



THE UNIVERSITY OF
WAIKATO
Te Whare Wānanga o Waikato

Research Commons

<http://researchcommons.waikato.ac.nz/>

Research Commons at the University of Waikato

Copyright Statement:

The digital copy of this thesis is protected by the Copyright Act 1994 (New Zealand).

The thesis may be consulted by you, provided you comply with the provisions of the Act and the following conditions of use:

- Any use you make of these documents or images must be for research or private study purposes only, and you may not make them available to any other person.
- Authors control the copyright of their thesis. You will recognise the author's right to be identified as the author of the thesis, and due acknowledgement will be made to the author where appropriate.
- You will obtain the author's permission before publishing any material from the thesis.

**An Analysis of the Impact of Reinforcement on Response
Variability on Multiple Behavioural Dimensions**

A thesis
submitted in fulfilment
of the requirements for the degree
of
Doctor of Philosophy
at
The University of Waikato
by
Xiuyan Kong



THE UNIVERSITY OF
WAIKATO
Te Whare Wānanga o Waikato

2016

Abstract

A series of four experiments were carried out to examine whether increases in variability in responding on some dimensions of a behaviour as a result of reinforcement generalised to other dimensions of the same behaviour. In Experiment 1, adult participants drew 300 rectangles on a computer screens; variability in the Area, Shape and Location dimensions of the rectangles drawn was measured. One group of participants received reinforcement when they drew rectangles varying on all three dimensions (VAR) using a threshold contingency to determine reinforcement while another group received reinforcement regardless whether they varied on any of the dimensions of the rectangles drawn (YOKE). Results showed that the variability of each dimension, as measured by the U-value, was higher for the VAR group than for the YOKE group. In Experiment 2, three groups of adult participants received reinforcement when they varied on two of the three dimensions of the rectangles drawn. Results showed that reinforcing variability in the Shape and Location dimensions and in the Area and Location dimensions resulted in higher variability in these dimensions compared to the dimension where variability was not directly reinforced, as measured by U-values. However, for the group who received reinforcement when varying in the Area and Shape dimensions, the variability across the three dimensions did not differ. Generalisation of reinforced variability was examined by comparing the variability in dimensions that were not subject to reinforcement for varying from Experiment 1 and Experiment 2. For the Shape and Location dimensions, results showed that the variability was higher when the two dimensions occurred with other dimensions that were subject for reinforcement for varying compared to when they occurred with other dimensions that were not required to vary. However, no significant difference was found for the Area dimension. Therefore, there appeared to be generalisation of reinforced variability for only two of these three groups of participants, suggesting the dimensions were not orthogonal. The non-orthogonality of the three dimensions of the rectangles was suggested as a possible confounding variable so that it not be concluded that generalisation of reinforced variability across dimensions had been obtained. Experiment 3 used a task with more independent dimensions (Shape, Colour and Pattern) to control for the non-orthogonality issue found in Experiment 2. Adult participants created

objects by selecting elements from three dimensions on the computer screen. One group of participants received reinforcement when they varied on the use of elements on all three dimensions using a threshold contingency while the other group received reinforcement independent of the elements chosen. Results showed that the variability in using the elements for all three dimensions was similar between the two groups; whether the group received reinforcement when the dimensions varied did not make a difference. Inspection of the elements selected on each trial revealed that participants from both groups, in general, responded in stereotypic patterns; however, the stereotypy in the responses was not picked up by the measure of U-value. In Experiment 4, a new task asking participants to colour t-shirts on the computer screen was created to examine the effect of reinforcement on the variability in colours used. Adult and adolescent participants were instructed to colour t-shirts on the computer screen in three phases. In Phase 1, no feedback was provided; in Phase 2, positive feedback was provided when participants used a colour they had not used before (including those used in Phase 1); and in Phase 3, no feedback was provided. Results showed that, after reinforcement for using novel colours, participants who used only a small number of colours initially used more colours; however, the increase was not found for participants who initially used a moderate or high number of colours. U-value as the measure of variability was examined by using two sets of simulated data. It was found that highly stereotypic response patterns can result in extremely high U-values. Also, the number of options/categories used affects the values of U, which makes comparisons between groups difficult. Overall, results from these experiments support the notion that the level of variability in responses can be controlled by reinforcement. It is possible that learned variability can generalise over unreinforced dimensions. However, the examination of generalisation of reinforced variability across dimensions is not conclusive when the dimensions are interdependent. The results and simulation showed that U-value as a measure of variability needs to be used with caution because it does not capture potential stereotypy in responses and the values of U can be ambiguous.

Acknowledgement / 致谢

First and foremost, my deepest thanks go to my supervisors Dr. James McEwan, Prof. Mary Foster and Dr. Lewis Bizo who have shared with me their wisdoms and have guided and supported me through the whole journey. It has been a journey with joy, frustration and many other emotions; I would not have been able to get to the the finishing line without the understanding, patience and encouragement I have received from them. I would also like to thank Professor Allen Neuringer for his kindness and for his passion for research in behaviour variability that has been inspiring me from the very early days of my PhD journey.

The PhD life will be so much different without the understanding and support from you all, Sandy, Surrey, Kathleen, Leanne and Pim. Special thanks to Kath and Leanne with whom I felt had fought a war together. I'm so grateful for Surrey Jackson's help with proofreading this thesis and for other students in the PhD office for sharing their laughs and cries with me.

I would also like to thank Andrew Malcolm, our super awesome programmer, for writing the computer programmes for my experiments. Thank you Allan Eaddy for the computer and software support. Also thank you Jan Cousens and Joy Fellows for your administration supports which made the running the experiments easier. And finally thank you to all who participated in the experiments; the thesis would not have been completed without your input.

特别感谢我的父母，没有您们无条件的爱、包容、理解和支持，我不可能走完这段求学路程。同样感谢一直爱护、关心和给我温暖的姐姐、弟弟、小婉、涛姐和排骨。当然还有家里几个给我无限快乐的小鬼头 Owen 哥，枫枫和圣圣！我还要感谢在新西兰如同我家人的朋友们，Daniel，Fiona，Annie 和 Andy，感谢你们一直以来对我和阿珍的照顾和关心。非常感谢一直给我精神支持的于我亦师亦友亦偶像的蔡芳钿老师，感谢你对我的信任、理解、支持和鼓励。我还要感谢在等待论文结果和准备答辩这段时间给我地方冷静思考、分散注意力缓和紧张情绪的 Jett Leung 和 Eric 摄摄下先生。最后再一次谢谢所有这些年来关心支持我的家人朋友们，生命里面有你们是我的福气和幸运！

Table of Contents

Abstract	ii
Acknowledgement / 致谢	iv
Table of Contents	v
List of Figures	vii
General Introduction	1
Experiment 1	36
Introduction	36
Methods.....	37
Measures	40
Results and Discussion.....	41
Experiment 2	46
Introduction	46
Method	46
Measures	47
Results and Discussion.....	48
Experiment 3	59
Introduction	59
Methods.....	60
Results	64
Discussion	68
Experiment 4	74
Introduction	74
Methods.....	75
Results and Discussions	79
An Analysis of U as a Measure of Variability	87
Introduction	87
Simulation 1 – Stereotypical Responding	88
Simulation 2 – U-values from using different number of the available options	90
SUMMARY	93

General Discussion.....	95
References	103
Appendix I.....	109
Appendix II	113
Appendix III.....	127

List of Figures

- Figure 1.1. The moving average across 30 trial blocks of the percentage of participants in the VAR and YOKE groups that met the variability criteria during Experiment 1. The solid lines are the moving average and the dotted lines are lines fitted by the method of least squares. ... 42
- Figure 1.2. Mean U-values for each dimension for the VAR and YOKE group. High U-values indicate high variability. Error bars show standard errors. Values below .6 have been omitted from the y-axis..... 43
- Figure 1.3. U-values plotted as a function of scores for Area, Shape and Location dimensions for both VAR (filled triangles) and YOKE (empty triangles) groups. U-values of VAR and YOKE that correspond to the same score were from individuals who were yoked to each other and received the same number of points..... 45
- Figure 2.1. Reinforcement arrangements. Filled boxes show dimensions that were required to vary for reinforcement; empty boxes show dimensions that were not required to vary for reinforcement. 48
- Figure 2.2. Mean U-values for Area, Shape and Location dimensions when they were not required to vary across Experiment 1 and Experiment 2. Error bars represent standard errors. U-values below .7 were omitted from y-axis. 49
- Figure 2.3. Moving average of the cumulative number of trials WRF less than .0825 for Area, Shape and Location dimensions for Non_Area, Non_Shape and Non_Location groups. 52
- Figure 2.4. Mean number of trials WRF < .0825 for Area, Shape and Location dimensions for Non_Area, Non_Shape and Non_Location groups. Error bars represent standard errors. 53

Figure 2.5. Mean U-values for Area, Shape and Location dimensions for Non_Area, Non_Shape and Non_Location groups. Error bars represent standard errors.	54
Figure 3.1. Percentage of participants meeting reinforcement criteria over 300 trials for both VAR (solid line) and YOKE (dotted line) group. Each datapoint represents the moving average of 30 trials.....	64
Figure 3.2. Mean U-values for the three dimensions for VAR and YOKE groups. The three sets of bars on the left were for U-values calculated based on 8 categories; and the three sets on the right were calculated based on 64 categories (Lag 1). Error bars represent standard errors.	65
Figure 3.3. U-values calculated based on 8 (filled circles) and 64 (empty circles) categories for Shape, Colour and Pattern dimensions for 17 pairs of VAR-YOKE participants. Data points corresponding to Random participant were U-values calculated from 300 computer generated random responses (between 1 and 8).	71
Figure 4.1. Mean number of colours used in Phase 1, Phase 2 and Phase 3. Number of different colours were represented by dark bars, number of new colours represented by light grey bars and the number of colours in Phase 3 that were new to Phase 1 was in dark grey bar. Error bars represent standard errors.	80
Figure 4.2. The number of colours used in Phase 1 and Phase 2 (panel 4.2a), in Phase 1 and Phase 3 (panel 4.2b.) and the number of colours used in Phase 1 and the number of colours used in Phase 3 but not in Phase 1; panel 4.2c).....	82
Figure 4.3. Mean number of different colours used in the three phases for groups of participants using least number of colours (Bottom), intermediate number of colours (Middle) and highest number of colours (Top) in Phase 1. Error bars represent standard errors.....	83

Figure 4.4. Mean number of colours used in Phase 1 (dark bars) and in Phase 3 (light grey bars) for Bottom, Middle and Top groups. Error bars represent standard errors.	85
Figure 5.1. U-values from four types of simulation of 100 responses over 4 categories and Frequencies of responses to each category from each type of simulation.	89
Figure 5.2. U-values resulting from 300 responses over different number of categories out of 16 used.	91
Figure 5.3. U-value from Experiment fitted in the figure with simulation data of 300 responses over 16 categories. Filled circles represent data from VAR group and empty circles represent those from the YOKE group.	94

General Introduction

Being able to generate variable responses is considered to be one of the key characteristics of creativity (Campbell, 1960; Neuringer, 2002). It is also considered to be an important aspect of the behaviour repertoire organisms play and explore the environment (Baldwin & Baldwin, 1978). Epstein (1996) suggested that it aids effective problem solving in everyday life; the absence of it is believed to impair functioning in everyday life significantly (e.g., depression and autism; Neuringer, 2002). For these reasons, variability in responding is an important area of study, especially studies are needed that attempt to find out what factors influence behavioural variability and how variability can be controlled so that new effective repertoires can develop (Neuringer, 2002, 2004; Neuringer & Jensen, 2013) or to better understand and enhance creativity (Runco, 1993; Stahlman, Leising, Garlick, & Blaisdell, 2013; Stokes, 2012).

The definition of variability is dependent on the context in which it is studied. Variability can be defined with reference to the number of different (within a defined period of time) or novel responses in some case while it can be defined with reference to response options that are used less frequently. Generally, variability is on a continuum; repetition and stereotypy represent the lowest end of variability as behaviour is highly predictable while stochasticity and randomness represent the highest variability as behaviour is highly unpredictable (e.g., Neuringer, 2002; Neuringer & Jensen, 2012). Researchers have been studying the control of reinforcement over behavioural variability, treating variability as an operant dimension. An aspect of behaviour is more likely to occur in the future if it has produced desirable outcomes for the individual; such outcomes are referred to as reinforcers (Skinner, 1938). The likelihood of this increased probability of behaviour occurring is higher when there are stimuli that signal the desirable consequences are available; these stimuli are termed discriminative stimuli (Skinner, 1938), and behaviour would be said to be under stimulus control if the probability of it occurring is affected by discriminative stimuli.

The extent to which variability in responses can be controlled by reinforcers and discriminative stimuli has been studied extensively over a number of decades (see reviews Neuringer, 2002; Neuringer, 2004; Neuringer, 2010;

Neuringer, 2012; Neuringer, 2013). One of the earliest studies that attempted to systematically test whether behavioural variability can be reinforced was done by Schwartz (1982). In his study, pigeons were trained to peck on two keys (left and right) to generate eight response sequences; each sequence included four pecks on the left key and four on the right one. Later, in order to encourage variability in the sequences, the pigeons received reinforcers only when the current sequence generated differed from the one produced in an immediately preceding trial. Schwartz (1982) concluded that requiring variability for reinforcement did not result in more variable sequences.

