannot send only the number that seems to be required at the point in time when the command is issued, for when the units reach the fire, it will have spread, and more units will be needed than were required when the units were sent out. Of course, the converse may also be true. If the fire is very far away, it may well have extinguished itself before the FFUs reach it.

In earlier studies, NEWFIRE has been used to investigate how people cope with the time constants, as well as with other forms of delays such as dead time and delayed reports from the FFUs (see Brehmer, 1995). The results have shown that while the participants are able to cope with the time constants, quickly learning to respond rapidly and massively to a fire, they have considerable problems with dead time and delayed reports (Brehmer, 1989). A possible explanation for this finding is that while the time constants can be seen to happen, dead time and delayed reports require inferences from the participants. However, even though the results show that people are able to cope with the time constants by means of sending out many FFUs to a fire so as to be sure to have enough FFUs on site, we do not know whether this reflects a more or less well calibrated strategy for handling the time constants, or whether it just represents a very general rule of responding rapidly and massively. To investigate this problem is the purpose of the present study.

The first question for the study concerned the precision with which the participants respond to a fire. Is this a well-calibrated response, tuned to the actual need of FFUs when they arrive at the fire, or is it just a very general response where the participant sends what he or she has as quickly as possible? To answer this question, we compared the response to fires requiring different numbers of FFUs in Experiment I. The results were surprising, at least when we first received them, albeit, perhaps not in hindsight, and led us to investigate the heuristics our participants used in the second experiment.

Experiment I

The purpose of Experiment I was to assess the precision in the participants' response to a fire by comparing the number of FFUs they sent out to a fire requiring only one unit and a fire requiring four units in NEWFIRE.

Participants

Ten undergraduate students were paid to take part in the experiment.

The simulation: NEWFIRE

NEWFIRE (Løvborg & Brehmer, 1991) is a general program for studying dynamic decision making in spatio-temporal decision problems such as forest fires. It is written in SMALLTALK and runs on a standard PC.

The general concept has been illustrated above (Fig. 1). Figure 2 shows what the participant sees. The figure shows that the participants face a computer generated map with 18 x 18 (=324) squares. This represents forest. In these experiments, the forest was homogenous, as indicated by all squares having the same shade of green. (It is possible to vary the terrain in the program.) In the forest, there is a base where the FFUs are located. To the left, there is a

