Controlling Stable and Unstable Dynamic Decision Making Environments

Magda Osman (m.osman@qmul.ac.uk)
Centre for Experimental and Biological Psychology, Queen Mary University of London,
London, E1 4NS UK

Maarten Speekenbrink (m.speekenbrink@ucl.ac.uk)
Cognitive, Perceptual and Brain Sciences, University College London
London, WCIE 6BT, UK

Abstract

In the present study we ask: Are people sensitive to the stability of a dynamic environment under short exposure to it? To examine this we investigate people's cue manipulation and strategy application when instructed to learn to control an outcome in a dynamic system by intervening on three cues. The system was designed in such a ways that in the Stable condition participants controlled an outcome that fluctuated steadily overall 40 trials, and in the Unstable condition the outcome fluctuated erratically over 40 trials. In the present study we show that people tended to intervene more frequently on all three cues when the system was Unstable compared to the when the system was Stable. Overall, the evidence from this study supports the general prediction made from the Monitoring and Control framework (Osman, 2010a, 2010b). It claims that people are sensitive to the underlying stability of dynamic environments in which they are required to control the outcome, but are insensitive to autonomous characteristics of the system.

Keywords: Dynamic; Control; Prediction; Decision making

Introduction

Complex dynamic environments come in many flavors (Cohen, Freeman, & Wolf, 1996; Klein, 1997; Lipshitz, Klein, Orasanu, & Salas, 2001; Lipshitz & Strauss, 1997), such as economic (e.g. stock exchange), industrial (e.g., chemical waste disposal), critical safety (e.g., automated-pilot systems) and biological (e.g., eco-systems). These situations differ from each other for a host of reasons, but crucially they share two fundamental features: they are dynamic and they are autonomous. That is, the outcome (e.g., state of the environment) fluctuates over time, whether rapidly (e.g., a sudden down pour of rain in an otherwise sunny day) or relatively slowly (e.g., steady increase in temperature over the spring months). Additionally, in both cases, changes in the outcome can occur independently of direct interventions made by decision makers. Given the probabilistic properties of these environments, an action may not reliably produce the same outcome each time, which raises the question: What are the differences in learning behaviors when attempting to control a highly noisy environment as compared with attempting to control a less noisy one? The aim of this study is to address this question in detail by examining control-based behaviors using a laboratory simulated complex dynamic task environment.

Uncertain Dynamic Environments: Often when determining the outcome in complex dynamic environments a series of inter-related decisions are made (Brehmer, 1992). That is, a future decision builds on the outcome of a previous decision and so on in order to work towards a goal. For instance, if we decide to take a couple of aspirin when we have a headache, we know that there is a variable delay in taking effect, and that the intensity of headaches changes over time. If after some period the headache persists, we may decide to take more aspirin, but without being sure that it will take effect, and if so, when it will do so. In this case, our decision making requires a series of choices to act towards achieving a specific goal (alleviating the headache), but there is uncertainty attached to our choice of actions (when to take aspirin, and what dosage), because we cannot be sure we will reliably produce the desired effect. Typically, people are required to interact with an environment by deciding from various cues (e.g., Drug A, Drug B, Drug C) actions that are relevant (e.g., selecting Drug A at dosage X) to changing the outcome (e.g., reduce the spread of disease). To introduce complexity into the task environment, the cue-outcome associations are probabilistic, and the environment dynamic, which ensures that from trial to trial the effects on the outcome will change. Moreover, it encourages people to adapt their decision making in order to

Many have used complex dynamic control tasks (CDC) as a way to examine the effects of varying the specificity of the goal under which the individual is instructed to learn about a complex environment (Burns & Vollmeyer, 2002; Geddes & Stevenson, 1997; Miller, Lehman & Koedinger, 1999; Osman, 2008; Vollmeyer, Burns & Holyoak, 1996). In this way, it is possible to examine the best conditions under which to learn to control an uncertain dynamic environment. Much of the evidence suggests that people are able to control complex systems successfully after sufficient opportunity to explore the environment first. However, training them to learn to control the system to a specific criterion can impair their ability to successfully develop flexible knowledge of the system that they can transfer to a different goal structures (Burns & Vollmeyer, 2002; Osman, 2008; Osman, 2010a).