Page and Neuringer (1985) used similar procedures to those used by Schwartz (1982) but found contradictory results. In one condition, in order to increase the variability in the eight response sequences, Page and Neuringer (1985) provided reinforcers for sequences that consisted of four responses to the left key and four responses to the right only when a sequence differed from the one generated in the immediate preceding trial; while in another condition, they provided reinforcers for any sequence that was different from the one prior to the current one, without restricting the sequences to include exactly four responses to the left key and four to the right key. They found that pigeons generated more variable sequences under the condition where the requirement of exactly four responses to each of the left and right key was present. In the condition where the four responses to each key was required, pigeons failed to obtain reinforcement not because they repeated their sequences but because the trials terminated with more than four responses to one key. Page and Neuringer (1985) suggested that the failure to find increased variability in Schwartz's study was due to the additional requirement of four responses on each key — only 70 of the 256 possible sequences met this requirement; there were not enough opportunities to reinforce variable responses.

As well as showing that reinforcement contingencies can control response variability, Page and Neuringer (1985) also showed that the amount of variability can be specified by contingencies using different lag schedules. Under lag procedures, reinforcement is provided for responses that are different from a specified number of responses prior to the current response. Take a Lag 5 schedule in their experiment for example, if a pigeon created a sequence of

LRLRRRL on the 11th trial, the pigeon will receive a reinforcer only if the sequence of 'LRLRRRL' had not been generated on the 10th, 9th, 8th, 7th, and 6th trial. If the pigeon repeated any of the sequences from these five trials, no reinforcement would be provided. The higher the number of prior trials the current trial had to differ from, the higher the variability requirement was. In one experiment, Page and Neuringer (1985) manipulated the number of previous sequences the current sequence had to differ from for reinforcement in an incremental manner, using Lag 5, Lag 10, Lag 15, Lag 25 and Lag 50. The variability in the sequences generated by the pigeons increased as the requirement for reinforcement increased; the higher the variability requirement (lag requirement), the more variable the sequences were, from Lag 5 through Lag 25; while the variability decreased slightly when requirement was Lag 50. The control of reinforcement on variability was shown by comparisons to the results from a condition when reinforcement was provided independent of variability, under this condition, repetitive sequences were generated.

The control of reinforcement over behavioural variability has been demonstrated not only in sequences generated by pigeons and rats, but also in other behaviours by other species. For example, studies have shown that reinforcement can be used to increase the variability in inter-response times between panel-presses in rats (Schoenfeld, Harris, & Farmer, 1966) and between key pecks in pigeons (Blough, 1966), in pigeons' key pecking sequences (Doughty, Giorno, & Miller, 2013; Morris, 1989), in key press sequences on computer keyboards by college students (Da Silver Souza, Abreu-Rodrigues, & Baumann, 2010; Maes, 2003; Neuringer, Deiss, & Olson, 2000; Stokes, Mechner, & Balsam, 1999) and by primary school children (Stokes, Holtz, Carlis, & Eisenberg, 2008) and in sizes, shapes and locations on the screen of rectangles drawn by college students (Ross & Neuringer, 2002).

The control of variability by discriminative stimuli has also been demonstrated. In their sixth experiment, Page and Neuringer (1985) provided reinforcers when pigeons generated variable sequences under one coloured key light (VARY) and when they generated one particular sequence under another coloured key light (REPEAT). They found that pigeons learned to vary and repeat according to the different signals, even when the lights were reversed for the

conditions. In another study, Denney and Neuringer (1998) first compared the variability of lever press sequences produced by a group of rats between two conditions with different signals (discriminative stimuli); reinforcement was provided contingent on less frequent sequences under one condition, signalled by one stimulus, and reinforcement was provided probabilistically and independently of the sequence variability in the second condition, signalled by another stimulus. The amount of reinforcement between these two groups was equalised. It was found that the rats produced more variable sequences when the discriminative stimulus signalled the requirement to vary the sequences, compared to when the discriminative stimulus signalled no requirement to vary the sequences. Later, with the reinforcement contingencies unchanged for these two groups, Denney and Neuringer (1998) arranged the conditions so that the discriminative stimulus signalling the contingent reinforcement for variability was first present, then absent and finally present again. Rats learned to generate more variable sequences when the reinforcement was contingent on variability compared to when variability was not required only when the signals were present. Results from both Page and Neuringer (1985) and Denney and Neuringer (1998) had shown the control of variability by discriminative stimulus; not only can variability be reinforced, it can also be controlled as to when and where by using discriminative stimuli.

The studies discussed above have demonstrated the control of contingencies over variability in highly controlled and arbitrary contexts; whether to vary or not, when and where to vary, and how much to vary can all be controlled by reinforcement and discriminative stimuli. Because of the importance of variability in creativity, problem solving and in informing interventions for individuals with limited behavioural repertoires, studies have also examined the control of contingencies over variability in more naturalistic context and in practical and applied settings.

Being able to make variable or novel responses is believed to be related to creativity; therefore, reinforcing response diversity that could lead to novel responses could aid enhancing creativity (Goetz, 1989). A number of research studies had shown that social feedback (e.g., Fallon & Goetz, 1975; Goetz & Baer, 1973; Goetz, Jones, & Weamer, 1973; Kratochwill, Rush, & Kratochwill,

1979), commonly in the form of verbal praise, can be used to increase response diversity and the probability of occurrences of novel responses.

In one study, the number of different construction forms and the number of new forms in a block-building activity of three preschool children was recorded across different experimental conditions (Goetz & Baer, 1973). Children were first asked to play with the blocks without receiving any feedback so that baseline form diversity could be assessed. They then received descriptive praise every time they constructed a form that was new in a session. They were then praised for creating the same forms they had created in another session. Finally, they were again praised when they built new forms in a fourth session. All three children used more forms within the sessions when social feedback was contingent on new forms compared to their baseline performance and to when repetition of the same forms were reinforced. There was a big increase in the number of new forms in the condition where form diversity was first reinforced for two children; but the increase was apparent for only one child in the last condition when form diversity was again reinforced. In this study, Goetz and Baer had demonstrated praise delivered contingent on novel form creations can increase the diversity in the forms in block-building play in preschool children.

As well as block-building, reinforcement has been used to increase the number of forms created in children's drawing with felt-pen (Fallon & Goetz, 1975; Holman, Goetz, & Baer, 1977) and with easel painting (Kratochwill et al., 1979); the number of colours and forms in children's felt-tip pen drawing have also been shown to be increased by reinforcement (Ryan & Winston, 1978). In addition to the use of commonly studied species, reinforcement has also been shown to increase the variability of vocal repertoires in Budgerigars (Manabe, Staddon, & Cleaveland, 1997), of sounds produced by Pacific walrus (Schusterman & Reichmuth, 2008) and of swimming and jumping forms by porpoises (Pryor, Haag, & O'Reilly, 1969). Consistent findings from these studies suggest that variability, as an important aspect of creativity, can be reinforced.

The results from animal experiments showing that behavioural variability can be reinforced have not only been applied to study the potential for increasing creativity in typically developing individuals but also to better understand and to inform interventions for behavioural problems in people with restricted

behavioural repertoires, such as individuals with depression and those with autism. For example, when asked to generate sequences on computer keyboards without any feedback, moderately depressed college students generated less variable sequences than the non-depressed students (Hopkinson & Neuringer, 2003). When a high level of variability was required for reinforcement, the sequences generated became increasingly variable for both the depressed and non-depressed students; the variability in the sequences generated by the depressed students ended as high as they were for the non-depressed students. In summary, reinforcement enable moderately depressed students to respond as variable as normal students.

Another area where increased variability in responding is desired is research with individuals with autism. Individuals on the autistic spectrum typically show limited behaviour repertoires; research has studied the role reinforcement can play in expanding these repertoires (for a review, refer to Wolfe, Slocum, & Kunnavatana, 2014). It has been shown that direct reinforcement on variability increased the variation in sequence generation on computer keyboards in adolescents and children (Miller & Neuringer, 2000; Murray & Healy, 2013), in the number of phonemes used in children (Koehler-Platten, Grow, Schulze, & Bertone, 2013), in the number of different vocalisations in children and adults (Esch, Esch, & Love, 2009; Lee, McComas, & Jawor, 2002; Susa & Schlinger, 2012), in block-building constructions (Napolitano, Smith, Zarcone, Goodkin, & McAdam, 2010) and also in the imaginative play and drawing of children (Newman, Reinecke, & Meinberg, 2000). These studies, varying in reinforcement schedules and targets of tasks, show that reinforcing variability can expand the behavioural repertoires in individuals with autism.

In addition to showing reinforcing variability can increase one aspect of creativity and can expand the behavioural repertoires in individuals with depression and autism, the potential advantage of enhancing behavioural variability for learning and exploration in new environment and problem solving has been also explored, although not extensively. It was found that reinforcing variability in sequences could potentially facilitate the learning of a difficult sequence in rats (Grunow & Neuringer, 2002; Neuringer, 1993; Neuringer et al., 2000). More research, however, is needed in this regard because results from

human studies have been contradictory (Doolan & Bizo, 2013; Maes & van der Goot, 2006). It was also found that after being trained to vary their ways of interacting with some objects in one environment, rats were more successful in finding food hidden in novel objects in a different environment (Arnesen, 2000; Weiss & Neuringer, 2012). Results from the studies mentioned above suggest that reinforcing behavioural variability could potentially increase one's tendency to explore in new environments. Finally, reinforcing variability is likely to benefit problem solving was shown in a study with children's tool use and improvisation (Parsonson & Baer, 1978). Parsonson and Baer (1978) assessed children's problem solving skills in three different tasks, pounding with different objects on a pounding toy, storing marbles in containers with a hole and finding suitable items to substitute shoelaces. Both the number of different tools use and the number of improvisation (combination of tools) increased when the children were praised every time they used a different tool within a session. Reinforcement seemed to have enhanced children's problem solving skills by expanding the solutions they could use. Learned variability has been shown to have good resistance to change in face of distractors (Doughty & Lattal, 2001) or delays of reinforcement (Odum, Ward, Barnes, & Burke, 2006); and it was not impaired under the influence of drugs like d-amphetamine (Ward, Bailey, & Odum, 2006), or alcohol (Cohen, Neuringer, & Rhodes, 1991).

Schedules of reinforcement

Various contingencies have been used to directly increase or maintain variability in responses; the choice was usually made based on the behaviour or dimension of a behaviour under investigation. This section is going to review the methods that have been used when studying reinforcement on variability.

Reinforcing novel responses. Reinforcers are provided for responses that are new to the individual, either new in a session or across all sessions. Pryor et al. (1960) in their study that reinforced variable jumping forms in porpoises, provided reinforcers for only forms that had never occurred in any previous sessions. The porpoises produced increasingly different forms in their jumps, some of which had never been observed in that species. In some other studies, reinforcers were provided to any new responses within a session, regardless of

whether those responses had occurred in previous sessions. For example, Goetz and Baer (1973) reinforced children every time they constructed a form that was new within a session. The children produced more forms within a session and produced more new forms over sessions when the reinforcement contingency was in place. Similar results were found in training two Pacific walruses to produce more variable sounds by providing fish for making sounds that were new in a session (Schusterman & Reichmuth, 2008). Reinforcing novel responses within a session has also shown to be effective in increasing the variability in mand frames in two children with autism after the frames were introduced (Betz, Higbee, Kelley, Sellers, & Pollard, 2011). The advantage of reinforcing novel responses for increasing behaviour variability is that the sets of responses can be produced is not necessarily known to the researchers; the reinforcement contingency can help expanding the behaviour repertoires for those individuals to get in touch with reinforcers. However, this method suffered from serious methodological limitations. One limitation is that when reinforcers are provided only for responses that are new compared to all previous responses, it could keep participants from getting into contact of reinforcers at all (Neuringer, 2002). Also, as the number of sessions increases, the difficulty in making a response that is new to all previous responses increases (Neuringer & Jensen, 2013). A second limitation is that, practically, as the number of different responses increases over sessions, it becomes more and more difficult for the observer to determine whether a response has been produced before (Neuringer & Jensen, 2013). This could raise concerns about the ability of the observer to accurately and immediately deliver reinforcers for appropriate responses or to withhold reinforcers for responses that do not meet reinforcement criteria; such a concern had been raised by Holman et al. (1977). Nevertheless, it seems that, for individuals starting with extremely low variability and with the increase in new responses not being large, and with the help of technology (e.g., computer programmed reinforcement deliveries) to prevent human errors, it is still an appropriate method to use in practical settings.

Reinforcing switching and staying. Under schedules that reinforce switching and staying behaviours, reinforcers are delivered sometimes for switching and sometimes for staying to encourage response variability. For

example, Bryant and Church (1974) reinforced responses on two levers (left and right) three groups of rats; one group of rats received reinforcement every time they switched between levers, another group received reinforcement on 75% of switches and on 25% of staying on the same lever (no switching), and the last group received reinforcement on 50% of switch responses and on 50% of stay responses. Bryant and Church (1974) found that the group receiving reinforcement for 75% switching and 25% staying produced response patterns that were most similar to random responses; these rats responded to the left or right lever equally often and they were equally likely to switch between levers or stay on the same one. In another study, Machado (1992) found that reinforcing the inclusion of switches in pigeons' 8-response sequences over two keys was sufficient for increasing the variability in the sequences generated.