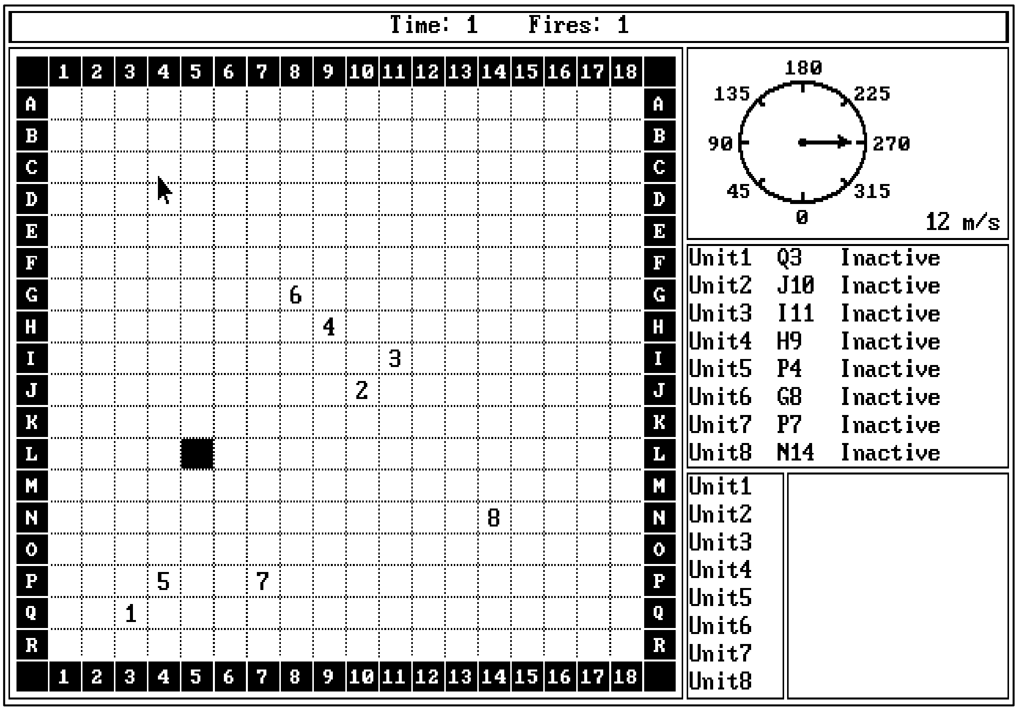

Figure 2. The NEWFIRE interface. Forest is represented by green squares, fires by the green aquares turning red (e.g., L5), FFUs by blue squares and the base, if there is one, by a brown square.

message box which gives messages from the FFUs, a box displaying the commands that have been given to the various FFUs and a weather report which gives the strength and direction of the wind.

In the middle of the map there was a base (a brown square marked with B) where eight FFUs were located at the start of a trial. After some period of time (in these experiments 40 sec. in Experiment I and 20 sec. in Experiment II), a fire starts, and this is indicated by a square on the map turning red. As the fire spreads, more and more squares turn red, and as fire is extinguished, the squares turn black. The participant sends out a FFU by first clicking on it and then on the location to which he or she wishes it to go. A message appears in the message window to the left saying that the unit is on the move. As the unit reaches its destination, it first mobilises (which takes one time unit) (and it sends the message “Mobilizing”) and then starts to fight fire (and it sends the message “Fighting fire”), if there is a fire in its location. After the fire has been extinguished, the FFUs demobilize for one time unit (sending the message “Demobilizing”). The time required for extinguishing a fire is a variable that can be set by the experimenter, as is the speed with which the FFUs travel to the fire. The participant can send as many units as he or she wished to a fire (out of a total of eight units) but putting more than one unit in the same location does not result in faster extinction. The number of fires is also a variable in the program. In the experiments reported here, there are two kinds of scenarios: one-fire scenarios and two-fire scenarios. A trial lasts until the fire has been extinguished or until all of the forest has been lost, or until the base has been lost. The fire spreads in the direction of the prevailing wind and its shape and the speed with which it spreads depends on the strength of the wind according to a fire model based on results from actual studies of forest fires.

The FFUs moved at the rate of 1 square/time unit.

NEWFIRE is a clock driven simulation. The picture is updated once every twenty seconds. (in Experiment I, in Experiment II the screen was updated every 10 sec.) An update interval is called a “time unit”.

Design

The design was a within-subjects design with distance between fire and the base as the factor. The factor had two levels: *close fires*, which required only one FFU, i.e., they could be extinguished simply by sending one unit directly to the location of the fire, and *distant fires,* which required four units, i.e., the participant had to predict which four squares would be on fire when the FFUs reached them, and position the FFUs accordingly. There were eight fires of each kind. In addition, there were 12 “filler items”, as described below. They were introduced to control the participants’ expectations. In all, the participants extinguished 28 different fires.

The time constant in this task is a function of the speed at which the FFUs move, the size of the fire, and the distance between the FFUs and the fire. The latter aspects vary as the scenario develops, and so do, of course, the requisite numbers of FFUs needed for a given fire. The definition of the independent variable in terms of close and distant fires holds only for the moment in time when the fire starts, given that the FFUs have not been moved from their initial positions. As we shall see, there is an important an important consideration in these experiments.

Fire scenarios. In all scenarios, the base was located in one of the middle four squares in the map, I9, I10, J9 or J10, se Figure 1. The FFUs were placed in the surrounding eight squares. All test scenarios involved a fire starting on the diagonal, with close fires starting at a distance 1 or 2 squares from the closest FFU and distant fires starting at a distance of 5 or 6 squares from the closest FFU (if the units were in the starting positions around the base). This yielded eight different starting positions for close fires and eight different starting positions for distant fires. In additions to these 16 test scenarios, there were 12 filler scenarios. Seven of these had two fires rather than one and five had a single fire starting in a position between the close and the distant fires. These scenarios were included to control for the participants’ expectations and make it impossible for them to position their FFUs in appropriate locations in advance. The two-fire scenarios also served the function to make the participants conserve their resources, and not send all their FFUs to the first fire they saw.