Present Study

CDC tasks come in many varieties (for review see, Osman, 2010a, 2010b), but crucially, they tend to fall into two categories of systems, namely those that are dynamic, by which we refer to Funke's (1993) definition "An endogenous variable [that] at time t has an effect of its own state at time t+1 independent of exogenous influences that might add to the effect", and those that are static; in which the state of the system between t and t+1 is only dependent on exogenous influences on the system. Studies using CDC tasks with actual dynamic systems have thus far not systematically examined the effects of varying the endogenous variables on control performance. In other words, there has been no direct comparison of the effects of instability – in which the fluctuations in state as a result of the influence of the endogenous variable are high, and stability – in which the fluctuations in state as a result of the influence of the endogenous variable are low, on generating specific outcomes reliably in uncertain dynamic environments.

Until now, there has only been one previous comparison of the effects on cue utilization when controlling a system to a specific criterion under conditions in which the system is either stable or unstable (Osman & Speekenbrink, 2011). In their study they examined the influence of instability on cue utilization in a complex dynamic control task. Their participants received extensive training (200 trials) to one of two types of environments (Unstable, Stable), from which they were required to learn to control the system. Osman and Speekenbrink (2011) reported that people behaved differently according to the stability of the environment. Here stability was manipulated according to the level of noise (probabilistic relationship between cues and outcomes) in the system. The critical difference between the Stable and Unstable groups concerned the frequency of cue interventions and the range of cue values that were chosen in order to bring the outcome value in line with the criterion. Those in the Stable environment made conservative changes to the cue value and tended to change one cue at a time, whereas in the Unstable environment people tended to intervene on all three cues across most of the training trials, while also making full use of the range of the cue values. Thus, the pattern of behavior suggests that both groups adapted their decision making to the dynamic properties of the environment based on exogenous changes to the system. However, they failed to detect the endogenous feature of the system. One of the cues (Null cue) did not have any impact on the outcome value and when manipulated the outcome would simply reflect the internal perturbation in the system. Neither group was sensitive to the fact that changes to the outcome when manipulating the Null cue reflected an autonomous change in the system.

Given the limited research on the effects of the stability of the CDC task environment on knowledge acquisition, the present study aims to further explore decision making behavior in detail by measuring control performance, cue utilization and strategy application in the same control system in which the context, structure and instructions were identical. The critical difference was that in one condition the system was stable and in the other condition the system was unstable. In so doing we aim to replicate and extend Osman and Speekenbrink's (2011) findings.

Thus, in our system (Unstable, Stable) there were three cues which could be manipulated. One had a positive effect on the outcome, one had a negative effect on the outcome, and the third was a null cue, which had no effect on the outcome. When the null cue was manipulated the observed changes to the outcome in the system simply reflected the perturbation inherent in the system which would either make the outcome fluctuate in an unstable way (i.e. Unstable system), or in a stable way (i.e. Stable system). Thus of critical interest would be whether participants would be sensitive to the null cue with little exposure to the task environment. That is, unlike Osman and Speekenbrink (2011), we are concerned with whether people establish the same pattern of behavior under limited exposure to the environment and extensive exposure to the environment. For this reason we present people with only 40 trials in order to learn to control a Stable or Unstable CDC task.

We base our predictions on Osman's (Osman, 2010a, 2010b; Osman & Speekenbrink, 2011) Monitoring and Control framework (hereafter MC framework). The MC framework proposes that dynamic and autonomous properties in a system contribute to it being subjectively experienced as uncertain. In uncertain dynamic control environments, when learning to control outcomes, people judge the success of their performance according to the discrepancy between the achieved and target outcome. Thus, under conditions in which there are endogenous as well as exogenous influences (i.e. direct changes to the outcome through cue manipulation) on the outcome, the relation between achieved and target outcome is difficult to interpret because of the source of change to the outcome is not only self initiated. There are two different types of influences on the outcome, those that are initiated by the decision maker, and those that are independent of the actions of the decision maker.