Reinforcing least frequent responses. Under this schedule, reinforcers are provided for responses that occur least often among the already emitted responses; it is used to encourage the use of less preferred responses so that responses are distributed to available options evenly, just like the random generator would do (Blough, 1966). For example, Blough (1966) created 16 bins with equal number of interresponse times (IRTs) based on random responses. Interresponse times (IRTs) between two pecks in pigeons were monitored; reinforcers were delivered to IRTs that fell into the bins which had occurred least frequently. It was found that the IRTs were distributed exponentially over the 16 bins, similar to random responses. The same schedule of reinforcement was also used to reinforce least frequent response sequences in pigeons (Shimp, 1967). A similar reinforcement schedule was used to increase the variability in communicative gesture use in individuals with mental retardation (Duker & van Lent, 1991). Instead of reinforcing least frequent responses, researchers withheld reinforcers for responses that had occurred most frequently in a preceding condition; reinforcers were delivered to only gestures that had been trained but had not been used frequently. Increased number of different gestural requests were used by all six participants as a result of the reinforcement contingency.

In summary, there are a few characteristics for schedules that reinforce least frequent responses or least frequent sequences. One is that a predefined set of responses or sequences is required. Second, responses under the reinforcement

are expected to be similar to random responses. Third, constant update of the relative frequencies of all available response options or sequences is required to determine whether current response or sequence qualify for reinforcement.

Lag procedures. Another method to reinforce variability that has been used is the lag procedures. Under lag schedules, reinforcers are provided for a response that has not occurred in a specified number of trials immediately prior to the current trial. The number of trials the current trial has to differ from is usually specified by the number stated after the Lag symbol. For example, a Lag 1 schedule requires a response to be different from the one from the previous trial for reinforcement; a Lag 5 schedule requires a response to be different from five preceding trials to the current one. Lag procedures have been used in both experimental and applied settings to increase variability in responses or response sequences. Under experimental settings, studies that used lag procedures to increase variability in responses or response sequences have been shown to successfully increase the variability in pigeon's key pecking sequences under different sizes of Lag procedures (Page & Neuringer, 1985; Ward, Kynaston, Bailey, & Odum, 2008) and in rats' lever pressing sequences (Cherot, Jones, & Neuringer, 1996). Higher Lag requirement was shown to result in higher level of variability in response sequences (Page & Neuringer, 1985). Lag procedures have often been used to increase the variability of functional responses in individuals with autism. For example, Lee et al. (2002) reinforced variable and appropriate vocal responses to a social question under a Lag 1 schedule and found that the appropriate responses were more variable for two of the three individuals when the lag schedule was in place; the overall number of novel appropriate responses increased for the same individuals. Lag 1 schedule has also been found to increase the variability in the overall number of novel appropriate verbal responses to the question "what do you like to do" in two teenagers with autism (Lee, Sturmey, & Fields, 2007) as well as to increase the variability in toy plays in three individuals with intellectual disability (Baruni, Rapp, Lipe, & Novotny, 2014).

Similar to schedules that reinforce novel response, lag schedules do not require a predefined set of responses; new responses that are unknown to the observer can occur as the result. Also, lag schedules that used low lag requirement have been shown to be sufficient in increasing response variability in some

applied research. Furthermore, lag schedules with low lag requirements do not require a large number of responses which is practical for applied settings where large numbers of responses within a short period of time are not viable. Lag schedules are also said to be relatively easy to implement in applied settings (Lee et al., 2007). The advantage of lag schedules over schedules that reinforce novel responses is that the level of variability required for reinforcement can be manipulated by adjusting the sizes of the lags. However, there are also a number of limitations of the lag procedures. First, stereotypical rather than variable responding can develop under low lag requirements. For example, under a Lag 1 procedure, one can obtain reinforcers on all trials by switching between two options (e.g., ABABABAB...); and under a Lag 3 procedure, reinforcers can still be obtained in every trial by cycling through four options, e.g., producing response sequences of ABCDABCDABCD. Another example was that one of the participants from both Lee et al. (2002) and Lee and Sturmey (2006) alternated between two responses to obtain reinforcers. A second limitation is that repetitions can occur in random responses while repetition is never reinforced under lag procedures (Neuringer & Jensen, 2013); the resulting behaviour patterns are therefore different from random responses.

Although a low lag requirement has been shown to produce strategic stereotypical responding in some studies, higher lag requirements could prevent this. For example, Page and Neuringer (1985) used high lag requirements up to Lag 50 with rats, strategic responding under such a high lag requirement would require superior memory which was not possible for the species. In Manabe et al.'s (1997) study with song production in budgerigars, strategic and stereotypical behaviours such as switching between two songs and cycling through three songs were observed as the result of using Lag 1 and Lag 2 schedules; however the stereotypy disappeared when a Lag 3 procedure was used. In another study, a boy with autism was successfully trained to use variable and appropriate responses to the question "How are you?" using Lag 1, Lag 2 and Lag 3 schedules; and his responses did not show stereotypical patterns (Susa & Schlinger, 2012). Therefore, for these two studies, at least with the budgerigars and the individual with autism, a Lag 3 schedule was sufficient to prevent strategic and stereotypical responding. Also, the variability of toy plays in three children with intellectual

disability increased under a Lag 1 procedure, and this did not result in stereotypical responding (e.g., alternating between two responses; Baruni et. al., 2014). Another method to prevent strategic stereotypical responding proposed by Lee et al. (2007) was to use variable lag procedures. Progressive lag procedures (e.g., Susa & Schlinger, 2012) may also work as starting from a low lag maintains responding and increasing the requirement gradually could prevent strategically repeating a response pattern.

Using the lag schedules method in applied studies has shown possibility of stereotypical responding under low lag requirement (e.g., Lag 1); however results have been mixed. Although lag schedules with low lag requirements could lead to strategic and stereotypical responding, this could be prevented by adjusting the schedules according to the behaviour under investigation or according to the different capacity in different species.

Percentile reinforcement schedules. Percentile reinforcement schedules probabilistically reinforce responses or sequences of responses based on the number of trials between two occurrences of the same response or sequence of response (Machado, 1989). The number of trials between the current occurrence and its last occurrence is calculated to form a variability score; these scores are momentarily arranged in an ascending order. Reinforcement for variability is then based on where on the rankings (percentile) the variability score falls. The level of variability can be determined by the criteria specifying which percentile the variability score should fall into or higher; the higher the percentile, the higher the variability requirement is. For example, the number of trials before a response or sequence can be repeated on the first percentile (25%) is smaller than it is on the third percentile (75%). Machado (1989) manipulated the number of trials a sequence emitted by pigeons can repeat for reinforcement when the probability (or intermittency) of reinforcement was kept constant. Pigeons emitted more variable sequences when the requirement for variability was high (e.g., on the third percentile) compared to when the requirement was low (e.g., on the 0 percentile). Miller and Neuringer (2000) also used percentile reinforcement schedules to increase variability in four-response sequences generated over two computer keyboards by individuals with and without autism. After each sequence was emitted, the computer instantly updated the relative frequencies of each of the 16

possible sequences; all relative frequencies were multiplied by an exponential coefficient after a reinforcer was obtained to weight sequences that were reinforced much earlier less. The weighted relative frequencies from the most recent 20 sequences emitted were constantly ranked; and if the weighted relative frequency of the newly emitted sequence was lower than the one ranked 11th lowest from the preceding 20 trials and it was also lower than the threshold value 0.35, a reinforcer was delivered. The percentile reinforcement contingency based on weighted relative frequencies in Miller and Neuringer (2000) successfully increased the variability in sequence generations in all participants, including individuals with autism.

The advantage of percentile reinforcement schedules over lag procedures is that by requiring the same response to be emitted after a certain number of trials, strategic and stereotypical responding, such as alternating between responses or cycling through response options which can occur under Lag 1 and Lag 2 schedules, can be prevented.

Frequency dependent schedules. Under frequency dependent schedules, probabilities of the occurrence of each response option are constantly calculated within a moving window of a predefined number of trials; reinforcers are delivered if the probability of current response is relatively low. Machado (1992, 1993) reinforced pigeons to peck on two keys variably using frequency dependent schedules. In one experiment, reinforcement was contingent on responding on either left or right key, whichever resulted in lower probability of occurrence; in another two experiments, reinforcement was contingent on pairs of responses on the left and right key when the probability of occurrence of the pairs was relatively low; and finally on a fourth experiment, reinforcement was contingent on triplets of responses over the left and right keys. Pigeons developed strong sequential dependent responding patterns (e. g., alternating between the two keys) even though their responses satisfied the reinforcement criteria when the criterion was based on only one response. When reinforcement was based on pairs of consecutive responses, the switching pattern disappeared; most of the pigeons generated random like responses over the two keys. However, some of the pigeons in this condition developed double alternation response pattern (e.g., RLLRLL sequence; Machado, 1992). When reinforcement was dependent on

triplets of key pecks, pigeons generated responses that were closest to random responses without the stereotypical patterns found when reinforcement was contingent on one or pairs of responses. Therefore, frequency dependent schedule based on triplets of responses was shown to be effective in increasing the variability in key peck sequences in pigeons.

Threshold schedules. Under the threshold schedule, the weighted relative frequency of a current response is compared to a threshold value which specifies a desired level of variability by the researcher; only responses resulting in weighted relative frequencies that were lower than or equal to the threshold value will be reinforced. Relative frequencies are calculated by dividing the number of occurrences of an option by the number of trials up to the current trial. The relative frequencies are then multiplied by a weighting coefficient to weight recent responses more. To encourage variability in responding, the threshold value is usually set low so that only infrequently occurring responses are reinforced. For example, Denney and Neuringer (1998) reinforced rats to vary the lever press sequences using threshold schedules. In a condition where variability was reinforced, weighted relative frequencies of the current sequence had to be equal to or lower than the threshold value .09 – only sequences that occurred equal to or less than 9% of the time were reinforced. The rats generated more variable sequences when the threshold schedule was in effect compared to conditions where variability was not required for reinforcement.

In another study, Grunow and Neuringer (2002) studied the control of threshold schedules over the variability of sequences generated by rats by manipulating the threshold values. Rats with reinforcement under a low threshold value produced most variable sequences; the variability in the sequences decreased as the threshold requirement increased. Threshold schedules were also used to increase the variability in the sizes, shapes and locations of rectangles drawn on computer screens by college students (Ross & Neuringer, 2002). In Ross and Neuringer's (2002) study, 16 arbitrary categories were created for each dimension of the rectangles so that relative frequencies for each categories could be calculated. Only rectangles currently drawn with weighted relative frequencies for all three dimensions that were lower than the threshold value resulted in a reinforcer. Participants under the threshold contingency were more likely to draw

rectangles that varied on all three dimensions, Size, Shape and Location, compared to participants who received reinforcement independent of the variability of the rectangles drawn.

Threshold schedules have also been used to increase variability in sequences of key pecking in pigeons (Doughty et. al., 2013) and to increase the variability in sequences generated on computer keyboards by students (Neuringer et al., 2000). Threshold schedules have been shown to be effective in increasing the levels of variability in responses (e.g., Ross & Neuringer, 2002) or sequences of responses (e. g., Denney & Neuringer, 1998; Grunow & Neuringer, 2002); and the levels of variability can be controlled through changing the threshold value.

Statistical feedback method. Under statistical feedback schedules, response distributions are compared to those generated from random generator; feedback of how these responses resemble those calculated from different statistical descriptors are provided. Reinforcers are provided after responses are shown to be indistinguishable from those of random responses. In one study, when college students were asked to generate random sequences from 0 and 1 without receiving any feedback, the sequences generated were significantly different from the sequences of responses from a random generator (Neuringer, 1986). However, when these students were provided with feedback on how their responses assembled the 5 statistical descriptors in one experiment and 10 statistical descriptors in a second experiment based on random responses, they produced sequences that were indistinguishable from those generated by a random generator. A statistical feedback schedule was used in another study in which students were given feedback on how well their predictions of spots occurring on a horizontal line on the computer screen and on their predictions of the three digit sequences fitted the iterations from random generators. Responses from all but two of those students closely, but not perfectly, approximated responses generated randomly. Both of these studies showed that statistical feedback schedules could be used to shape variable sequences that are similar to random responses.

Summary

The schedules used to increase variability in responses depend on the nature of the behaviour being investigated, as well as whether randomness is required for the functional definition of variability. For studies that considered

highly variable responses to be similar to random responses, schedules used usually reinforced responses or sequences that were not frequently used and in a non-stereotypical way; sometimes multiple schedules had to be incorporated (e.g., using threshold contingency and lag schedules in Maes, 2003 and in da Silva Souza & Abreu-Rodrigues, 2010). For studies that did not consider randomness a requirement for a definition of variability, such as the creativity related studies and studies with individuals with autism, reinforcing novel or different responses and lag schedules with low lag requirements have been used successfully.