The first fire ignited 40 seconds (2 time units) after the start of the scenario in all scenarios. In the two-fire scenarios, the second fire ignited 70-100 seconds after the ignition of the first fire.

Instructions. The participants received standard instructions (see Løvborg & Brehmer, 1991, for details). They were told that their task was to act as fire chiefs and that they would receive reports about fire on the screen from a spotter plane that reported the position of a fire accurately and without delay. The screen and the mechanics of giving commands and the way in which the FFUs would respond were explained. They were further told that their task was to preserve the base and as much forest as possible. Thus, they had two goals: To protect the base and to extinguish the fire(s) as quickly as possible. These goals were sometimes in conflict, but the participants’ handling of these conflicts was not examined in this experiment.

Procedure. After receiving the instructions, the participants were given an opportunity to practice, using the mouse, and to familiarize themselves with the experimental setting for about 20 minutes. They then went through the 28 scenarios. They experiment required about 1½ hr. to complete.

Results

To assess possible learning effects, the scenarios were divided into four blocks, with 7 (in the analyses involving all 28 fires) or 4 (in the analyses involving only the test scenarios) scenarios in each block.

Number of units sent to different fires. When the number of commands issued after the start of a fire was examined for the two kinds of test scenarios, the results showed that the participants quickly learned to discriminate between scenarios requiring 1 unit and scenarios requiring 4 units. There was no significant change over blocks, but the number of commands issued for the distant fires exceeded that for the close fires (8.35 vs. 6.36, F 1/143 = 16.23, $p < .0001$). Remembering that two commands are required to reposition a FFU, this means that the participants used slightly more than the required number of 4 (The mean number was 4.2) for the distant fires and considerably more than the required number of 1 (The mean number of units sent out was 3.2) for the close fires.

Overall performance. The primary dependent variable to assess performance in NEWFIRE is the number of cells lost to fire. An analysis of all 28 scenarios showed that there was a learning effect (F 3/238 = 4.52, $p < .004$). Performance was worse in the first block than in the remaining three blocks, which did not differ significantly according to a Scheffe post hoc analysis. Only two of the 280 (10 participants x 28 scenarios), or 0.7% of the scenarios ended with a loss of base. These two scenarios were excluded from the analyses.

Next, the results for close and distant fires were examined. These results could not be subjected to an analysis of variance because of extreme differences in within-cells variances between the two conditions. However, the results, are clear enough also without statistical analysis. Thus, for close fires, the number of cells lost to fire was 1 over blocks and for distant fires, it was 5.7 with no systematic trend over blocks for either kind of fire.

If the participants had left the FFUs in their initial positions and then issued the correct commands to their FFUs after the fire had started, we would have expected that they would have lost 1 square to fire for the close fires and 4 squares for the distant fires. The results showed that although the participants lost the expected number of one square for the close fires, they lost an average of 5.7 squares for the distant fires, i.e., more than the expected number.

The participants repositioned their FFUs in the interval between the scenario started and the time when the fire started as shown by the fact that they issued an average of 15.03 commands during this interval. Since it takes two commands to move a FFU, this means that they managed to move almost all FFUs during the interval. Despite this, their performance was worse than expected for the distant fires. This shows that they were not very efficient in repositioning their units, or that they were unable to translate the new positions into effective fire fighting commands. The fact that the participants send out four FFUs to the distant fires required, but nevertheless lose more than the expected number of cells for these fires supports the second of these hypotheses.