Osman (2010b) also proposes that the greater the flexibility and range of outcomes generated by the control system, the greater its instability, and the greater the demands it places on exerting control on the system. Therefore, by increasing the endogenous influences on the outcome (i.e., increase instability), it is expected that the cue-outcome associations will be harder to detect, and therefore cue-outcome knowledge will be less accurate and will in turn impair control performance. To complement this, studies of motor control propose that learning cue-outcome relations in dynamic tasks is based on the congruency between one's own actions and the observed effects on the system. Therefore to increase one's control in a system that appears to be unstable, people will increase their interventions on it in order to establish a closer association between their actions and the outcomes in the system.

Method

Participants

Thirty (12 Male) graduate and undergraduate students from University College London and University of Surrey volunteered to participate in the experiment for reimbursement of £6. The assignment of participants to the two groups was randomized with 15 participants in the Unstable Condition and 15 participants in the Stable Condition. Each participant was tested individually.

Materials and Design

The study included one between subject variable which compared the effects of the stability of the system that participants were required to control (Unstable, Stable). With the exception of stability, the interface, cover story, and goals of the system were identical for both conditions. The design of the environment involved four continuous variables, three of which were cues and one of which was the outcome (see Figure 1).

The cues varied in their relation to the outcome in the following ways: one was positively associated, the other negatively associated, and a third was unrelated to the outcome (null).

Structure of System: $y(t) = y(t\text{-}1) + b_1\, x_1(t) + b_2\, x_2(t) + e_t$

Note that the Positive cue = x_1, Effect of positive cue = b_1 = 0.65, Negative cue = x_2, Effect of negative cue = b_2 = -0.65. Random perturbation = e_t, (the random perturbation component, is normally distributed, with a mean of 0), Outcome value = y(t), Previous outcome value = y(t-1).

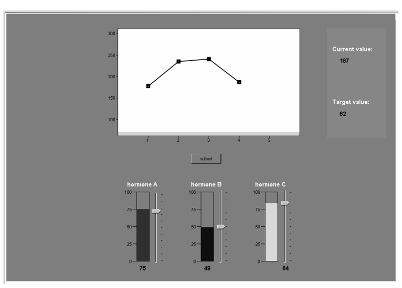

Figure 1: Screen Shot of Dynamic Control Task

To vary the stability of the system for the Random perturbation component we used a standard deviation of 16 (Stable condition) and to make it unstable we doubled the standard deviation to 32 (Unstable condition).

Successful control of the system: To learn to effectively control both stable and unstable versions of the system the endogenous influences on the outcome need to be distinguished from the exogenous influences on the outcome. To achieve this, the fewer and the more systematic the interventions made, the easier the cue-outcome associations are to learn for both versions of the system. Having accurate knowledge of the cue-outcome associations will in turn lead to successful control because as the outcome fluctuates, the subject will be able to identify the corresponding intervention on the cue necessary to bring it closer to target. E.g., if the outcome increases, then the subject needs to intervene on the negative cue on the next trial to bring the outcome value down towards the target value.

The visual layout of the screen, cover story, and the main instructions were identical for all four groups. Participants were presented with a summarized report of an article appearing in a medical journal.

> *It has recently been reported in The Lancet (##/##/##) "Patients under stress" (pp23-29) Special issue, that the Neurotransmitter (N) is released when patients are experiencing intense stress-related symptoms that slow down recovery. In addition, the research reported that three different naturally occurring hormones A, B, C also affect the release of the same neurotransmitter N. The basis of the research that you will be taking part in is to look at the relationship between the three different hormones A, B, C and their affects on the neurotransmitter N.*

Participants were informed that as part of a medical research team they would be conducting tests in which they would inject a patient with either one, or any combination of the three hormones, with the aim of maintaining a specific safe level of neurotransmitter release. The system was operated by varying the cue values (hormones A, B and C) that would affect the level of neurotransmitter release. The screen included the three labeled cues, and the outcome which was presented in two ways, as a value presented at the top right of the screen, and also in a small progress screen in which a short trial history (5 trials long) of outcome values was presented. The progress screen included a bar which highlighted the target value to which the outcome needed to be maintained. Thus, for each training trial participants received feedback concerning the current level of the neurotransmitter (i.e. achieved outcome) and the target value.

