

Cue utilization and strategy application in stable and unstable dynamic environments

Action editor: Cleotilde Gonzalez

Magda Osman^{a,b,*}, Maarten Speekenbrink^b

^a *Biological and Experimental Psychology Group, School of Biological and Chemical Sciences, Queen Mary, University of London, Mile End Road, London E1 4NS, UK*

^b *Cognitive, Perceptual and Brain Sciences, Division of Psychology and Language Sciences, University College London, Gower Street, London WC1E 6BT, UK*

Received 24 March 2010; received in revised form 20 October 2010; accepted 20 October 2010
Available online 25 December 2010

Abstract

We took a novel Empirical approach to investigating dynamic decision making behavior by examining the profiles of individuals' information sampling behavior and strategy application under conditions in which the control task was unstable as well as stable. Participants were presented with a dynamic system which they interacted with by intervening on three cues in order to reach and maintain a specific outcome (goal). The system was manipulated so that in the Stable condition participants controlled an outcome that fluctuated steadily overall trials, and in the Unstable condition the outcome fluctuated erratically over trials. In general, unstable fluctuations in the outcome led people to sample all the cues most of the time, even those which had no effect on the outcome. In contrast, under Stable conditions people were more conservative in their cue sampling behavior. The implications of these findings are discussed with respect to previous work on dynamic decision making and the Monitoring and Control (Osman, 2010a, 2010b) framework.

© 2011 Elsevier B.V. All rights reserved.

1. Introduction

We often find ourselves in complex control situations in which we are required to develop plans of actions to generate desirable outcomes under conditions of uncertainty. Of interest to research in problem solving and decision making are the ways in which people identify information that is relevant for controlling outcomes. This is often observed in high stake situations when there are costs (time, money, effort, fatalities) attached to the outcomes people generate (Cohen, Freeman, & Wolf, 1996; Lipshitz, Klein, Orasanu, & Salas, 2001; Lipshitz & Strauss, 1997) such as economic (e.g. stock exchange), critical safety (e.g., automated-pilot

system) or medical (e.g., circulatory system) situations. These applied situations have many different characteristics that define them, but they share two important common features: they are dynamic and they are autonomous. That is, the outcome may fluctuate rapidly over short periods of time (e.g., a sudden down turn in the stock market) or remain relatively stable over long periods of time (e.g., growth of an economy). Additionally, in both cases, changes in the outcome can occur independently of direct interventions made by decision makers. So, in order to control outcomes in an environment which is autonomous and in which the outcomes fluctuate at different rates, do people interact with systems differently when they differ according to their stability? The aim of this study is to address this question. More specifically, this study examines whether peoples' information utilization and strategy application is sensitive to differences in the stability and autonomic features of the control task environment.

* Corresponding author at: Biological and Experimental Psychology Group, School of Biological and Chemical Sciences, Queen Mary, University of London, Mile End Road, London E1 4NS, UK.

E-mail address: m.osman@qmul.ac.uk (M. Osman).

1.1. Issues in the study of dynamic decision making

From what we understand of research on dynamic decision making (Brehmer, 1992; Osman, 2010b), typically, when we are required to control an outcome in a complex dynamic system a series of non-independent decisions and actions are made. That is to say, a future decision builds on the outcome of a previous decision and so on in order to work towards reaching a particular goal. However, control systems generate uncertainty because the system's states can change at different rates (stable, or unstable), and because the system itself may have properties that are autonomous – that is, changes in the outcome can occur independently of any action taken by the decision maker (Osman, 2010a, 2010b). Both these properties of the control system may obscure the cue-outcome relations making it difficult to discover information that is relevant to generating actions that can produce a specific outcome regularly (Lipshitz & Strauss, 1997; Osman, 2010a, 2010b). Given these features of uncertainty, it is important to understand how sensitive people are to situations in which their interactions with the control systems may have minimal effects on changing the outcome (e.g., misdiagnosing a disease and prescribing the wrong medication), or when the rate of change of the outcome is erratic (e.g., trying to control the sudden spread of wild fire across a forest). Until now, much of the focus of research on complex decision making is concerned with the impact of different methods of training on improving control ability, and less on studying the influence of system characteristics on knowledge and strategy application, which is the focus of the current study.