Measures of variability

While traditionally in psychology, variability is measured by descriptive statistics such as variances and standard deviations, studies of reinforcing variability typically use other types of measures. For the studies where variability is directly reinforced, the dimensions are usually not quantifiable in that measuring the distance of a response from the mean or median is not possible or meaningful, e.g., sequences of left and right key pecks. Studies investigating reinforcing behavioural variability have employed various methods to measure variability based on the nature of the dimension of the behaviour under investigation as well as the context in which the research was conducted.

The current section reviews the different measures of variability that have been employed in the studies of direct reinforcement of variability. Each of the methods is limited in measuring only one level or aspect of variability, however, behaviours can appear to be highly variable under the analysis of one measure while appearing to be highly stereotypical under the analysis of another measure that looks at a different level (Neuringer, 2002), therefore, more than one measure has been commonly used.

Percentage of trials meeting variability criteria for reinforcement. The percentage of trials meeting variability criteria is calculated by dividing the number of trials meeting reinforcement criteria by the total number of trials times 100. The reinforcement criteria usually require a relatively high level of variability through the reinforcement procedure (e.g., lag procedures and threshold procedures); thus having higher percentages of trials meeting variability criteria would mean having higher temporal variability. The percentage of trials

meeting reinforcement criteria is expected to be higher for conditions where reinforcement was contingent on variability compared to conditions where variability was not required. For example, Denney and Neuringer (1998) reinforced sequence variability in rats under threshold contingencies and compared the performance of these rats with those who received reinforcement independent of sequence variability; the percentage of trials with varying sequences were higher for rats in the variability contingent group than in the variability independent group towards the end of the experiment. The higher percentage of trials meeting variability contingencies showed that when under the variability contingency, the rats produced more less frequently used sequences. Ross and Neuringer (2002) reinforced one group of students to draw rectangles that varied on Sizes, Shapes and Locations while reinforced another group to draw rectangles without any variability requirement for any of the dimensions. The percentage of trials meeting variability requirement was higher for the group receiving contingent reinforcement for varying on all three dimensions compared to the group. Percentage of trails meeting variability criteria has also been used as a measure of variability in a number of other studies (e.g., Cohen et al., 1991; Maes, 2003; Mook, Jeffrey, & Neuringer, 1993; Neuringer & Huntley, 1991). This measure shows the temporary variability of responses under the control of reinforcement schedule and is often used in combination of other measures that shows other levels of variability.

U-values. When the measure of percentage of trials meeting variability criteria measures temporary (or molecular) variability, U-value has been considered to measure a more molar level of variability (Denney & Neuringer, 1998). Derived from information theory, U-value was adopted from Miller and Frick's (1949) measure of entropy to show the variability of overall response distributions to predetermined sets of response options (Denney & Neuringer, 1998). For example, if 300 responses were distributed over 16 bins, the number of responses made to each bin will be counted; a U-value will be calculated based on these counts. The resulting U-value will show the evenness of the response distribution over these 16 bins. U-value is calculated from the following formula (Ross & Neuringer, 2002):

$$U - value = - \sum_{i=1}^{\beta} \frac{\alpha_i \times \log(\alpha_i)}{\log(\beta)}$$

In the equation, α_i represents the relative frequency of response i , β represents the total number of possible responses (or available response options). U-values ranges from 0 to 1; a U-value of 0 represents extremely low variability while a value of 1 represents extremely high variability. Continued with the earlier example with the 300 responses, if equal number of responses were made to each of the 16 bins, the resulting U-value will be close to 1; if all responses were made to only 1 bin, the resulting U-value will be 0. If the 300 responses were to be distributed to the 16 bins randomly, a U-value close to 1 (but not equal to 1 because 300 responses cannot be completely evenly distributed over 16 bins) will result, as random responses are considered to be highly variable.

For the studies of direct reinforcement of behaviour variability, U-value has been one of the most commonly used measure (Da Silva Souza & Abreu-Rodrigues, 2010; Denney & Neuringer, 1998; Hopkinson & Neuringer, 2001; Maes, 2003; Murray & Healy, 2013; Grunow & Neuringer, 2002; Page & Neuringer, 1985; Neuringer & Huntley 1991; Paeye & Madelain, 2011; Ross & Neuringer, 2002; Stokes, 1999; Ward, Kynaston, Bailey & Odum, 2008). For example, in Denney and Neuringer's (1998) study, generating four response sequences on two levers with rats was reinforced; reinforcers were sometimes provided for highly variable sequences while they were sometimes provided without any variability requirement. U-values were calculated based on the relative frequencies of each of the 16 possible 4-response sequences over two levers. U-values as an indication of the level of variability were higher when the reinforcement was contingent on variability (U-value close to .9) than when it was independent of variability (U-value equalled around .6 and .7; Denney & Neuringer, 1998). U-values have also been used in measuring the variability in sequences generated by humans (Maes, 2003), in topographies in rats' bar press behaviour (Stokes, 1995) and in the sizes, shapes and locations of rectangles drawn by humans (Ross & Neuringer, 2002).

Studies that used U-value as a measure of variability have usually also used percent of trials meeting variability contingencies so that variability as measured at both a molar and a more temporary level. Although the variability

contingencies usually aimed at producing highly variable responses, stereotypical responding that meets the reinforcement criteria could result (e.g., alternating between two options in a Lag 1 procedure or cycling through options in a Lag 3 procedure and a threshold procedure). Therefore, having higher percentage of trials meeting variability criteria does not guarantee higher variability. Also, as U-value is calculated at the end of a session or from all responses within a defined period of time (e.g., last 300 responses in a session), it does not give information about the sequence of responding. For example, when reinforcement is provided for a child touching 4 coloured disks randomly under a Lag 2 schedule, s/he can simply cycle through all four disks to get reinforcement on all trials after the initial two. The percent of trials meeting variability criteria will be close to 100%. Additionally, because each disk would have been responded to equally, U-value will be close to or equal to one. Therefore, for this particular task, using the combination of two measures would not be sufficient for a full picture of variability in responding.

For some tasks when such stereotypy is not possible, such as for pigeons to cycle through all possible 8-response sequences to obtain reinforcers, checking whether they have cycled through the sequences may not be necessary. However, for those tasks that are at risk of stereotypical responding which can result in high number of reinforcers, alternative or additional measures that examine the sequences of responding should be sort after.

Sequential dependency measures. Measures of sequential variability in responses including autocorrelations, conditional probability, Markov Chain analysis and the frequencies of switching have been used in addition to percent of trials meeting variability criteria and/or U-values. Studies that use sequential dependence analysis methods assessed whether highly variable responses were similar to random responses where responses would have been distributed evenly in a non-stereotypical way.

One of the measures used is autocorrelations. Autocorrelation shows the correlation between current responses and a preceding response. If responses are made randomly, autocorrelations of any lengths (the number of preceding trials that is compared to with the current response) will be relatively low; if responses are made in a repetitive pattern (e.g., alternating between two options,

“12121212” or cycling through options, “123412341234”), autocorrelations will be relatively high. Autocorrelations of different numbers of preceding trials would be analysed so that repeated patterns of any length of sequences can be picked up. Autocorrelations have been used to measure whether responses had made in a fixed pattern that was not picked up by U-value (Maes, 2003). Maes (2003) compared the variability in 3-response-sequences between two groups of students, one received positive feedback contingent on varying the sequences used and the other received no feedback at all. U-values obtained from both groups were similarly high, an indication that both groups used the 27 possible sequences equally often. When the data were analysed using autocorrelations (analyses from Lag 1 through to Lag 27) , Maes’s (2003) results showed that some students generated relatively random responses as indicated by relatively low autocorrelations while some students had repeated long or short sequences to obtain reinforcers, indicated by the relatively high correlations.

Another set of measures of sequential variability used are conditional probability and Markov chains. Conditional probability of responses shows the probability of a response occurring given the occurrence of a previous response – (Machado, 1992; Stokes, 1995). If responses were randomly emitted, the probability of a response will be independent of any particular previous response; the probability of the current response given a specific previous response will therefore be unconditional. If the probability of current responses is higher given the occurrence of a particular responses previously emitted – the current response can be reliably predicted by a previous one– the conditional probability will deviate from the unconditional probability, which indicates there is a sequential dependency in the responding. For example, although the pigeons pecked equally often on left and right keys and they alternate between these two keys more during the condition where sequence variability was reinforced compared to a condition where sequence variability was not required, conditional probability analysis revealed that some pigeons showed a repetitive pattern of repeating the just emitted response or alternating between the two keys (Machado, 1992).

Both autocorrelations and conditional probability analysis examine the sequential dependency of only two trials; however, a repetitive sequence of more than two responses could occur (e.g., LLRLLRLLR). Markov chains analysis was

used for this type of sequential dependency in responding (Machado, 1992; 1993; 1997). Markov chains analysis detects possible repeating patterns of two or more responses; the number of responses a pattern consists of determines the order of the chain – first order refers to situations when the current response depends on the preceding one, and second order refers to situations when the current response depends on the preceding two consecutive responses and so on. Markov chain analysis revealed that the pigeons in Machado's (1992) study developed first-order and second-order stereotypy in their sequences to maximise reinforcement; however, third order stereotypy was not observed.

Higher order stereotypy identified by the sequential dependency measures showed that responses that met variability criteria for reinforcement were not necessarily as variable as random responses. Rather than suggesting variability cannot be trained, it points to the need to re-examine the reinforcement contingency as the higher order stereotypy observed was likely the result of the reinforcement contingency (Machado, 1992). Machado adjusted the reinforcement contingency to require variability in pairs of responses (1992) and in triplets of responses (1993), and the higher order stereotypy then disappeared. It is all about adjusting the reinforcement contingency based on the species' ability to learn to generate stereotypical response patterns rather than the desired response pattern, in this case, randomly, to obtain reinforcement.

Percentage of different responses. Percentage of different responses is calculated by dividing the number of unique responses within a given session by the total number of responses in that session. When the schedules that directly reinforce variability are permissive, such as a Lag 1 schedule, using only small number of different responses (e.g., two for a Lag 1 schedule) can result in high number of reinforcers obtained. Calculating the percentage of different responses will inform us whether these schedules produce only small number of different responses or whether they are sufficient in producing a varying pattern of results that result in more different responses than is needed for reinforcement. Page and Neuringer (1985) calculated the percentage of different 8-response sequences produced by pigeons in sessions with different lag requirements. Pigeons produced over 65% different sequences under a Lag 5 schedule; the percentages increased as the lag increased to 10, 15 and peaked at Lag 25. These results

suggested that pigeons produced more different sequences than the amount required for reinforcement. Other studies using this or a similar (number of different responses) method have generally found similar results; more different responses (sequences) usually resulted from direct reinforcement of variability (e.g., Machado, 1997; Morris, 1989; Stokes et al., 2008). The percentage of different responses (or sequences) is a useful measure to show the diversity of responses made.

The number of novel responses. This measure simply counts the number of unique responses throughout all sessions. Studies that reinforced behavioural variability by reinforcing novel responses within a session usually count the number of different responses within a session as well as the accumulative number of novel responses over different sessions. Recording the cumulative numbers of novel responses in addition to the number of different responses within a session is needed because repeatedly using same set of different responses in different sessions can result in high number of reinforcers under the schedule that reinforces novel responses within a session. The number of novel responses allows us to see whether the reinforcement contingency simply produces a small set of different responses or whether the set of different responses keeps expanding. The number of novel responses as a measure of variability is usually compared between a baseline condition in which no reinforcement is provided and an experimental condition in which reinforcement is contingent on producing different responses; the number is expected to be low and stable in baseline but higher and keeps increasing in experimental condition. When children were praised for constructing new forms in a block-building task, they produced increasingly more forms that were completely new to all previous sessions (Goetz & Baer, 1973). The same results were found in the training of diversity in the forms in easel painting; children painted more new forms during the period when novel forms within a session were reinforced (Goetz et al., 1973). Counting the number of novel responses as a measure of variability has also been used in studies with individuals with autism (Betz et al., 2011; Napolitano et al., 2010).

The number of novel responses is an appropriate measure of variability when the main purpose of increasing variability in responses is to expand the size

of the sets of responses. However, one should be careful as this measure is susceptible to ceiling effect, which could result from schedules such as reinforcing novel responses within a session; the number of novel responses stops increasing towards the end of the last sessions as individuals exhaust response options within their repertoires.

Other measures. The number of switches has been used to measure the variability in 8-response sequences generated by pigeons (Machado, 1997). In Machado's (1997) study, generating sequences over two keys that included at least one switch by pigeons in one experiment was reinforced and sequences including 3.5 switches was reinforced in another experiment. The average number of switches included in sequences generated by the pigeons in the experiment that reinforced at least one switch was found to be higher than that predicted by a random model; whereas the average number of switches for some of the pigeons in the other experiment where sequences including 3.5 switches were reinforced most often was close to those predicted by a random model. Although some of the birds were found to generate sequences similar to the random model, further analysis of the sequences of their responses revealed strong preferences for particular keys in some pigeons; response patterns that were different from random responses. Another method used to measure variability is to compare the differences between observed data and random data (Neuringer, 1986); however, this method is not commonly used.