Discussion

The purpose of this experiment was to investigate whether the participants would learn the time constants for the firefighting task. Evidence for such learning would be that the participants would have used the appropriate number of FFUs for different fires. The results indicated that they were able to distinguish between the two kinds of fire at last to some extent, and that the difference in response for the two fires was in the appropriate direction. Their response was not very precise, however, with close to the optimal number for the distant fires and considerable overreaction to the close fires to which 3 units, rather than 1, were sent out. As a consequence, the participants actually reached the optimal number of one square lost to fire for the close fires although at considerable cost. For the distant fires, the participants performed worse than expected, despite their repositioning of the FFUs in the interval between the scenario started and the fire ignited. They sent out slightly more than four units that would have been required had they left the units at the base, rather than repositioning them. This suggests that the repositioning of the units was not as efficient as it could have been, or that the participants were notable to translate the new positions into effective fire fighting commands. Both close and distant fires start on one of the diagonals, the distant fires starting farther out than the close ones. (If the participants sent one FFU out along each dianoal (NNW, NNE, SSW and SSE) as far as possible before the fire ignited, they would have been able to reach both kinds of fires with that FFU.) As noted above, the fact that they loose more fire than expected despite sending out the appropriate number of units supports the second of these hypotheses. This suggests that the participants do not only have problems with the temporal aspect of the task (i.e., finding the appropriate number of units to send out at the correct point in time) but also with the spatial aspects, that is, to which locations that they should send the units. This is in agreement with other results obtained with the fire fighting simulation (Brehmer, 1998), which show that participants in these kinds of experiments have problems predicting how the fire will spread.

These results, then, give little evidence of any precise learning of the time constants of the task. The attempts to move the units into a better position before the fire may actually be seen as an attempt to *avoid* learning these time constants, and to reduce all fires to fires of the same kind, that is, fires to which the participants could send out the maximum number of FFUs available to them as quickly as possible. This is, of course, an effective strategy in that achieves the goal of extinguishing the fire. But it is not efficient in terms of cost. Specifically, it is too costly for the fires requiring only one FFU, and in view of the fact that the participants have moved the units before the fire has started, it is also costly for the fires requiring 4 units. For the latter fires, the participants do not seem to be able to profit from their repositioning of the FFUs and reduce the number of FFUs sent to these fires. One possible explanation is, of course, that the high number of FFUs sent to these fires reflects the same tendency to send many FFUs to a fire as the high number of FFUs sent to the fires that would have required only one FFU.

The fact, that the participants repositioned the FFUs in the interval between the start of the scenario and the ignition of the fire, makes it hard to assess the extent to which they learned the precise time constants of the task. In Experiment II, therefore, we examine how the participants handle the time constants under conditions when they are not allowed to reposition their FFUs in the interval between the start of the scenario and the ignition of the fire.

Experiment II

Participants

Forty undergraduate students were paid to participate in the experiment.

Simulation

This experiment used NEWFIRE, the same simulation as that used in Experiment I. However, compared to the version used in Experiment I, that used in Experiment II also showed the accumulated cost in a trial at the top of the map display, Cost was displayed in terms of money units, and the cost of the various activities were as follows: A unit being inactive cost nothing, when it moved that cost 2 money units per time unit, mobilizing, fighting fire, demobilizing cost 4 money units/time unit and watching (= when a unit had been sent out to a location where there was no fire (yet)) cost 1 money unit/time unit. The cost factors were the same in the two conditions of the experiment (see below).

The update rate was set to 10 seconds rather than 20 seconds as in the first experiment simply so that that the participants would not have to be inactive for so long before the first fire. While considerably faster than the update rate in the first experiment, experience with NEWFIRE shows that a 10 second update rate in no way makes excessive demands on the participants, and it is not likely that the difference in update rates is the explanation for the differences in results between Experiments I and II.

Design

There were two experimental conditions. In the first, the *Variant* condition, the strength and direction of the wind was determined randomly for each scenario. In addition, the initial positions of the FFUs varied randomly. Four of the scenarios in this condition involved two fires, the remaining 16 one fire. In the two-fire scenarios, the second fire ignited in some location of the forest 70-100 seconds after the first had ignited. The second condition was called the *Constant* condition, and in this condition, the direction and strength of the wind was the same in all scenarios. The FFUs were always located in the eight squares surrounding the base as they were in Experiment I. Ten of the 20 scenarios were two-fire scenarios of the same kind as those in the Variant condition.