1.2. System characteristics of complex dynamic control (CDC) tasks

Typically, control tasks require that an individual interact with an environment (e.g., simulated medical context) by deciding from various cues (e.g., drug A, B, C) the actions that are relevant (e.g., selecting drug A at dosage X) for changing the outcome to a specific criterion (e.g., reduce the spread of disease). The control aspect of the task is that some property of the system entails that the outcome/outcomes will fluctuate from one time step to the next. To achieve this, there are many different types of cue-outcome associations which have been incorporated in the system. Some are linear but noisy (Burns & Vollmeyer, 2002; Osman, Wilkinson, Beigi, Parvez, & Jahanshahi, 2008), non-linear (Brehmer, 1992; Broadbent & Ashton, 1978), or probabilistic (Kerstholt, 1996), but this is by no means an exhaustive list. In addition, the structure of the system may also have decay functions so that the outcome from trial to trial will change (Hagmayer, Meder, Osman, Mangold, & Lagnado, 2010), or there are delays so that the effects of cue manipulations on the outcome are not observed in one time step, but over several time steps (Diehl & Serman, 1995). All these different types of cue-outcome associations still have the same general effect on

our decision making behavior; they encourage the individual to regularly intervene on the system by choosing actions that will maintain the outcome at a desired state.

As mentioned, the cue-outcome associations in control systems come in many varieties (for review see, Osman, 2010a, 2010b). Nevertheless, they can be loosely categorized as those that make the system dynamic, by which we refer to Funke's (1993) definition "An endogenous variable [that] at time t has an effect on its own state at time $t + 1$ independent of exogenous influences that might add to the effect", and those that make the system static; in which the state of the system between t and $t + 1$ is only dependent on exogenous influences on the system. In general, there is little to suggest that a truly dynamic control task (e.g., Gonzales, Lerch, & Lebiere, 2003; Hagmayer et al., 2010; Kerstholt, 1996) is harder to control as compared to a static control tasks (e.g., Berry & Broadbent, 1988; Burns & Vollmeyer, 2002; Dienes & Fahey, 1995; Osman et al., 2008). However, it is difficult to survey the literature and make cross comparisons concerning the general effects of static and dynamic systems on control ability; this is because the types of systems that are used vary in context, structure, and instructions. Moreover, it is not clear whether we can draw from existing literature to comment on whether there are systematic effects on control performance as a result of different types of dynamic properties of the control system.

For instance, take Kerstholt's (1996) and Hagmayer et al.'s (2010) demonstrations of poorer control performance in their studies, in both cases decrements in control performance may not have been simply the result of the fact that the systems in both studies are dynamic. In Kerstholt's (1996) dynamic system the outcome value (i.e. health of an athlete) could decline rapidly (unstable), or else fluctuate (i.e. within healthy bounds) (stable) over a course of trials. In addition, Kerstholt manipulated the financial cost of seeking information and financial penalties were attached to delays in implementing actions. As a result of these manipulations, Kerstholt found that stability was less of an influence on control ability as compared to whether there were costs attached to failing to reach and maintain target states in the system. Although they did not include penalties for poor control performance, Hagmayer et al.'s (2010) system was also dynamic, though it used a decay function to generate changes from trial to trial so that the outcome value would automatically reduce by half (i.e. uptake of a neurotransmitter in a rat brain). In their study they found that this encouraged participants to pay closer attention to their interventions because the changes to the outcome were dramatic. In one version of the system the decay was not well separated in time between decision and outcome, and as a result, control performance was poor. In another version control accuracy improved when the decay function was easier to observe in time after an action was taken. The evidence from Kerstholt's (1996) and Hagmayer et al.'s (2010) studies suggest that outcome changes in real time are easier to control

when they are unstable as compared with when the outcome changes in real time, but the changes are hard to identify. However, the findings are silent with respect to claims concerning the effect of real time changes on the outcome in terms of cue utilization and the strategies developed to control the system.

Similarly, the relationship between events in the system is hard to interpret in studies using control systems with feedback delays between actions and outcomes. These studies provide insights into the effects on decision making when the system is perceived to be autonomous. Kerstholt and Raaijmakers (1997) reviewed a variety of dynamic control systems that incorporated delays between the action taken by the individual and an associated outcome. They claimed that feedback delays are difficult for people to integrate into their understanding of cue-outcome associations in the system. Often when people make an intervention in the system they expect immediate feedback from it, when they do not receive it immediately, the delayed feedback tends to get forgotten or ignored, which in turn adds to the misrepresentation of the system. Diehl and Serman (1995) also showed that people's knowledge of the cue-outcome associations in the system can degrade when more delayed feedback is encountered, which in turn leads to poorer control ability.