Besides the measures used in the study of directly reinforcing variability mentioned above, there are other measures used in studies that did not directly reinforce variability which can also be useful. For example, in a study that examined the variability in locations of pigeon's pecking on a screen under conditions with different probability of reinforcement, Stahlman, Young and Blaisdell (2010) and Stahlman, Roberts and Blaisdell (2010) calculated the distance between the median location of a pigeon's pecks and the location of each individual peck from the same pigeon. The greater the distances from each peck to the median, the higher the variability in the pecking locations were. This task and measure may help to move the experimental studies of reinforcing variability from focusing on sequence learning to include wider dimensions, such as enhancing variability in spatial learning. Other measures used include bar-press

duration in rats (Gharib, Derby, & Roberts, 2001; Gharib, Gade, & Roberts, 2004; Leising, Ruprecht, & Stahlman, 2014) and response rates (Leising et al., 2014).

In terms of the measure of variability, there is no one perfect measure. The types of measures used depend on the definition of variability the researcher took on. For example, when highly variable responses is considered to be similar to random responses (e.g., Denney & Neuringer, 1998; Maes, 2003), measures that show how well the responses meet the variability criteria (e.g., percent of trials meeting reinforcement criteria when the criteria was based on high level of variability), how evenly the distributions of the responses are over available options (e.g., U-values), and whether these responses are made in stereotypical ways (e.g., autocorrelations) are needed. On the other hand, if randomness is not necessary in defining variable responding (e.g., use of colours; drawing of new forms); counting the number of different/new forms can be sufficient.

Summary

Each method described above is limited as they only show variability at only one level. Most of the studies mentioned above used a combination of different measures that examined variability at different levels. Some studies using different measures found similar variability at different levels (e.g., Neuringer & Huntley, 2001; Paeye & Madelain, 2011) while some studies found variability of the same response to be different at different levels (e.g., da Silva Souza & Abreu-Rodrigues, 2010; Maes, 2003; Machado, 1992; 1993; 1997). Most of these studies used a task related to sequence learning for which randomness was required to be considered highly variable; for these studies, the use of multiple measures of variability is important. For the study of variability in other dimensions such as variability in the colours used in drawing, choices of activities at schools, or variable functional verbal responding in individuals with autism, when expanding the response diversity is more important than making these different responses randomly, the sequential dependence methods (e.g., autocorrelations and Markov chains) might not be necessary nor appropriate.

Identifying the gap

The control of contingencies on variability has been consistently demonstrated in both highly controlled experimental settings as well as in naturalistic and applied settings; and in various species across a wide range of tasks using different reinforcement schedules and by various measures. Most of the studies investigating the impact of reinforcement on behavioural variability looked at only the variability of one dimension of a response on which the reinforcement was contingent, sequence is one example. However, in reality, behaviour is multidimensional and more than one dimension of a behaviour can occur simultaneously. For example, in an experiment when a rat is pressing a bar, we may measure the rate of responding, which provides one dimension, yet force and duration are also available simultaneous dimensions on which the response may vary. Thus, for each response, there are at least three simultaneous dimensions (rate, force and duration). Skinner (1953) suggested that the frequencies of responses that have not been directly reinforced increase as the result of directly reinforcing another. It is not clear to which degree reinforcing variability in one dimension affects variability in the others. More specifically, would reinforcing variability in one dimension increase the variability of another dimension of which the variability was not directly reinforced?

Generalisation

Generalisation is said to have occurred if the frequencies of a response increase under untrained stimuli as the result of that response being reinforced under a stimulus which shares some common features of those untrained stimuli (Skinner, 1953). Generalisation is traditionally demonstrated by the generalisation gradients; the frequencies of responses would be highest to the trained stimulus and decrease as the dissimilarity (or spatial space) of untrained stimuli increases (Brown, Bilodeau, & Baron, 1951; Skinner, 1953). According to Skinner (1953), because of the multidimensional nature of any response, to establish strong discriminative stimulus control over responses, one has to vary properties of other dimensions while the property of the target dimension is remained unchanged. That is to say, the control of a stimulus over a response, although targeting on only one dimension of that response, could end up having similar effects over other non-target dimensions.

Despite its theoretical significance, generalisation is often (if not more) sought after in applied settings. When a behaviour is trained in a well-controlled experimental setting, it is desirable that it will also occur in the more natural environment when reinforcement from the original training setting is unavailable or is available to a less extent. Generalisation is said to be occurring if a trained behaviour occurs when the reinforcement used during training is no longer available or available only to a lesser extent (Stokes & Baer, 1977). Generalisation of the training effect can be assessed across behaviours, individuals or contexts (Stokes & Baer, 1977). In applied settings, generalisation is usually considered to be a criterion for successful outcomes for interventions individuals with various disabilities (Stokes & Baer, 1977; Stokes & Osnes, 1989), including individuals with autism (Arnold-Saritepe, Phillips, Mudford, De Rozario, & Taylor, 2009). Therefore, investigation of the generalisation of the control of a stimulus over responses across stimuli, contexts and individuals has both theoretical and applied significance.

Behavioural variability has been shown to be under the control of contingencies and discriminative stimuli, just as other operant dimensions; whether trained variability will generalise is therefore also of interest. One of the suggestions on the methods of promoting generalisation by both Stokes and Baer

(1977) and Stokes and Osnes (1989) was to train diverse responses that would lead to the increased contact of reinforcement in non-experimental environments. Reinforcing individuals to vary among the acquired responses can therefore be seen as an extended step to this training of diverse responses. The training of variability does not usually precede the training of the prerequisite behaviours or responses, which was referred to as “training diverse exemplars” in Stokes and Baer (1977). Schmit and Bjork (1992) suggested that training variable responses to one dimension of a behaviour facilitated the generalisation of such behaviour in novel contexts. In regard to generalisation of learned variability, the questions being asked would be: after being reinforced to vary responses along one dimension of a behaviour, will the individual vary their responses on other dimensions of the same behaviour? Or after being reinforced to vary responses in one behaviour, will the individuals vary their responses in other related behaviours? Or after being reinforced to respond variably in one context, will the individuals respond more variably in another context?

Despite the theoretical and applied significance of the generalisation of learned variability it has not been extensively investigated. In order to identify research that have investigated generalisation, possible tests that show the existence of generalisation are proposed based on the definition of generalisation from Skinner (1953) and Stokes and Baer (1977). The term generalisation has been found to be used inconsistently across researchers (Edelstein, 1989), therefore, it is necessary to clarify the definition used in this thesis. Traditionally, when generalisation was assessed using generalisation gradients, frequencies of responses to stimuli that were around the target stimulus were examined; responses to these stimuli had not been directly reinforced (Skinner, 1953). Therefore, there should be a difference in the reinforcement contingencies between the behaviour targeted for training and for the behaviours which generalisation of the training that is to be assessed later. Specifically, when reinforcement is provided for the target response or the target dimension of a response, the same amount of reinforcement should not have been provided for the responses or dimensions of responses for which generalisation are to be assessed. In some instances, the difference could be having a lessened amount of

reinforcement in the condition where generalisation is assessed (Stokes & Baer, 1977).

Generalisation of learned variability

Bearing the above mentioned criteria for a successful demonstration of generalisation of trained behaviour in mind, the study of generalisation of learned variability should have the following characteristics. First, relatively high level of variability should have been established for one behaviour or for at least one dimension of a behaviour by directly reinforcing variability. Second, variability of the to-be-generalised behaviours or dimensions of a behaviour or of the same behaviour in the to-be-generalised contexts should not have been directly reinforced in related contexts. Third, the changes in the level of variability, if any, in the variability of the to-be-generalised behaviour or dimensions of behaviour or in the to-be-generalised contexts should be assessed. Finally, generalisation of learned variability in the to-be-generalised behaviour or dimensions of a behaviour or in the to-be-generalised contexts is demonstrated if there is increase in the variability assessed after variability is directly reinforced in another behaviour or in other dimensions of a same behaviour or in other contexts.

There have been studies investigating whether learned variability would generalise across behaviours, individuals or contexts. In one study, Stokes et al. (2008) trained participants to move a light box from the top of a pyramid to different endpoints at its bottom. Three sizes of pyramids were used (5-pyramid, 7-pyramid and 10-pyramid); the size of the pyramid reflected the number of responses required to reach its bottom endpoints. Each of the endpoints could be reached in a number of possible paths. In their last experiment, variability in the paths taken to reach designated endpoints of a 10-pyramid was required for reinforcement for one group of participants while any paths taken to reach the designated endpoints were reinforced, regardless of whether participants varied the paths or not. Participants from the group in which variability was required used more different paths than participants from the other group during training. All participants were then tested after training on a different pyramid (7-pyramid) and on endpoints that had never been trained on the 10-pyramid; both groups received reinforcement for using any paths to designated endpoints. The group

who had been reinforced for varying remained more variable in the paths use than the other group in the new pyramid and to the new endpoints (Stokes, et. al., 2008). These results show that after experiencing the reinforcement contingency that required variability, the trained level of variability generalised to a different (but similar) task and to new elements of the same task.

Evidence of generalisation of learned variability also comes from studies in the exploratory behaviour in rats. Long-Evans strain (LEs) rats are considered bold as they would spend relatively long time staying in the centre of a space and rearing more to inspect their surrounding environment while PVG strain (PVGs) rats are considered shy as they tend to spend little time in the centre of a space and rear little (Weiss & Neuringer, 2012). Weiss and Neuringer (2012) trained two groups of rats, one of the LEs and the other PVGs to vary their ways of interactions with novel objects to see whether they responded to the contingency similarly. For each of the experimental groups which received food pellets for varying their responding, another group of rats of the same strain - sisters of the experimental groups - received food pellets yoked to their sisters from the experimental condition, independent of response variability. Both the LEs (bold) and the PVGs (shy) rats reinforced for varying their responses spent more time interacting with the objects and used more variable ways, ways that were less commonly used by other rats interacting with the objects than their sisters in the other condition in which variability was not required for reinforcement. The level of variation in the responses from the genetically shy PVGs rats came to reach to a level that was the same as the genetically bold LEs rats.

All the rats, two groups of LEs and two groups of PVGs from the previous experiment and two new groups of experimentally naïve PVGs rats, were later tested for generalisation of the trained level of variation in responding. The rats were put into a completely different environment with a set of new objects they had never encountered in their life. Food pellets were hidden in each of these objects; the more the rats vary their interaction with the objects, the more likely they would find the food pellets. The rats who received reinforcers for variable interactions during training successfully found and consumed more food pellets in the new environment hidden in new objects than the number found by the other groups; the experimentally naïve rats found the least number of food pellets. The

bold LEs rats found and consumed more food pellets than the shy PVGs rats in all pairing groups (e.g., Bold reinforced – Shy reinforced, Bold independent – Shy independent, and Bold naïve – Shy naïve). Results also showed that the time spent interacting with the new objects was longest for the rats who had been reinforced for varying their responses than the other two groups; and the naïve group spent the least time interacting with these objects. The strain of the rats did not affect the interaction times between experimental conditions. These findings were exciting as they demonstrated that the increased level of variability from training would generalise to new environments with new objects; absence of the training resulted in less exploration in the new environment and less different ways interacting with new objects which in turn lower the probability of obtaining reinforcers. The test of generalisation to some extent parallels natural environments (especially in situations involve problem solving) where reinforcers were not obvious, one has to explore and try different methods to obtain them. The finding that variability in responding in one environment can generalise to new environments which in turn enhances the probability of being reinforced would have significant implication for intervention for individuals with depression.

There is evidence of possible generalisation from two studies that did not directly investigate generalisation. Maes (2003) studied the effect of contingent and non-contingent reinforcement on response variability in three response sequences over three computer keyboard. One of their experiments included two groups of participants; one group first received contingent reinforcement on variability of the sequences (VAR1) followed by reinforcement that was yoked to the number and distribution of reinforcers received previously in VAR1 but independent of variability of current response (YOKE1). Participants from the other group first received reinforcement that was yoked to the VAR condition of the first group, but independent of the variability of the current sequences (YOKE2) followed by reinforcement contingent on the variability of the sequences (VAR2). It was found that the percent of sequences that would have met the variability requirement showed an increasing trend for YOKE1 when this condition followed the variability required condition (VAR1); however, the percent of sequences showed a decreasing trend for YOKE2 when this condition

preceded the condition where variability was required (VAR2). The results suggested that after students learned to vary under contingent reinforcement on variability, their tendency to vary persisted later when the reinforcement was no longer contingent. Students who had no previous experience of being reinforced to vary, non-contingent reinforcement did not have the same effect.

Possible generalisation was also found in Da Silva Souza and Abreu-Rodrigue's (2010) study using a task similar to Maes' (2003) three response sequence generation task. In one phase of a condition in Da Silva Souza and Abreu-Rodrigue's (2010) second experiment, participants received reinforcers probabilistically (50%) and independently of the variability of the sequences. Performances of two groups of participants were compared. One group experienced the variability independent phase after a phase in which variability of the sequences was required for reinforcement. The other group experienced the variability independent phase prior to the variability phase. They found that participants produced more variable sequences in the variability independent phase if this phase was preceded by a phase in which variability was reinforced. Therefore, similar to results found in Maes (2003), having experienced contingent reinforcement on varying the sequences had positively influenced the variability of sequences even when it was not required for reinforcement. Results from both the studies showed that the acquired level of variability from contingent reinforcement could have generalised to a condition where variability was not directly reinforced.