Procedure: The task included a total of 40 trials. Participants were presented with a computer display with three cues (hormones A, B, C) and the outcome (neurotransmitter). Each trial consisted of participants interacting with the system by changing cue values using a slider corresponding to each cue with a scale that ranged from 0-100. On the start trial, the cue values were set to '0' and the outcome value was 178. This means that for each trial people had to remember the interventions they made, but the effects of their interventions were presented graphically on screen in a small window (see Figure 1). Participants were instructed to maintain the outcome within a safe range (+/-10) of the target value, which was set at 62 throughout. After making their decisions, participants clicked a button labeled 'Submit' which made the cues inactive, and revealed on the progress screen the effects of their decisions on the outcome. The effects on the outcome value were cumulative from one trial to the next, and so while the cue values were returned to '0' on the next trial, the outcome value was retained from the previous trial. The cumulative effects on the outcome value were presented as a trial history on screen which contained the outcome values of the last five trials. When participants were ready to start the next trial, they clicked a button labeled 'Continue', after which the cues became active and were reset to '0'. After they completed the learning phase, participants then proceeded to the test phase.

Scoring: The training trials of the two different conditions were scored according to three different criteria (control performance, cue utilization, and strategy application). Control performance was based on error scores calculated as the absolute difference between the achieved and desired outcome value on each trial for each participant. Cue Utilization was scored in two ways: Cue manipulation and Parameter setting. For each participant, *Cue manipulation* was based on calculating the proportion of occasions that each of the three cues was manipulated. Second, *Parameter setting* was calculated based on the mean cue value that participants chose for each of the three cues. The *strategies* were based on calculating for participant the proportion of trials across blocks of training in which no cue was changed (No-intervention strategy), one cue was changed (One-cue-strategy), two cues were changed (Two-Cue-strategy), and all three cues were changed (All-Cue-strategy).

Results

The 40 control trials were divided in four blocks of 10 trials each and control error scores were averaged across each block for each participant. The following analyses were based on the mean error scores by block presented in Figure 2 for control error scores.

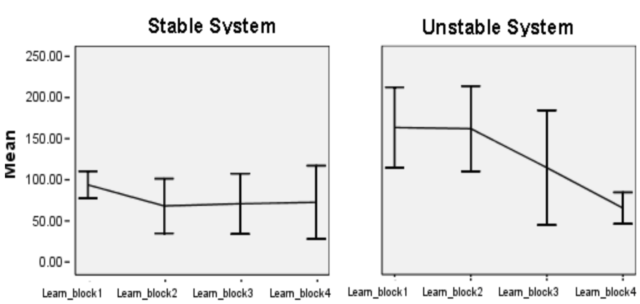

Figure 2: Mean SE (+/-) Control Performance by Condition

Control performance: The following analysis compared control performance by perturbation level (i.e. Stable vs. Unstable). A 4x2 ANOVA was conducted on control performance scores using Block (Learning Block 1, 2, 3, 4) as within subject factor, and Stability (Unstable, Stable) as between subject factor. As indicted in Figure 2, generally, for both conditions performance increased as familiarity with the system increased, confirmed by a main effect of Block, $F(3, 84) = 9.02$, $p < .0005$. There was a main effect of Stability on error scores $F(1, 37) = 10.42$, $p < .005$. Overall control performance was poorer in the Unstable condition compared with the Stable condition. A Block x Stability interaction, $F(3, 84) = 6.49$, $p < .0005$, was investigated further and was located in the first two blocks of the task, thereafter there was no difference in control performance between conditions (F<1).