One reason for the poor control performance observed as a result of delays in feedback is that the outcome of the system appears to be changing in ways that are unexpected and seemingly uncorrelated with the actions the decision maker is initiating. This results in the system being perceived as behaving autonomously, when in fact it is still under the direct control of the individual (Diehl & Serman, 1995; Gibson, 2007; Kerstholt & Raaijmakers, 1997). In a related domain, studies on perception action associations and motor control suggest that learning the relationship between our actions and the outcomes we observe depends on the degree to which there is congruence between the predicted and actual outcome of one's actions (Blakemore, Frith, & Wolpert, 2001; Blakemore, Wolpert, & Frith, 1998), and between predicted and observed outcomes (Osman et al., 2008). For this reason, when the congruency is disrupted, as is reported in studies in which there is delayed feedback, it is difficult to learn the relationship between one's own actions and the observed effects because it is harder to predict. Moreover, evidence from studies on motor control also suggest that if the action-outcome associations in a control system are obscured (by manipulating endogenous variables in a control task, or introducing autonomous properties into the system), then one might expect that people will have greater difficulty learning cue-outcome associations, and will show poorer ability in controlling an outcome (e.g., Blakemore et al., 2001; Osman et al., 2008). In addition, from this it follows that, to increase the congruency between their actions and the outcomes they observe in the system, it is likely that people will increase their interventions with the system. The approach taken in the present study allows these predictions to be examined.

1.3. Present study

Until now, there has been no direct 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. Moreover, there has been no dedicated investigation into people's sensitivity to different types of cue information in the system when the system is stable or unstable. Thus, given the limited research focus on examining the effects of the stability of the CDC task environment on knowledge acquisition, to explore these outstanding issues, the present study examined decision making behavior in detail by measuring control performance, cue utilization, strategy application and strategy development in the same control system in which the context, structure and instructions were identical. The critical difference was that the system was varied in such a way that in one condition the system was stable and in the other condition the system was unstable.

In each version (Unstable, Stable) of the system 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 such that they would intervene less on it regardless of the stability of the system.

We base our examination of the effects of stability on Osman's (2010a, 2010b) 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

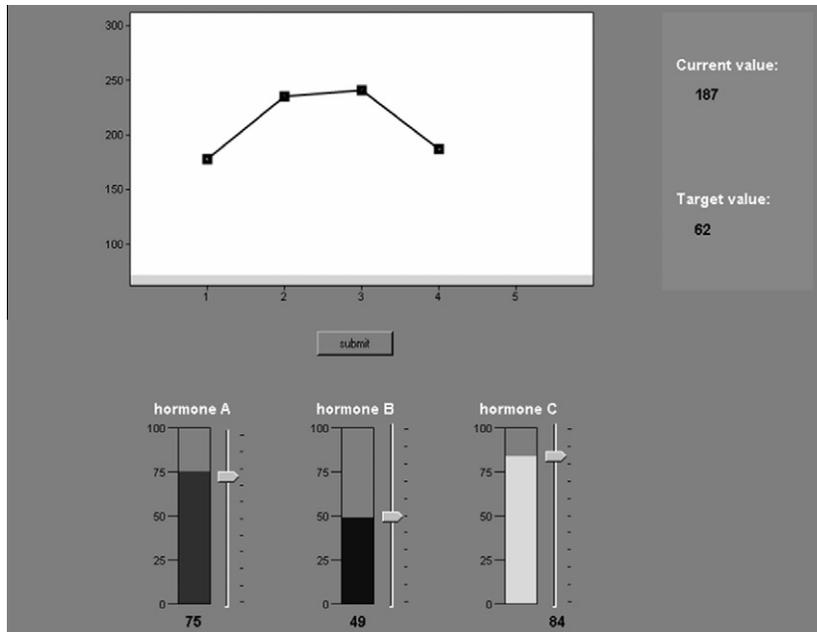


Fig. 1. Control task.

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.