In each condition, the participants received a total of 20 scenarios. Half of the single fire scenarios in each condition, i.e., 8 out of 16 in the Variant condition and 5 out of 10 in the Constant condition were one-FFU test scenarios of the same kind as those used in Experiment I while the remaining were four-FFU-scenarios, as in Experiment I.

Procedure

The participants received the standard instructions for NEWFIRE used in Experiment I, but the instructions also included the information about the cost of different activities described above, and the instructions emphasized that the participants' goals were to put out the fire as quickly and with as little cost as possible. The need to protect the base was emphasized, and the participants were informed that if the fire reached the base, the trial was over. The participants were also told that they were not allowed to move the FFUs until a fire had started. As in Experiment I, participants were allowed about 20 minutes to familiarize themselves with the simulation. After that they received the 20 trials in their condition in a random order unique for each participant. After 45 minutes those who wanted were allowed a break. The experiment took between 2 and 4hr to complete.

Results

To determine if there were learning effects, the data from the experiment were analyzed combining scenarios into four blocks of five scenarios each.

All trials in which the participants violated the rule forbidding them to issue commands before a fire had started were excluded from the analyses. This happened in four out of the total of 40 participants x 20 scenarios (= 800 scenarios) data set. These scenarios were treated as missing data.

The base was lost to fire in 84 of the total of 800 scenarios (11%). For trials in which the base was lost, the total loss of forest was set to 324 cells. The logic here was that when the base was lost, the participant could no longer control the FFUs and, consequently, all forest would be lost. As in Experiment I, there was significant improvement over blocks of scenarios ($F\ 3/117 = 10.77$, $p < .000005$). Performance was better in the Variant condition than in the Constant condition ($F\ 1/19 = 4.42$, $p < .005$), probably because there were more two-fire scenarios in that condition, but, more important there was no interaction between conditions and blocks suggesting that learning proceeded in the same way in both conditions.

The final level of performance in the fourth block was an average of 38 cells lost to fire, compared to an average of 2 cells lost to fire in Experiment I. Thus, both the results with respect to cells lost and bases lost show that performance in Experiment II was worse than that in Experiment I.

Close and distant fires

There was no difference between the number of FFUs sent out in one-FFU and four-FFU scenarios to the fire during the first time period after the ignition in either condition, and no difference between the conditions. That is, the participants did not discriminate between the two kinds of scenarios. They learned to respond rapidly and massively, increasing the number of units sent out to the fire during the first time period from 1.6 units in the first block to 2.5 units in the fourth block ($F\ 3/117 = 25.97$, $p < 000001$), with no main effect of conditions and no conditions by blocks interaction. The final value of 2.5 is what should be expected if the participants had sent 1 FFU to the one-FFU fires and 4 FFUs to the four FFU-fires.

Individual analyses

The participants were divided into “Good” and “Bad” participants according to their performance over the 20 scenarios in their condition, the “Good” participants were those loosing the least forest in each condition, and the “Bad” participants being those who lost most forest. An analysis of variance showed that the “Good” participants sent out significantly more FFUs to the fire in the first time unit than the bad participants (F 1/19 =

64.48, p < 0005). This was true both for close and distant fires, there was no participants by fires interaction. Thus, the “Good” participants had not learned the time constants better than the “Bad” participants.

Discussion

The results from Experiment II show that the participants did not learn to discriminate between the two fires. This suggests that they do not learn the time constants of the fire-fighting task under conditions when they are required to do so, but not allowed to rely on the heuristic of repositioning the FFUs before the fire breaks out. This, in turn, suggests that the adaptation to the time constants, seen in this and earlier experiments with NEWFIRE, expresses a general strategy of massive and rapid responding rather than an adaptation to the specific time constants of the task. It is interesting to note that this is true also on the individual level: successful participants differ from less successful participants only in that they employ the strategy of responding rapidly and massively to a greater extent than less successful participants. The successful participants did not learn to discriminate between the two finds of fire better than the less successful participants, so there is no evidence that they learned the precise time constants better.