Cue Manipulation: To examine the general patterns in the way people in Stable and Unstable conditions manipulated the three cues (positive, negative, null) we conducted a coarse analysis simply based on the proportion of manipulations made collapsed across blocks (See Figure 3).

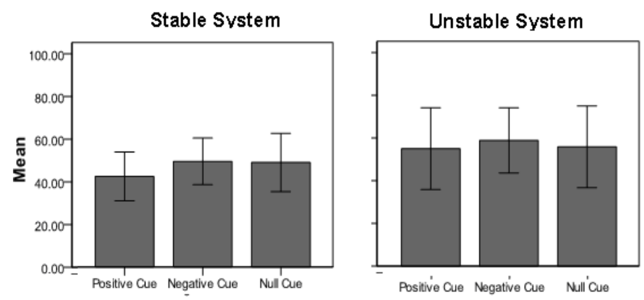

Figure 3: Mean SE (+/-) Cue Manipulation by Cue by Condition

A 3x2 ANOVA was conducted on the mean proportion of changes to cues across the all 40 trials. We use Cue (Positive, Negative, Null) as the within subject factor, and Stability (Unstable, Stable) as the between subject factor. There was no main effect of Cue, $F(2, 56) = 1.03$, $p = .37$, implying that the occasions on which the three different cues were intervened upon was equally distributed across the 40 trials. As indicated in Figure 3, there appeared to be an influence of stability on cue manipulation, which was confirmed, $F(1, 28) = 7.47$, $p < .01$. Thus, when the environment was Stable the three cues were manipulated less frequently than in the Unstable condition.

Parameter Setting: As a further method of examining people's sensitivity to the underlying stability of the system, we examined the range of values selected for each of the three cues. Figure 4 suggests that the overall values for the three cues appear to be lower in the Stable condition as compared with the Unstable condition.

Confirming this trend, a 3x2 ANOVA on mean values for the cues with Cue (Positive, Negative, Null) as within subject factor and Stability (Unstable, Stable) as the between subject factor, revealed a significant main effect of Stability, $F(1, 28) = 14.44$, $p < .0001$. No other effects were significant.

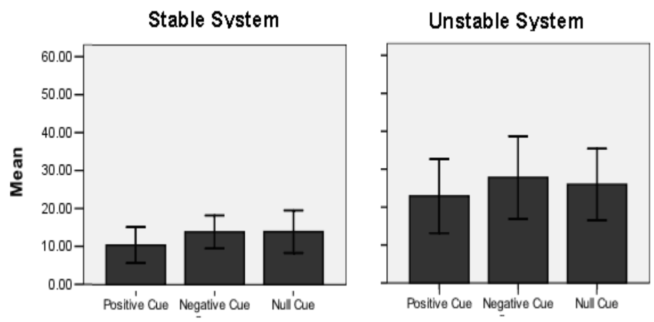

Figure 4: Mean SE (+/-) Cue Value By Cue By Condition

Strategy application*:* The following set of analyses examines patterns in the application of strategies in Stable and Unstable conditions. The first set of analyses is a coarse analysis of the general patterns across all 40 trials as shown in Figure 5.