2. Methods

2.1. Participants

Twenty-five graduate and undergraduate students from University of Surrey and University of London volunteered to participate in the experiment for reimbursement of £6. The assignment of participants to the two groups was randomized. There were two conditions (Unstable ($n = 13$), Stable ($n = 12$)). Participants were tested individually.

2.2. Design and materials

The study included one between subject variable that 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 outcome (see Fig. 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: $y(t) = y(t-1) + b_1x_1(t) + b_2x_2(t) + e_t$

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 = $y(t-1)$. To vary the stability of the system for the Random perturbation component we used a standard deviation of 16¹ (Stable condition) and to make it unstable we doubled the standard deviation to 32 (Unstable condition).

2.3. 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 experienced separately from the exogenous influences on the outcome. To achieve this, a series of trials in which no-interventions on the cues will reveal three critical aspects of the system: (1) that the outcome can change autonomously, (2) that it will fluctuate to some degree, though unbeknownst to the subject the fluctuation was experimentally manipulated as either stable or unstable,

¹ Osman and Speekenbrink (submitted for publication) includes two studies which varied the stability of the system for the purposes of investigating the effects on accuracy of prediction and control processes. Though a system with a random perturbation component of 16 SD was found to be difficult to control, this was demonstrated with only 40 control trials. However, by the last 10 trials participants were reaching asymptote. Therefore, in the present study we included 200 control trials and could comfortably assume that over this length of trials the system would be experienced as relatively stable and easy to learn as compared to a random perturbation component with 32 SD.

and (3) in which direction the outcome will change i.e. only in a positive direction, only in a negative direction, or both. To learn the endogenous influences on outcome, each cue needs to be sampled individually for a series of trials in order to learn the cue-outcome relations separately. In sum, the fewer interventions made and the more systematic the interventions made are, 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 positively 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” (pp. 23–29) 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 effect 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 in the top right of the screen, and also in a small progress screen in which a short trial history (five 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 their current level of the neurotransmitter (i.e. achieved outcome) and the target value.

2.4. Procedure

The task included a total of 200 trials in the extensive training condition. 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 to 100. On the start trial, the cue values were

set to ‘0’ and the outcome value was 178. 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.

2.5. Scoring

The training trials of the four different groups were scored according to three different criteria (control performance, control optimality, cue utilization, and strategy development). *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. *Control optimality* was based on how much participants’ cue manipulations deviated from the optimal cue settings. In the control task used here, for a given (previous) outcome value and goal, the optimal cue settings define a line in a two-dimensional plane. E.g., if the deviation between the previous outcome and goal is 50, then the optimal cue settings are all values for the positive cue x_1 and negative cue x_2 such that $50 = 0.65x_1 - 0.65x_2$, for instance a value of $x_1 = 77$ and $x_2 = 0$, or $x_1 = 78$ and $x_2 = 1$, $x_1 = 87$ and $x_2 = 10$, etc. Control optimality scores are computed as the (shortest) distance between a participant’s actual settings for these two cues and the line defining the optimal cue settings. For the environment used here, this distance is computed as in the following equation:

$$D(t) = |x_1(t) - x_2(t + [y(t - 1) - g]/0.65)|/\sqrt{2}$$

in which $g = 62$ denotes the target outcome value. On those trials in which the difference between previous outcome and target was larger than 65, or smaller than -65 , it was not possible to reach the target outcome in a single trial, thus the value $y(t - 1) - g$ was replaced by either 65 or -65 to compute the distance. Note that as the null cue has no effect on the outcome, it is not taken into account in the control optimality scores. 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 across all training trials that each of the three cues was manipulated. Second, *Parameter setting* was calculated based on the mean cue

value that participants chose across all training trials for each of the three cues. The strategies that were identified during training 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).

3. Results

The 200 control trials were divided in four blocks (block = 50 trials), control error scores and optimality scores were averaged across each block for each participant. The following analyses were based on the mean error scores by block presented in Fig. 2 for control error scores, and Fig. 3 for optimality scores.

3.1. Control performance

The following analysis compared control performance by perturbation level (i.e. Stable vs. Unstable). A 4 × 2 ANOVA was conducted on control performance scores using Block (learning block 1, 2, 3, 4) as within subject factor, and Stability (Unstable, Stable) as the between subject factors. As indicated in Fig. 2, generally, for both conditions control performance increased as participants became more familiar with the task. This was confirmed by a main effect of Block, $F(3, 75) = 16.57, p < .0005, \text{partial } \eta^2 = .41$. There was a main effect of Stability on error scores $F(1, 25) = 22.27, p < .001, \text{partial } \eta^2 = .49$. Overall control performance was poorer in the Unstable condition compared with the Stable condition.