Generalisation of trained level of variability has also been investigated in applied settings with individuals with autism. Lee et al. (2002) trained three individuals with autism to use more variable and appropriate responses to questions from adults. The number of appropriate responses started low for all three individuals and it increased greatly for two of them during intervention when a Lag 1 schedule was used. There was also an increased number of novel verbal statements observed for these two individuals. When later the reinforcement was no longer in effect, the same two individuals continued to use more variable and appropriate responses to new therapists in the same setting or to the same therapist in a different setting. Lee et al.'s results (2002) demonstrated

that the level of variability in functional verbal responding in individuals with autism acquired under training could generalise across people and settings.

Miller (2012) examined whether variability in shapes built with wood blocks or patterns created using pegs on a foam board trained under a Lag 3 schedule would generalise to variability in the sequence of colours used in painting in two squares on a paper. Participants were first reinforced for building/making any shapes or patterns in baseline, with the instruction “Build/Make something”. They were then reinforced to build/make different shapes and patterns under a Lag 3 schedule, with the instruction “Build/Make something different”. After that, they were reinforced to repeat any shapes or patterns built or made in the previous three trials (Repeat 3). Another Lag 3 and Repeat 3 schedules (same order) followed once for two participants and twice for the third. Then the reinforcement contingencies alternated between Lag 3 and Repeat 3 for all participants for three following sessions. Clear stimulus control of varying or repeating was shown in the final alternating session in which the shapes or patterns were more variable in the Lag 3 component than in the Repeat 3 component. Whether this stimulus control would generalise to the painting task was later assessed. During the generalisation phase, participants were trained for the painting task under the same alternating schedules as in the block-building task. Results showed that the variability in the sequences of colours used was similarly and relatively low for two participants; for the third participant, variability in the Lag 3 component was only slightly higher than the Repeat 3 component.

Based on these results, Miller (2012) suggested that the stimulus control of the variability contingency and the repeat contingency established for the block-building task did not generalise to the new painting task. The instruction used in the generalisation phase was “Make something different” for the VAR component and “Make the same” for the Repeat 3 component. When “Build/Make” the same or different was more straightforward as what to do when shapes and patterns were assessed, “Make” the same or different might not be as clear in terms of what the behaviours were expected in the painting task. The lack of generalisation of the stimulus control could have been due to the lack of appropriate discriminative stimuli (S^d) used. If the instructions (S^d) were “Paint” the same or

different, could there be more stimulus control from the Lag 3 and Repeat 3 contingencies?

Furthermore, participants had training for the Lag 3 and Repeat 3 contingencies before the phase with alternating schedules started. Could the stimulus control observed in the first alternating session for the block-building task be weaker if participants had not had the previous training? By including a baseline alternating session for both the tasks before training would allow a more accurate assessment of generalisation across task. Assessing the changes in the variability in the painting task before and after experiencing the reinforcement contingency in the block-building task or shape making with pegs task would give a clearer idea whether there was generalisation. An example of the procedure would be having baseline for block-building, making shapes with pegs and painting, and immediately follow an alternating session for all three tasks before the above training procedures for the block-building/shape making with pegs task; and finally test the generalisation to the painting task under an alternating session. If the stimulus control for varying and repeating was stronger for the alternating session before training for the painting task in Generalisation phase than it was for the alternating session before training for the other tasks, there would be generalisation. Therefore, the claim that generalisation of trained variability and the stimulus control did not generalise to a new task was not conclusive because of the limitation in the instructions and experimental design.

Generalisation of learned variability has also investigated in the creativity studies; mixed results were found. In one study, two children received descriptive social reinforcement for painting different forms in each paintings while they received no reinforcement for a block-building task which occurred in the same experimental session as the painting task (Goetz et al., 1973). The number of different forms in the paintings increased when the reinforcement contingency was in place for both of the children; novel forms were more likely to be observed during these sessions. Variability of the forms used in the block-building task was never reinforced; however, the number of different forms built in each session increased across the sessions where variability of shapes in the painting was reinforced (1973). Training variability in responding in one task had increased the

variability in another task that occurred in the same session; thus generalisation was observed.

While generalisation of learned variability across tasks was said to have occurred in Goetz et al.'s (1973) study, contradictory results were found in another study. Holman, Goetz, and Baer (1977) reinforced the form diversity in the shapes of felt-pen drawing by three children and then assess the form diversity of three other tasks completed on the same day but those received no reinforcement. These three tasks included one topographically similar easel drawing and two topographically different tasks, wooden block and Lego™ building. For felt-pen drawing, where form diversity was directly reinforced, the number of different forms increased when the reinforcement contingency was in place for all three children; so did the number of new forms for two children. For easel painting, there was a clear increase in the number of different forms observed for one child, a reasonable increase for another, but little increase for the third. The number of new shapes accumulated over sessions also increased for these children. Therefore, there seemed to be generalisation of trained variability from the felt-pen drawing to easel painting. For block-building and Lego™ building, there was no increase in the number of different forms built within a session nor in the number of new forms accumulated over sessions. Generalisation of the trained variability seemed to have occurred for tasks that are topographically similar but not to topographically different ones.

This correlation between generalisation and topographical similarity has been observed in another study. Goode, Geraci, and Roediger III (2008) trained three groups of college students to solve word anagrams. Two of the groups were trained using only one of four anagrams of a word while the third group was trained using three of the four anagrams of the same word. All participants were tested immediately and two days after the training on half of the studied words and the same number of unstudied words. One group was tested using the same anagrams as used in training, a second group was tested using a different anagram to the one used in training, and the third group was tested using the fourth anagram that was not used in training. All three groups showed increased success rate in the immediate test session than in their first practice session; the group that was trained using different anagrams had the largest increase compared to the

other two groups. However, there was no difference among the three groups in success rate of solving anagrams of unstudied words. Training using different anagrams appeared to have facilitated solving new anagrams of the words already studied but not to solving anagrams of new words.

Conclusion and Aim

Generalization of learned variability to date have been investigated both in experimental and applied settings across tasks and across contexts. In addition to the lack of consistent results for generalisation of learned variability across tasks in the creativity related studies (e.g., Goetz, et. al., 1973; Holman, et. al., 1977) and in the studies with individuals with autism (Lee, et. al., 2002; Miller, 2012), the results from these studies were also at risk of having been confounded by other variables. For example, Holman et al. (1977) commented that making an immediate judgement on whether a form built was different and whether it qualify for reinforcement was difficult and could result in errors. The fact that to determine whether a response qualify a reinforcer is relatively subjective and takes time; the immediacy of delivery of reinforcement could also have been compromised. Studies in the more controlled experimental environments (e.g., Arneson, 2000; Stokes et al., 2008; Weiss & Neuringer, 2012) as well as studies that did not directly investigated generalisation of variability (Da Silva Souza & Abreu-Rodrigues, 2010; Maes, 2003) show more promising results.

Though generalisation across behaviours and/or settings is the ultimate goal, it may be more rigorous to investigate such effect within setting in a more controlled experimental environment. Clear demonstration of generalisation in controlled experimental environments with minimum confounding variables will be a first step. Generalisation across behaviours and contexts are important, as well as generalisation across dimensions of the same behaviour. The present thesis attempts to investigate generalisation of learned variability across dimensions of a behaviour by firstly using a modified version of the rectangle drawing task mentioned earlier. Once confirmation of generalisation of reinforced variability was obtained, the investigation would extend to examine whether same results would be obtained when dimensions of behaviour that were more closely related to everyday living were used.

Experiment 1

Introduction

To study whether learned variability generalises across dimensions of a task requires a task with two characteristics. First, the task needs to provide readily measurable multiple dimensions that occur simultaneously, with this it would be possible to differentially reinforce variability in each dimension. Second, the task should allow for accurate and immediate delivery of reinforcers (e.g., more objective and standardised ways of determining whether a response qualifies for reinforcement). Such a task was devised by Ross and Neuringer (2002) who used a computerised rectangle drawing task to examine the effect of reinforcement across multiple dimensions of the response. Each response created a rectangle that had three measureable dimensions: Area, Shape and Location on the screen. In one experiment, one group of participants received reinforcement only when each response varied on all three dimensions; that is, they needed to create rectangles of different sizes, shapes and at different locations. Another group received the same amount of reinforcement for drawing rectangles without the requirement of varying any of the dimensions.

Ross and Neuringer (2002) found that the sizes, shapes and locations of the rectangles were more variable for the group when reinforcement was contingent on variability on all dimensions than when it was not. In a second experiment in their study (Ross & Neuringer, 2002), a reinforcer was provided when participants varied their responses on two dimensions and repeated on the third. Variability was shown to be high for the dimensions that were required to vary and low for the dimension that was required to repeat. Ross and Neuringer (2002) demonstrated that reinforcement can influence multiple dimensions of a response simultaneously; and that it can also be used to produce different levels of variability over each of these simultaneous dimensions. Although not investigated by Ross and Neuringer (2002), the rectangle drawing task provides an opportunity to examine whether reinforcing variability on some dimensions would influence the variability of another dimension of which variability was not directly reinforced.

This first experiment was a direct replication of Ross and Neuringer's (2002) first experiment in which they reinforced participants to vary on three dimensions (Area, Shape and Location) of rectangles drawn on a computer screen. Success in replicating their results would provide a methodology to study variability further as well as providing data for the examination of the generalisation of learned variability across dimensions.

Methods

Participants

Forty adults (23 females and 17 males), aged between 17 and 42 years, participated in this experiment. They were recruited by advertisements posted on notice boards around the university and online. The first 20 participants recruited were assigned to the experimental group; and the next 20 were assigned to the control group. Participants who were first year psychology students were given course credit.

Apparatus

A PC computer with a 20-in monitor, a keyboard and a mouse were used to present instructions and record participants' responses.

Procedures

Establishing reinforcement criteria

A computerised rectangle drawing task was created based on Ross and Neuringer's (2002) study. Three simultaneous dimensions – Area, Shape and Location – of the rectangles drawn were measured and monitored for variability. Area was defined by height times width; Shape was defined by the height to width ratio; and Location was defined by the (x, y) position of the centroids of the rectangles on the computer screen.

Creating 16 categories

A computer was programmed to randomly generate 500,000 rectangles within the confines of the computer screen; four measures were recorded for each rectangle: the sizes (Area), the ratio of width to height (Shape) and the x and the y coordinates of the centroid (Location). These four measures were then sorted independently of each other in an ascending order. Then, two of the lists, for Area and Shape, were divided into 16 sets with equal number of values ($31,250 * 16 =$

500,000); while the rest of the two lists, x and y values, were divided into four sets of equal number of values. For Area and Shape dimensions, the upper and lower bounds of each set for these two dimensions defined the 16 categories used for assessing variability, which will be described later. Using this method, combined with x and y data, 16 categories with similar sized regions on the screen were defined.

Reinforcement of Variability (VAR)

In order to encourage variability across responses, a threshold contingency was used. Threshold contingency reinforces responses that had occurred relatively less frequently in the recent past. For the rectangle drawing task, this means that reinforcement would be given to categories that had been used less frequently in the recent past. To determine how frequently categories had been used, relative frequencies were calculated for all categories for each dimension after each response. Relative frequencies were calculated by dividing the number of occurrences of a category by the total number of trials up to the current trial. For example, if Category 1 for the area dimension had been responded to 18 times over 300 trials, the relative frequency would be .06. When a rectangle was drawn, three categories, one from each dimension would result. The relative frequencies of these categories from the immediately preceding trial would be compared to a threshold value to assess whether the rectangle drawn on the current trial qualified for reinforcement.

For the experimental group where variability on all three dimensions, Area, Shape and Location was desired (VAR), participants were awarded a point only when the relative frequencies of the categories for all three dimensions were lower than a threshold value. An exponential weighting function was applied to trials that received points by multiplying the relative frequencies by .99 so that recent categories were weighted more than earlier ones. A detailed description and an example for the calculation of relative frequencies and the weighting procedure can be found in Denney and Neuringer (1998) and Ross and Neuringer (2002). These weighted relative frequencies were used to assess whether the categories used in a response were less frequent than the level specified by a threshold value. Ross and Neuringer (2002) suggested that responses of optimal variability would be similar to those seen in random responses. Thus the threshold value was

calculated based on what a random generator would do when drawing 300 rectangles. If the responses were evenly distributed across all 16 categories, the resulting threshold value would be .0625 ($(300/16)/300$). However, a threshold value of .0825 was chosen instead by Ross and Neuringer who found that this level of difficulty for this multi-dimensional task was appropriate after piloting different threshold values with humans. In the present experiment, we used the same threshold value as Ross and Neuringer (2002), thus for the VAR group, in order to get reinforcement (points), participants had to draw rectangles of which the weighted relative frequencies (WRF) of categories for Area, Shape and Location dimensions were lower than .0825. In other words, participants had to vary all three dimensions of the rectangles to get reinforcement.

Reinforcement Independent of Variability (YOKE)

A yoking procedure was used to guarantee participants in the control group (YOKE) received the same amount of reinforcement, at the same time as the experimental group. Each participant in the YOKE group was paired with a participant in the VAR group. The trials on which points were delivered to a participant in the YOKE group were identical to that of their assigned partner in the VAR group. For the participants from the YOKE group, points were delivered independent of the types of rectangles they drew. For example, if responses made by a participant from the VAR group met reinforcement criteria on the sixth, eleventh, thirteenth ... n^{th} trials, the participant from the YOKE group whom was paired with this particular participant from the VAR group would receive reinforcers on the sixth, eleventh, and thirteenth ... n^{th} trial.