Whether people are actually unable to adapt to the specific time constants, we do not know, of course. We can only note that they did not do so in this experiment. The fact that they use the heuristic of repositioning their FFUs when allowed to do so, suggests that they may be aware of the difficulties that they have with the precise time constants and seek alternatives. Specifically, the alternative chosen would be to adopt a strategy that makes it unnecessary to learn the time constants by repositioning the FFUs so that all fires become similar and require the same response. Although this strategy is costly in that it requires them to send out their FFUs from their rest state at the base, it nevertheless gets the job done much better than if they had not employed this strategy. Evidence for this is the much better overall performance in Experiment I than in Experiment II (0,7% base losses compared to 11% base losses, and an average of 2 cells lost to fire compared to 34 cells lost to fire in the last block, dramatic differences both), despite that there were more of the more difficult two-fire scenarios in Experiment I (43% vs. 35%). The strategy employed by the participants is, of course, neither new nor unknown; many city ordinances require fire stations to be located so that any fire in the city can be reached within a specified period of time.

The participants in NEWFIRE thus discover a way of achieving their goal, without performing the task in an optimal way (hence the somewhat extravagant title of this paper). Of course, we do not know if this is because they cannot learn the time constants, or because they are satisfied with what they can achieve without learning these constants. The main effect of ignoring these time constants is, after all, only that it makes the fighting of the fires more expensive, not that it prevents them from actually extinguishing the fire. That is, it may be that the participants decide to achieve one of the goals: extinguishing fire, at the expense of the other goal: that of extinguishing the fire(s) at the lowest possible cost. Whatever the explanation, the participants find a way of performing the task that does not require them to cope with the exact time constants of the task. It is clear that they understand the general implications of the time constants, viz., that they must respond rapidly and massively, and the better they follow this simple rule, the better their performance. That is, it seems that they employ a simple qualitative heuristic rule, and avoid coping with the quantitative details. This suggests that it may be the quantitative aspects of the task, rather than its general dynamic nature, that gives problems. This may well be exacerbated in this task by the fact that these quantitative aspects have to be *estimated* by the participants from what they experience in the task; they are not displayed in numerical form. Perhaps this is the case also for other dynamic tasks. That is, it may well be that participants fail in such tasks, not because they do not understand the dynamics, but because they cannot handle the quantitative details required to demonstrate their understanding. This suggests that it may be important to distinguish between these aspects of performance, and that it is necessary to obtain separate indicators for the understanding of the general nature of the task, and for the ability to handle the quantitative details.

It may also be, that people rely on heuristic rules, such as the prepositioning of the FFUs in NEWFIRE (and of fire stations in cities) shown here, to handle other dynamic tasks as well. To find these heuristics, we need to study how people actually behave in dynamic tasks, not only whether they live up to what normative systems require.

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. (1998). Effects of time pressure in fully-connected and hierarchical architectures of distributed decision making. I Y. Waern (Red.), *Co-operative process management.* (pp. 131-143). London: Taylor & Francis.

Brehmer, B. & Allard, R. (1991). Dynamic decision making: The effects of complexity and feedback delays. In J. Rasmussen, B. Brehmer & J. Leplat (Eds.), *Distributed decision making: Cognitive models of cooperative work.* (pp. 319-334) Chichester: Wiley.

Hastie, R. & Dawes, R. M. (2001). *Rational choice in an uncertain world*. Thousand Oaks, Calif.: Sage.

Jensen, E. & Brehmer, B. (2002). Understanding and control of a simple dynamic system *System Dynamics Review*, *19*, 119-137.

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

Author note

The authors are indebted to Leif Løvborg and Anders Winman for their help with Experiment I and to Eva Jensen for her comments on earlier versions of the paper. The research reported here was support by grants from the Swedish Social Research Council and the Swedish Armed Forces (“The ROLF project).

Back to the Top