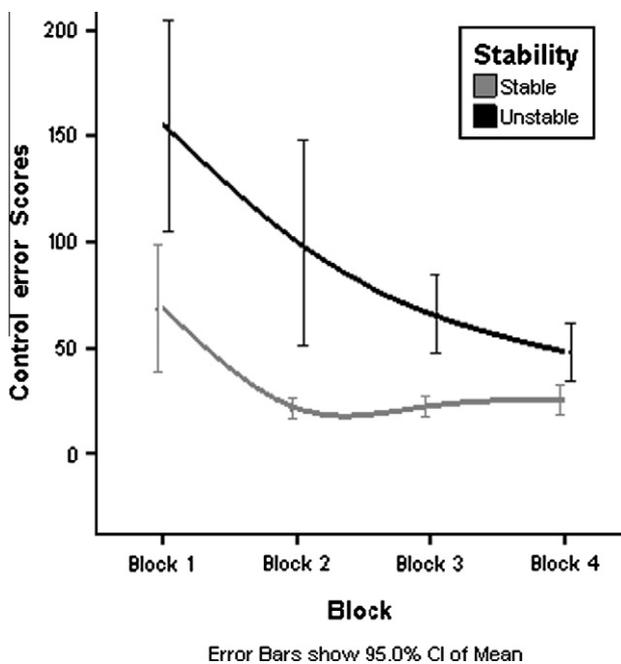


Fig. 2. Mean SE (±) control performance by condition.

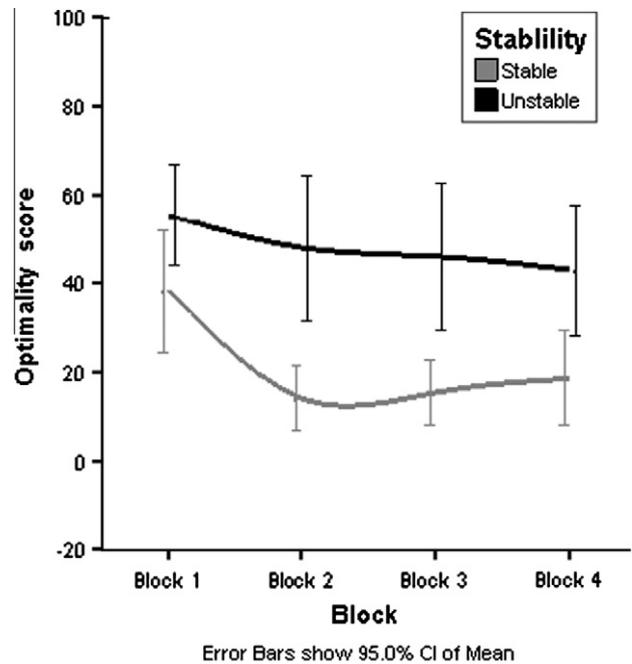


Fig. 3. Control optimization by condition.

There was also a significant Stability × Block interaction, $F(3, 75) = 3.276, p < .05, \text{partial } \eta^2 = .15$. Univariate analyses conducted comparing error scores separately for each block revealed that there were significant differences between both conditions for each of the four blocks of trials ($p < .005$).

3.2. Control optimality

A 4 × 2 ANOVA was conducted on control optimality scores with Block (learning block 1, 2, 3, 4) as a within subject factor, and Stability (Unstable, Stable) as a between subject factor. As indicated in Fig. 3, participants improved in their ability to select the optimal cue settings as they became more familiar with the task. This was confirmed by a main effect of Block, $F(3, 75) = 11.09, p < .0001, \text{partial } \eta^2 = .37$. There was a main effect of Stability on error scores $F(1, 25) = 18.38, p < .001, \text{partial } \eta^2 = .55$, indicating that those in the Unstable condition chose less optimal cue values than participants in the moderate perturbation condition.