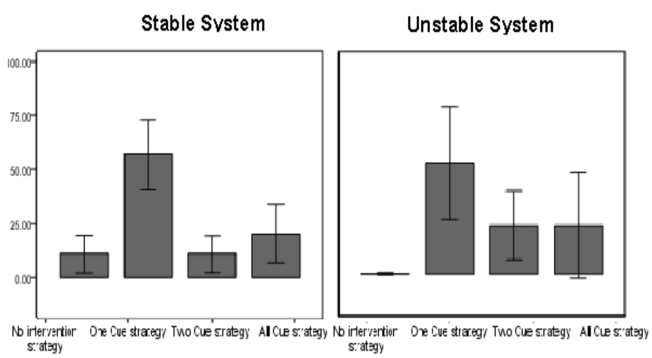

Figure 5: Mean SE (+/-) Proportion of the Four Strategies Employed by Condition The second set of analyses considers the profile of strategy development across blocks of control trials. To begin, a 4x2 ANOVA was conducted on the proportion of trials in which cues were varied using Strategy (No-Intervention-Strategy, One-Cue-Strategy, Two-Cue-Strategy, All-Cue-Strategy) as a within subject factor, and Stability (Unstable, Stable) and as the between subject factor. The analysis revealed a main effect of Strategy, $F(3,84) = 9.50$, $p < .0005$, suggesting that there were differences in the types of strategies favored overall, as indicated in Figure 5. There was also a main effect of Stability, $F(1, 30) = 18.85$, $p < .001$, and a significant Strategy x Stability interaction $F(3,84) =3.36$, $p < .05$. To locate the source of the Strategy x Stability interaction, univariate analyses revealed that the Stable condition used the No-Intervention-Strategy more often than the Unstable condition, $F(1, 28) =5.71$, $p < .05$. No other effects were significant.

General Discussion

The objective of this study was to examine in detail how people utilize information and develop strategies in a control system under conditions in which the outcome is either easy or difficult to control. This was achieved by keeping all other properties of the system the same but manipulating the endogenous properties of the system so that it was either experienced as Unstable or Stable. Overall, the evidence replicates and extends the findings reported by Osman and Speekenbrink (2011). The findings also supports the general prediction made from the MC framework (Osman, 2010a, 2010b), suggesting that people are sensitive to the stability of the environment, and that while people learnt to control an unstable as well as a stable condition, instability in the system is a source of uncertainty for people as indexed by the poorer control performance of the Unstable condition.

Differences between Unstable and Stable Conditions: The principle objective of this study was to examine if people were sensitive to both their effects on the outcome as well as the internal changes that could occur without their intervention (i.e. exogenous and endogenous influences on the outcome). We predicted that it would be harder to detect the endogenous effects on the outcome in the Unstable condition. The study found that even with such short exposure to the task environment, as compared with 200 trials that were used by Osman and Speekenbrink (2011) people increased their cue utilization as compared with the Stable condition. Second, the pattern of behavior for parameter setting of the three cues suggested that the values chosen for all three cues were consistently greater in the Unstable condition compared with the Stable condition. It appears that the fluctuations in the outcome value lead those in the Unstable condition to select more extreme cue values in an attempt to reduce the discrepancy between achieved outcome and target outcome from trial to trial. In turn this would also facilitate learning cue-outcome relations because by selecting extreme cue values that were easier to remember participants could have observed the effects of their interventions more clearly. Third, by intervening on the system more often there was less opportunity for people in the Unstable condition to uncover the dynamic and autonomous properties of the system, resulting in less accurate cue-outcome knowledge which impaired control ability. Fourth, the main difference between the types of strategies implemented between the two conditions was specific to the No-Intervention-Strategy, that is, those in the Stable condition employed this strategy more than the Unstable condition. In an earlier study (Osman and Speekenbrink, 2011), we found that the popular strategy used to control the Unstable condition involved varying all three cues, whereas in the Stable condition people experienced more trials in which the outcome of the system changed on it's own. People in the Stable condition made fewer but more systematic interventions, by varying one cue at a time. Crucially though, the main difference between that study and the current one is the number of trials that participants experienced (i.e. 40 vs. 200). Given the limited training it is likely that different patterns of behavior are revealed which then are likely to change with extensive exposure to the same environment. It appears that the default strategies that people employ are changing one cue at a time. Previous findings also suggest that varying one cue at a time is a more successful strategy to controlling a system as compared with varying all cues at the same time (Tschirgi, 1980; Vollmeyer, Burns, & Holyoak, 1996). However, this clearly changes when people have extended exposure to the environment, which implies that people adapt their strategies over time as they gain experience with a complex dynamic environment.