Experimental procedures

Sessions were conducted with one participant at a time. The experimenter greeted the participants and provided them with the information sheet and consent form. Participants were given time to read the information sheet and ask questions before the experimenter read the consent form to them. After signing the consent forms, participants proceeded to the computer on which the instructions were displayed. The instructions displayed were as follows:

“To play, simply click the mouse and drag on any diagonal to create a rectangle. Release the mouse button when you are satisfied with your

rectangle. The object of this game is to get the most points. You have received points for your actions whenever you hear the ascending tones. There will be 300 trials. It will take approximately 15 min. Feel free to stop the experiment at any time. Please ask the researcher if you have any questions. Click anywhere on the screen to start.”

Light grey rectangles appeared on a dark grey background when the participants dragged the mouse diagonally to move the cursor on the screen. The experimenter left the room after the participants made two or three correct responses. Ascending tones (100-ms 1500-Hz followed by 100-ms 2000-Hz) occurred immediately after rectangles that met the reinforcement criteria were drawn. Every rectangle drawn remained on the screen for 500 ms and mouse clicks during this period had no effect; the screen cleared after 500 ms and the next trial started. Participants were debriefed at the end of the experiment.

Measures

Number of trials meeting variability criteria

For both groups, the number of trials that met the variability criteria were counted. For the VAR group, the number of trials that met the variability criteria equalled the number of points they obtained in the session. For the YOKE group, although the points were not given based on the variability of their responses, the computer recorded whether the responses would have met the variability criteria if they were in effect.

U-values

U-value is a measure of overall variability across all responses without regard to the reinforcement contingency. U-value was calculated according to the following formula:

$$U - value = - \sum_{i=1}^{\beta} \frac{\alpha_i \times \log(\alpha_i)}{\log(\beta)}$$

In the formula, β equals the number of possible categories (16 in this case) and α equals the relative frequency of category i . The U-value shows how equally the responses were distributed over the 16 categories. U-values range from 0 to 1, where a value of 0 indicates strict repetition and 1 indicates responses are equally

distributed to all categories. By way of an example, if one examines the domain of Area, if a participant drew 300 rectangles of Category 4, the U-value for area dimension would be 0; if the participant drew about 18 rectangles in each of the 16 categories, the U-value will be 1. Higher U-values would reflect greater variability in the responses. U-values for the three dimensions were calculated separately.

Results and Discussion

Number of trials meeting variability criteria

The number of trials meeting variability criteria was measured for both groups under the different reinforcement contingencies. A one-way ANOVA across the two groups on the number of trials that met the reinforcement criteria showed the number was significantly higher for the VAR group ($M = 70.6$, $SD = 25.85$) than it was for the YOKE group ($M = 27.25$, $SD = 26.25$), $F(1, 39) = 27.69$, $p < .001$, partial $\eta^2 = .65$. This finding is consistent with Ross and Neuringer's (2002) results (VAR = 62.4, YOKE = 31.7, $F(1, 39) = 21.55$, $p < .01$).

The percentages of participants in each group meeting the variability contingency on each of the 300 trials are shown in Figure 1.1. The solid lines have been smoothed using a moving average across 30 trials. Dotted lines represent the linear regression for each group. On the first trial all participants met the criteria to earn a point because the relative frequencies of all categories were zero at the start of the experiment. As they completed more trials, the percentage of participants meeting the criteria decreased markedly before increasing after more and more trials were completed for VAR group only. As shown in Figure 1.1, participants in the VAR group increasingly met the reinforcement contingency (varying their responses on all three dimensions) as the session progressed, while the percentage remained unchanged for the YOKE group. This is also consistent with Ross and Neuringer's (2002) results.

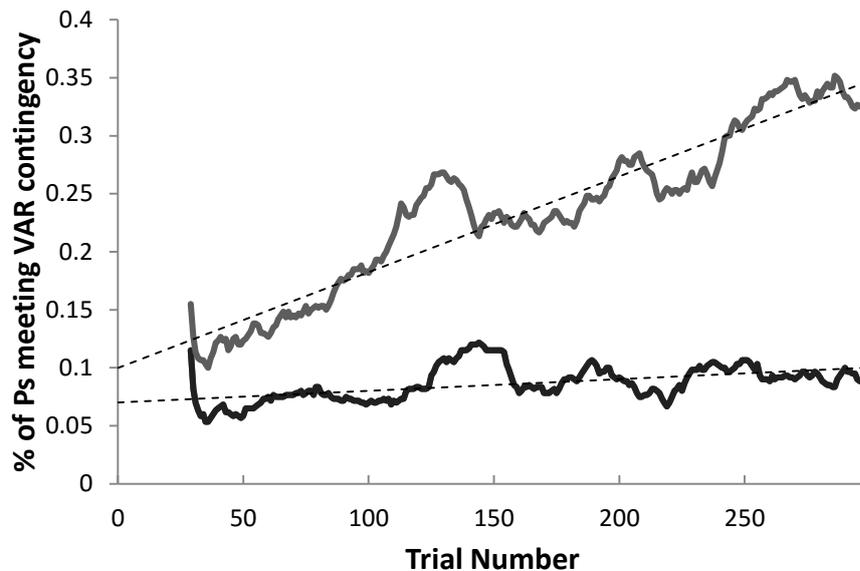


Figure 1.1. The moving average across 30 trial blocks of the percentage of participants in the VAR and YOKE groups that met the variability criteria during Experiment 1. The solid lines are the moving average and the dotted lines are lines fitted by the method of least squares.

U-values

U-values were calculated separately for each of the three dimensions based on the 300 responses for each participant. A mixed design repeated measures one-way ANOVA of the U-values revealed a significant main effect between the VAR and YOKE group, $F(1, 38) = 27.88, p < .001, \text{partial } \eta^2 = .423$. Mean U-values for each dimension for both groups were presented in Figure 1.2. Post-hoc t-tests were used to compare the mean U-values for each dimension between the VAR and YOKE groups. Results showed that mean U-value for the Area was significantly higher for the VAR group (.95, $SD = .047$) than it was for the YOKE group (.83, $SD = .146$), $t(38) = 3.57, p < .001$; mean U-value for Shape was significantly higher for the VAR group (.92, $SD = .088$) than it was for the YOKE group (.72, $SD = .181$), $t(38) = 4.39, p = .012$; and mean U-value for Location was significantly higher for the VAR group (.95, $SD = .33$) than it was for the YOKE group (.78, $SD = .153$), $t(38) = 4.939, p < .001$. Standard errors appeared to be much smaller in the VAR group than they were in the YOKE group.

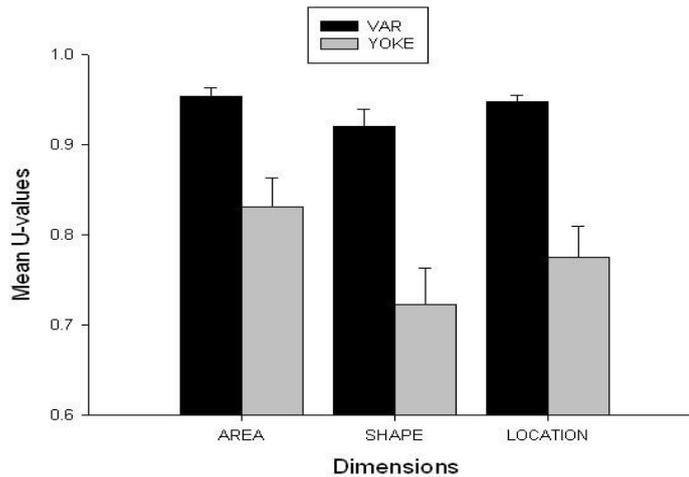


Figure 1.2. Mean U-values for each dimension for the VAR and YOKE group. High U-values indicate high variability. Error bars show standard errors. Values below .6 have been omitted from the y-axis.

These results were consistent with Ross and Neuringer's (2002) experiment in which mean U-values were found to be higher in the VAR group than the YOKE group for all three dimensions; and standard errors to be smaller for the VAR group than that of the YOKE group. The differences between VAR and YOKE groups for each of the three dimensions observed in the current study appeared to be bigger than those reported in Ross and Neuringer's (2002); and there seemed to be greater differences between mean U-values among the three dimensions for the YOKE group in the current study (between .12 and .21) than in Ross and Neuringer's (2002) study (within .01). In addition, a repeated measures ANOVA revealed that there was a significant difference in the U-values among the three dimension for the YOKE group, $F(2,38) = 4.55, p = .017, \text{partial } \eta^2 = .193$. Specifically, pairwise comparisons among the three dimensions revealed that U-values for Area were significantly higher than for Shape. This suggested that participants from the YOKE group had varied on the sizes (Area) more than on the shapes of the rectangles. No significant difference was found among the three dimensions for the VAR group.

The significant difference found between dimensions for the YOKE group suggests that the reinforcement delivered for this group, although independent of the variability of any dimensions, had impacted on the variability on the three dimensions differently. Specifically, there might be a tendency for participants to

vary the sizes (Area) of the rectangles but not to vary on shapes or locations. Figure 1.3 shows the scores and U-values for all participants in the two groups for each dimension. As can be seen, data shown in Figure 1.3 agree with what was seen from the mean U-values in Figure 1.2; most of the U-values for the VAR group were higher than those from the YOKE group for all three dimensions. The differences in U-values between the VAR and YOKE pairs appeared to differ between Area dimensions and Shape and Location dimensions for pairs of VAR-YOKE participants who scored over 80. Participants from the YOKE group appeared to have varied more on the sizes (Area dimension) only when they received higher number of points; even though these points were delivered independently of the variability in the sizes (Area); participants may have superstitiously attributed the reinforcement to varying on sizes. The fact that the same did not happen to Shape and Location dimension may suggest that when the reinforcement contingency was unclear to the participants, they tended to vary on the Area dimension rather than Shape and Location dimensions – showing a response bias to Area dimension.

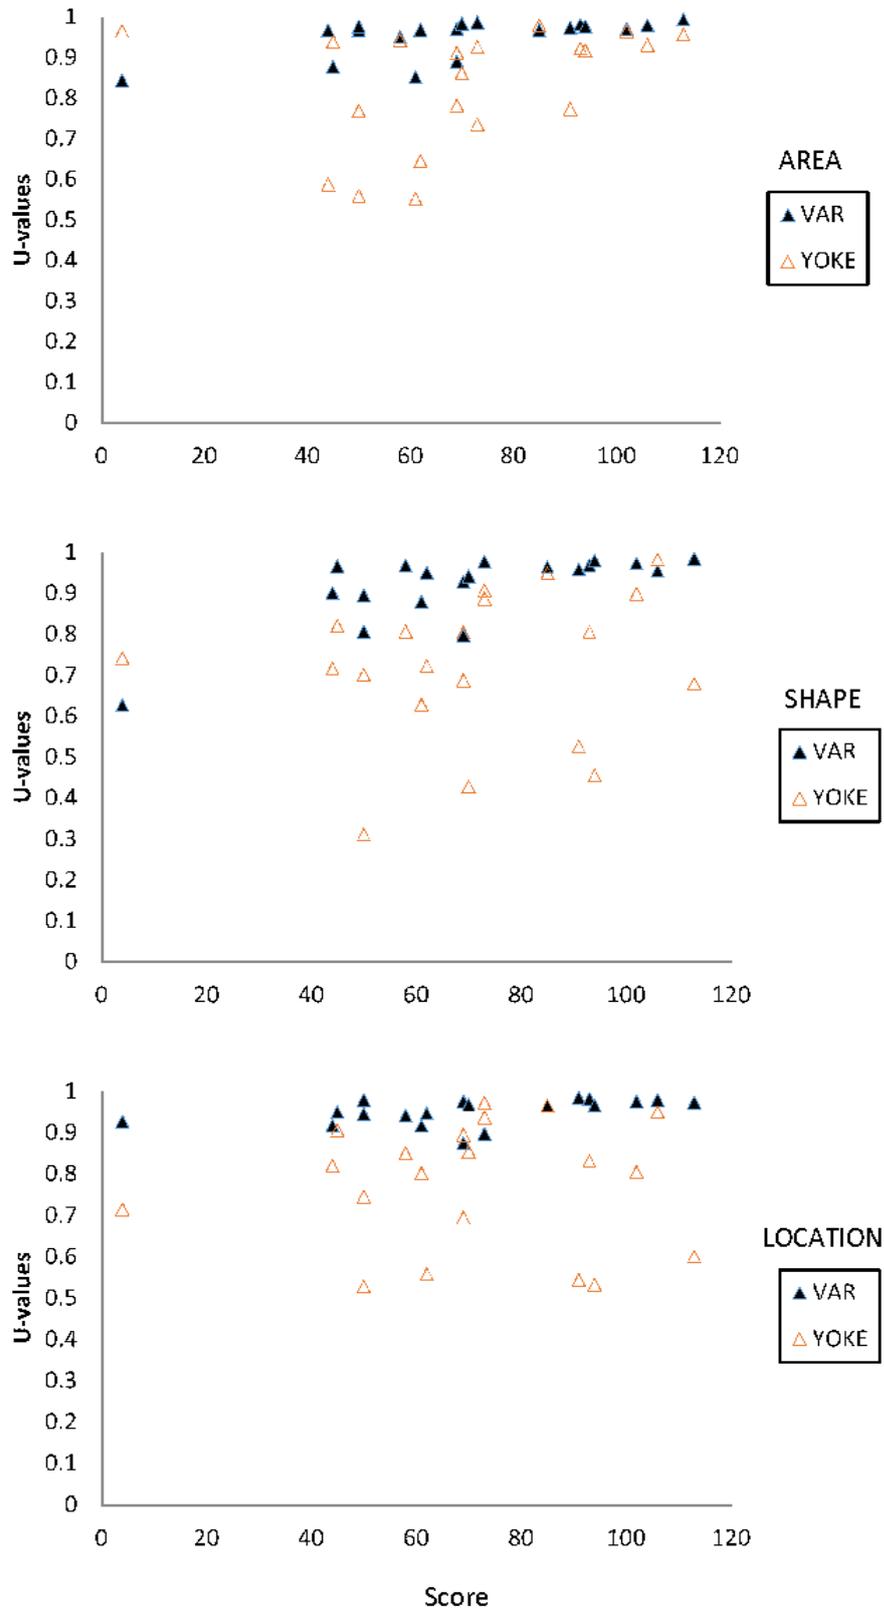


Figure 1.3. U-values plotted as a function of scores for Area, Shape and Location dimensions for both VAR (filled triangles) and YOKE (empty triangles) groups. U-values of VAR and YOKE that correspond to the same score were from individuals who were yoked to each other and received the same number of points.