3.3. 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. A 3 × 2 ANOVA was conducted on the mean proportion of changes to cues across the all 200 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,

$F(2, 50) = 0.18, p > .05$, partial $\eta^2 = .001$, implying that the occasions on which the three different cues were intervened upon was equally distributed across the 200 trials. Confirming the suggested trend in Fig. 4, there was a significant main effect of Stability, $F(1, 25) = 19.22, p < .0005$, partial $\eta^2 = .54$, showing that those in the Stable condition manipulated the three cues less frequently than the Unstable condition.

3.4. Parameter setting

In addition to examining the proportion of trials in which the cues were manipulated, as a further indication of sensitivity to the underlying stability of the system, we examine the range of values selected for each of the three cues.

Fig. 5 shows that overall the values for the three cues appears to be lower in the Stable condition as compared with the Unstable condition. Confirming this trend, A 3×2 ANOVA was conducted on mean values selected for the three 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, $F(1, 25) = 19.17, p < .0001$, partial $\eta^2 = .54$. No other effects were significant.

3.5. Strategy

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 the entire 200 trials as shown in Fig. 6.

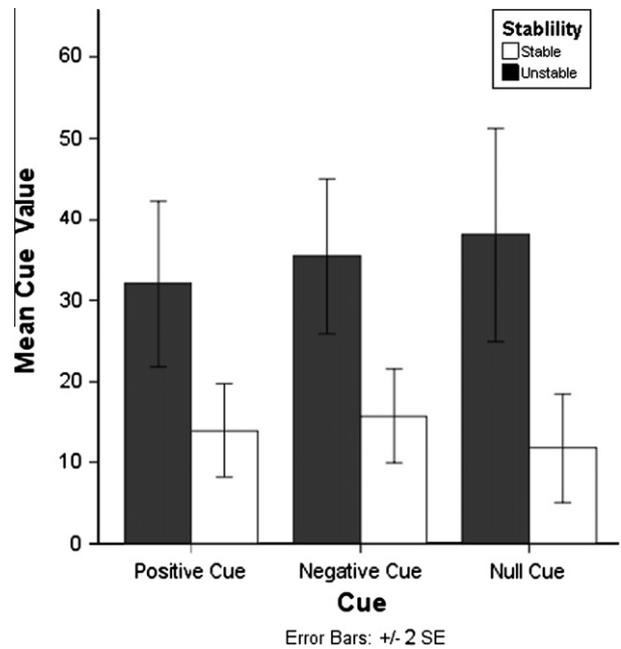


Fig. 5. Mean SE (\pm) values set for each type of cue by condition.

The second set of analyses considers the profile of strategy development across blocks of control trials. To begin, a 4×2 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, $F(3, 75) = 10.87, p < .0001$, partial $\eta^2 = .36$, suggesting that

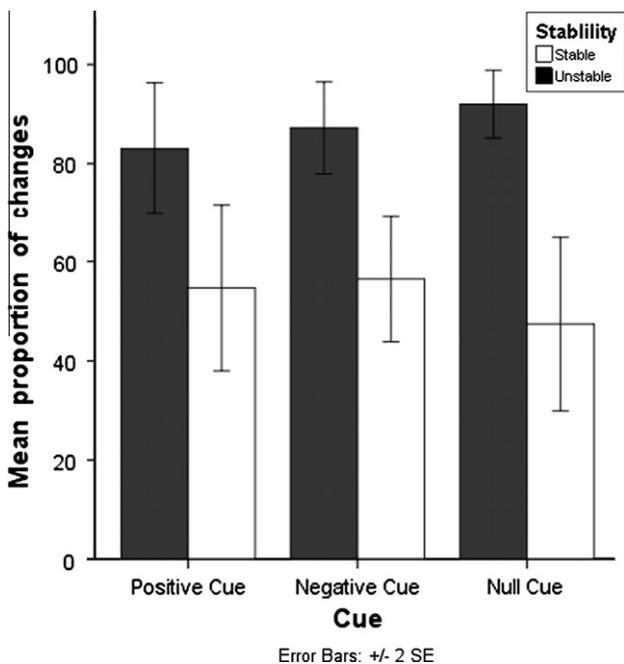


Fig. 4. Mean SE (\pm) proportion of occasions that cues were manipulated by condition.