Similarities between Stable and Unstable conditions: The general pattern of cue utilization and strategy application differentiated people in the Unstable condition from the Stable condition. Consistent with Osman and Speekenbrink's (2011) findings regardless of the stability of the system, people utilized all three cues equally, and the range of values that were set for each cue was approximately the same.

Thus, in agreement with previous evidence, people in stable and unstable conditions are sensitive to exogenous influences on a dynamic environment but have difficulty detecting endogenous changes in the environment. As mentioned previously the null cue had no effect on the outcome, and simply reflected the random perturbation component of the system. However this would be hard to discover unless people reliably selected extreme values for this cues over a series of consecutive trials. In this way it would be easier to detect the dissociation between actions and effects.

Did both conditions fail to detect the endogenous property of the system for the same reasons?

It may be the case that while both groups failed to detect the null cue, the reasons for this are different. In the Stable condition people tended to manipulate one cue at a time, but were conservative with the cue values they chose which is possibly why they failed to detect the null cue. In contrast, even though the Unstable condition tended to pick extreme values for the cues, they also manipulated all the cues at once most of the time, which again would have made the null cue hard to detect. Thus, while stability influenced control performance, cue utilization, and strategy application, it did not affect ability to detect the null cue. In general, it may be the case that because people do not expect there to be erroneous cue information, they would operate a system assuming that each cue had an effect on the outcome. Moreover, they may also make the assumption that their actions will reliably generate changes in the system, because this is an obvious bias which is maintained in control task situations (Osman, 2010b).

Acknowledgments

Preparation for this research project was supported by Engineering and Physical Sciences Research Council, EPSRC grant - EP/F069421/1 (MO) and the ESRC RES-062-23-1511 (MS).

References

Brehmer, B. (1992). Dynamic decision making: Human control of complex systems. Acta Psychologica, 81, 211-241.

Burns, B. D., & Vollmeyer, R. (2002). Goal specificity effects on hypothesis testing in problem solving. Quarterly Journal of Experimental Psychology, 55, 241-261.

Cohen, M. S., Freeman, J. T., & Wolf, S. (1996). Meta-cognition in time stressed decision making: recognizing, critiquing and correcting. *Human Factors*, *38*, 206-219.

Funke, J. (1993). Microworlds based on linear equation systems: A new approach to complex problem solving and experimental results. In G. Strube & K.-F. Wender (Eds.), The cognitive psychology of knowledge (pp. 313–330). Amsterdam: Elsevier.

Geddes, B. W., & Stevenson, R. J. (1997). Explicit learning of a dynamic system with a non-salient pattern. *Quarterly Journal of Experimental Psychology, 50A*, 742-765.

Klein, G. (1997). Developing expertise in decision making. *Thinking and Reasoning*, *3*, 337-352.

Lipshitz, R., & Strauss, O. (1997). Coping with uncertainty: A naturalistic decision-making analysis. Organizational Behavior and Human Decision Processes, 69, 149–163.

Lipshitz, R., Klein, G., Orasanu, J., & Salas, E. (2001). Taking stock of naturalistic decision making. Journal of Behavioral Decision Making,14, 332–351.

Osman, M. (2008). Observation can be as effective as action in problem solving. Cognitive Science, 32, 162-183.

Osman, M. (2010a). Controlling Uncertainty: A Review of Human Behavior in Complex Dynamic Environments. Psychological Bulletin, 136, 65-86.

Osman, M. (2010b). Controlling Uncertainty: Learning and Decision Making in complex worlds. Wiley-Blackwell Publishers, Oxford.

Osman, M., & Speekenbrink (2011). Information sampling and strategy development in complex dynamic control environments. *Cognitive Systems Research, 355-364*.

Tschirgi, J.E. (1980). Sensible reasoning: A hypothesis about hypotheses. *Child Development, 51,* 1-10.

Vollmeyer, R., Burns, B. D., & Holyoak, K. J. (1996). The impact of goal specificity and systematicity of strategies on the acquisition of problem structure. Cognitive Science, 20, 75–100.