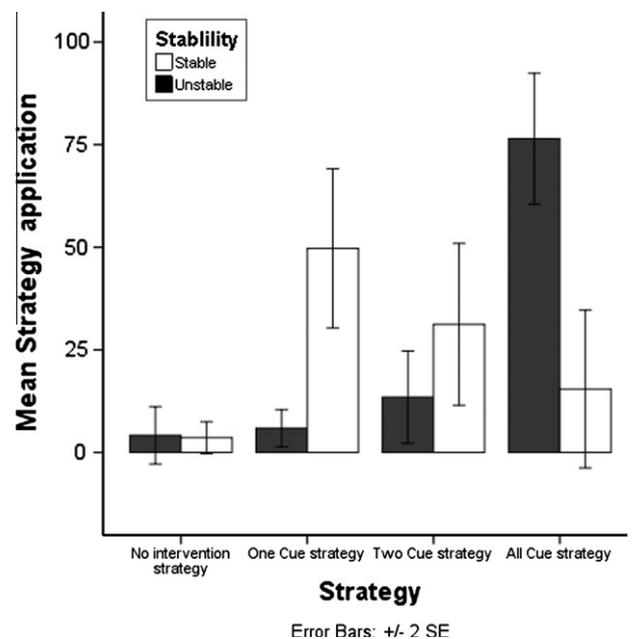


Fig. 6. Mean SE (\pm) proportion of the four strategies employed by condition.

there were differences in the types of strategies favoured overall, as indicated in Fig. 6. There was also a main effect of Stability, $F(1, 25) = 24.85$, $p < .0005$, partial $\eta^2 = .59$, and a significant Strategy \times Stability interaction $F(3, 75) = 14.98$, $p < .0005$, partial $\eta^2 = .49$. To locate the source of the Strategy \times Stability interaction, univariate analyses revealed that compared with the Stable condition, the Unstable condition applied the All-Cue-strategy more frequently, $F(1, 25) = 23.95$, $p < .0001$, partial $\eta^2 = .59$, whereas the One-Cue-strategy was applied less often, $F(1, 25) = 21.73$, $p < .0005$, partial $\eta^2 = .55$. No other analyses were significant.

3.6. Strategy development

The final set of analyses concern the way in which strategy application may have changed across trials. Here we focus on strategy application separately for each condition across the four blocks of trials. For the Unstable condition, a 4×4 ANOVA revealed that there was a main effect of Strategy, $F(3, 75) = 43.40$, $p < .0001$, partial $\eta^2 = .79$, however no main effect of Block suggesting that the distribution of the different strategies that were applied across blocks remained the same. Post hoc tests revealed that the All-Cue-strategy was applied more often than any of the other strategies, and that the remaining three strategies were equally unpopular ($p < .0005$). No other analyses were significant suggesting that most participants' preference was for the All-Cue-strategy and that it was applied consistently across trials.

For the Stable condition the extent to which the four different strategies were applied differed, $F(3, 75) = 20.89$, $p < .001$, partial $\eta^2 = .55$. Again Post hoc tests were conducted, and revealed that the One-Cue-strategy was the most popular compared with the other strategies, and the No-Intervention-strategy was the least popular as compared with all other strategies ($p < .05$).

4. 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 from this study 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.

4.1. Differences between Unstable and Stable conditions

The requirements of the task were such that participants had to gain knowledge of both exogenous and endogenous influences on the outcome, which was made harder to discover in the Unstable condition. The study showed that 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. As proposed by the MC Framework and studies of motor control, introducing instability into system may have resulted in poorer performance for the following reasons. Clearly, the fluctuations in the outcome value encouraged participants 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. 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. Forth, the popular strategy used by the Unstable condition involved varying all three cues, whereas in the Stable condition people tended to varying one cue at a time. Previous findings also suggest that varying one cue at a time is a more successful strategy for controlling a system as compared with varying all cues at the same time (Tschirgi, 1980; Vollmeyer, Burns, & Holyoak, 1996).

4.2. 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, suggesting that in general people are sensitive to the instability of the system and adapt their decision making behavior accordingly. However, unlike previous studies that have examined the effects of instability on decision making (e.g., by varying the dynamics of the system, feedback delays), the present study revealed new insights into control behavior. In general, regardless of the stability of the system, people utilized all three cues equally, and the range of values that were set for each cue were approximately to the same level. This suggests that people in Stable and Unstable conditions were insensitive to the null cue. 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 cue over a series of consecutive trials. In this way it would be easier to detect the dissociation between actions and effects. 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 one 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 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).

Taken together, some aspects of cue utilization and strategy application were similar in both groups, however we propose that the underlying reason for this is in fact different for the two conditions, and reflects sensitivity to the instability of the system, in line with the general proposals made by the MC framework. Over time, judgments of uncertainty will converge with the objective characteristics of the control system which make it uncertain (e.g., dynamic properties), because our learning and decision making mechanisms are driven by them, and are sensitive to them (Osman, 2010b).

5. Conclusion

Until now, the findings have been mixed concerning whether dynamic properties *per se* cause problems for people to the extent that they cannot control a complex system (Gonzales, 2005; Kerstholt, 1996; Hagmayer et al., 2010). Moreover, when the system operates autonomously, again, it is unclear from previous studies whether this is a problematic feature of the system that makes it hard for people to control. In the present study we aim to bring clarity to these issues. We show that over time people can successfully improve their ability to control a system which is dynamic and autonomous. We propose that the manipulation of instability affected both the perceived dynamics of the system (i.e. the fluctuation of the outcome from trial to trial) and the perceived autonomous nature of it (i.e. the extent to which it was perceived as operating on its own because of the lack of correspondence between actions and outcomes). We also show that people are sensitive to the underlying properties of the system to extent that their cue utilization and strategy application differed in Stable and Unstable conditions. However, there appear to be aspects of the decision making process that are preserved in both Unstable and Stable conditions, and this may reflect a more general underlying adaptive strategy to avoid shifting strategies at the cost of successfully controlling an outcome because control is a goal-directed pursuit.

Acknowledgments

The support of the ESRC Research Centre for Economic Learning and Human Evolution is gratefully acknowledged. 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

- Berry, D. C., & Broadbent, D. E. (1988). Interactive tasks and the implicit–explicit distinction. *British Journal of Psychology*, *79*, 251–272.
- Blakemore, S. J., Frith, C. D., & Wolpert, D. M. (2001). The cerebellum is involved in predicting the sensory consequences of action. *NeuroReport*, *12*, 1879–1884.
- Blakemore, S. J., Wolpert, D. M., & Frith, C. D. (1998). Central cancellation of self-produced tickle sensation. *Nature Neuroscience*, *1*, 635–640.
- Brehmer, B. (1992). Dynamic decision making: Human control of complex systems. *Acta Psychologica*, *81*, 211–241.
- Broadbent, D., & Ashton, B. (1978). Human control of a simulated economic system. *Ergonomics*, *78*, 1035–1043.
- 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.
- Diehl, E., & Serman, J. D. (1995). Effects of feedback complexity on dynamic decision making. *Organizational Behavior and Human Decision Processes*, *62*, 198–215.
- Dienes, Z., & Fahey, R. (1995). Role of specific instances in controlling a dynamic system. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*, 848–862.
- 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.
- Gibson, F. P. (2007). Learning and transfer in dynamic decision environments. *Computational and Mathematical Organizational Theory*, *13*, 39–61.
- Gonzales, C. (2005). Decision support for real-time, dynamic decision making tasks. *Organizational Behavior and Human Decision Processes*, *96*, 142–154.
- Gonzales, C., Lerch, J. F., & Lebiere, C. (2003). Instance-based learning in dynamic decision making. *Cognitive Science*, *27*, 591–635.
- Hagmayer, Y., Meder, B., Osman, M., Mangold, S., & Lagnado, D. (2010). Spontaneous causal learning while controlling a dynamic system. *Open Psychology Journal*, *3*, 145–162.
- Kerstholt, J. H. (1996). The effect of information cost on strategy selection in dynamic tasks. *Acta Psychologica*, *94*, 273–290.
- Kerstholt, J. H., & Raaijmakers, J. G. W. (1997). Decision making in dynamic task environments. In R. Ranyard, R. W. Crozier, & O. Svenson (Eds.), *Decision making: Cognitive models and explanations* (pp. 205–217). New York, NY: Rutledge.
- 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. (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*. Oxford: Wiley-Blackwell Publishers.
- Osman, M., Wilkinson, L., Beigi, M., Parvez, C., & Jahanshahi, M. (2008). The striatum and learning to control a complex system? *Neuropsychologia*, *46*, 2355–2363.
- Osman, M., & Speekenbrink, M. (submitted for publication). Learning to predict vs. learning to control a dynamic environment.
- 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.