Experiment 2

Introduction

The first experiment successfully showed that higher variability in responses over multiple simultaneous dimensions can be obtained by reinforcement, thus replicating Ross and Neuringer (2002).

Another aim of this thesis was to investigate whether learned variability generalises across dimensions of the same behaviour. To achieve this, one has to assess whether the variability in a dimension, where variability was not reinforced, changes as the result of reinforcing variability in other dimensions of the same behaviour. The current experiment employed a schedule that reinforced variability on only two dimensions. Variability on the third dimension was allowed but it did not produce reinforcement. It should be noted, however, the non-orthogonality of the three dimensions can be a confounding factor as variations in one variable can be restricted by the parameters of other variables. For example, for rectangles drawn on some locations of the screen (e.g., the corners of the screen), the variability of the sizes (Area) can be restricted, or if extremely large rectangles were drawn, variability in the selections of the locations will be extremely limited. As it is not clear how to overcome this confound, like Ross and Neuringer (2002), we utilised the task to attempt to examine generalisation of learned variability across dimensions.

Results from Experiment 2 were combined with results from the previous experiment to examine generalisation across dimensions. In the previous experiment, for one group of participants, variability in none of the three dimensions was required for reinforcement (YOKE). In the current experiment, reinforcement would be arranged to be contingent on variability of two of the three dimensions. Comparisons of the variability in a dimension that was not required to vary between the current experiment and the YOKE group in the previous experiment would allow examination of generalisation.

Method

Participants

Sixty students from the University of Waikato participated in the experiment.

Apparatus

The same rectangle drawing task as in Experiment 1 was used; the only difference was in the reinforcement contingency. Experimental sessions were conducted in small office rooms at the University of Waikato.

Procedures

Reinforcement of variability

The threshold value for the dimensions to be varied was kept the same as in Experiment 1. Weighted relative frequencies (WRF) for the categories used for each of the three dimensions were calculated automatically by the computer once a rectangle was drawn and then were compared to a threshold value (.0825) to determine whether that rectangle would give a reinforcer, i.e., a point would be given.

To earn a point, the WRF of the categories for two dimensions out of three had to be lower than the threshold value regardless of the WRF of the categories used for the third dimension. Participants were assigned to three groups, one group received points when they varied the categories used for both Shape and Location dimensions (Non_Area), another group received points when they varied on both Area and Location dimensions (Non_Shape), and the third when they varied on both Area and Shape dimensions (Non_Location).

Experimental procedures

The computer automatically assigned the first participant to the Non_Area group, the second to the Non_Shape group, the third to the Non_Location group, and the fourth to the Non_Area group and so forth until each group was assigned 20 participants.

The experimental procedure was the same as it was in Experiment 1 except that there was no YOKE group in this experiment.

Measures

U-values were calculated for each dimension independently for each group. Planned comparison statistics were used to compare performance between the YOKE group from Experiment 1 and each of the groups from Experiment 2 in the levels of variability on the different dimensions when variability was not required for reinforcement. Figure 2.1 shows whether variability was reinforced

for the three dimensions in each groups in Experiment 1 and Experiment 2. Dark boxes represent dimensions for which variation was required in separate conditions of Experiment 1 and Experiment 2. Comparisons were made of U-values for the dimensions where variability was not required – empty boxes; such as variability in the Area dimension between YOKE and Non_Area group, in the Shape dimension between YOKE and Non_Shape groups and in the Location between YOKE and Non_Location groups. Comparisons of U-values across dimensions for each of the groups in the current experiment were also made.

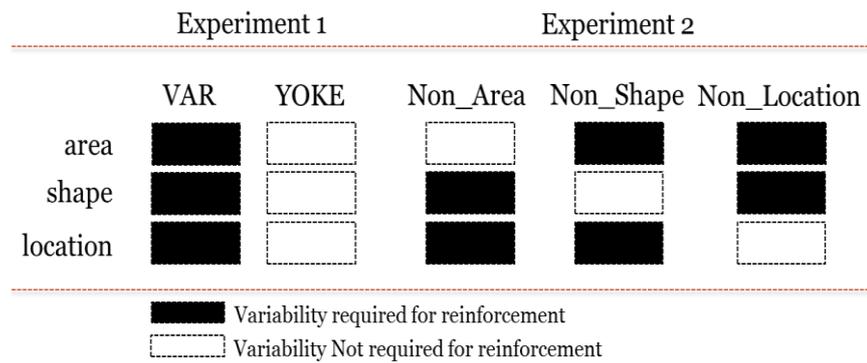


Figure 2.1. Reinforcement arrangements. Filled boxes show dimensions that were required to vary for reinforcement; empty boxes show dimensions that were not required to vary for reinforcement.

The numbers of trials meeting reinforcement criteria were calculated and compared across the three groups to examine whether performances differed because of the different dimensions required to vary. The number of trials with WRF less than .0825 was also counted for each participant for each dimension; this was used as a measure of temporal variability in the responses.

Results and Discussion

U-values – Experiment 2 and YOKE from Experiment 1

Figure 2.2 shows the mean U-values for the Area dimension (slashed bar) from Non_Area group, for the Shape dimension (dark grey bar) from Non_Shape group and for the Location dimension (dotted bar) from Non_Location group; U-values for those dimensions occurred with other dimensions that were required to vary for reinforcement from Experiment 2. Figure 2.2 also shows mean U-values for each dimension from the YOKE group in Experiment 1 (empty bars) where all three dimensions were not required to vary for reinforcement; each dimension occurred with two other dimensions that were not required to vary.

U-values for the dimensions that were not required to vary in Experiment 2 (e.g., Area dimension in Non_Area and so forth) were compared to U-values from the corresponding dimensions from the YOKE group in Experiment 1 in order to assess whether there was generalisation of variability from the reinforced dimension to those dimensions where variability was not directly reinforced. As can be seen in Figure 2.2, overall, the U-values for all three dimensions appeared to be higher from Experiment 2 when they occurred with other dimensions that were required to vary than they were from Experiment 1 when they occurred with other dimensions that were not required to vary.

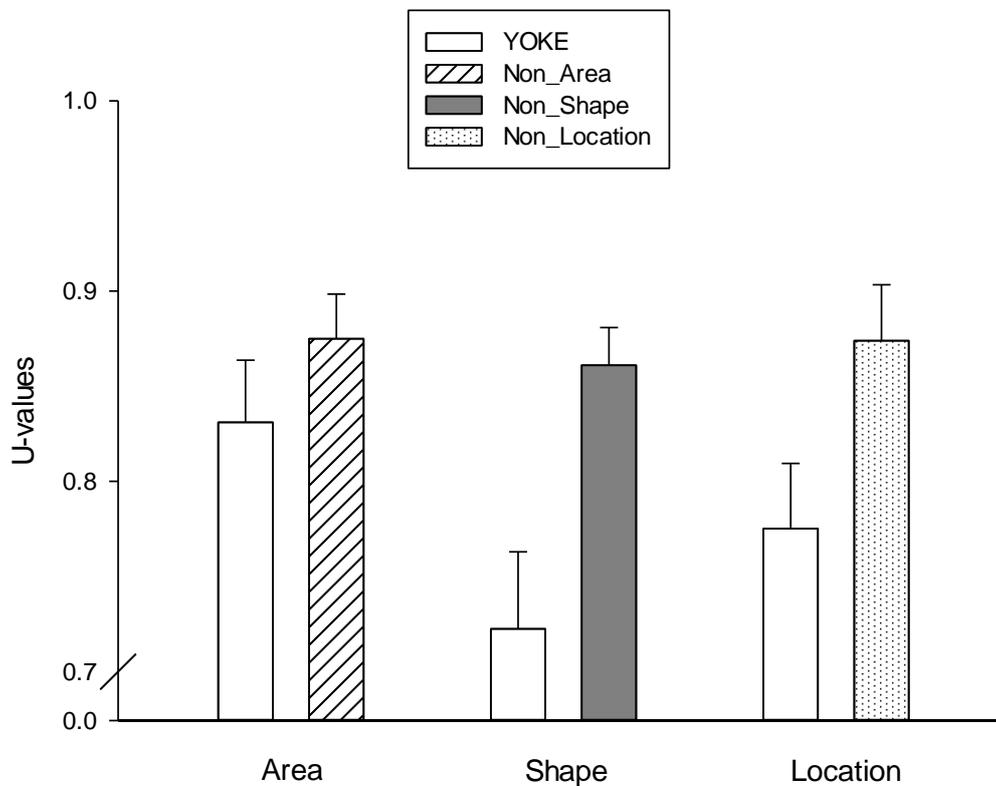


Figure 2.2. Mean U-values for Area, Shape and Location dimensions when they were not required to vary across Experiment 1 and Experiment 2. Error bars represent standard errors. U-values below .7 were omitted from y-axis.

Independent samples t-tests were carried out for the Area dimension between the YOKE and Non_Area groups, for the Shape dimension between the YOKE and Non_Shape groups and for the Location dimension between the YOKE and Non_Location groups. For the Area dimension, U-values were not significantly different between the Non_Area group and the YOKE group, $t(38) = 1.10$, $p = .281$, $r = .18$. For the Shape dimension, U-values were significantly

higher for the Non_Shape group than for the YOKE group, $t(38) = 3.09, p = .004, r = .45$. For the Location dimension, U-values were significantly higher for the Non_Location group than for the YOKE group, $t(38) = 2.19, p = .035, r = .33$. This analysis confirmed that the variability for the Shape and Location was significantly higher when these two dimensions occurred with other dimensions where variability was reinforced (as in the Non_Shape and Non_Location groups) compared to when they occurred with other dimensions that were not required to vary (as in the YOKE group).

These results suggested that for two of the three groups, Non_Shape and Non_Location groups, the reinforced variability on the two dimensions could have generalised to a third dimension where variability was not directly reinforced but occurred simultaneously with these two dimensions. These results also give support to the possible generalisation of learned variability interpreted in studies from Maes (2003) and Da Silva Souza and Abreu-Rodrigues (2010). The lack of difference in U-values between YOKE and Non_Area groups could be due to the relatively high U-values obtained by the YOKE group; this was raised as a concern in Experiment 1 where U-values being significantly higher for the Area dimension than for the other dimensions for the this group. Ross and Neuringer (2002) pointed out that the three dimensions could not be completely independent because of the constraints of the computer screen but they suggested that the correlation among the three dimensions was extremely weak. Results from current experiments suggested, that even though weak, the correlation could have been strong enough to have impacted the data collected from my sample.

Number of trials meeting reinforcement criteria

The numbers of trials meeting reinforcement criteria were counted for each of the three groups in the current experiment. These numbers were compared across the three groups to check if varying on any two of the dimensions for reinforcement was more or less difficult. One way ANOVA showed that the numbers (Non_Area $M = 114.10, SD = 39.96$; Non_Shape $M = 105.50, SD = 29.4$; Non_Location $M = 93.65, SD = 39.53$) were not significantly different, $F(2, 57) = 1.57, p = .217$. Therefore, although the three groups had reinforcement contingent

on varying different dimensions thus appear to have met the contingency on a similar number of trials.

Number of trials with Weighted Relative Frequencies (WRF) lower than threshold value

The numbers of trials with $WRF < .0825$ were counted for each dimension separately for each of the three groups. This measure shows whether the variability of categories use differed across dimensions within each group. Figure 2.3 shows the mean accumulative number of trials with $WRF < .0825$ over 300 trials for each dimension for the three groups. Data were smoothed by plotting the moving average of 30 trials.

As can be seen in Figure 2.3, for Non_Area and Non_Shape groups, there was a steeper increase for the two dimensions that were required to vary (Shape and Location for Non_Area and Area and Location for Non_Shape) than for the dimension that was not required to vary (Area for Non_Area and Shape for Non_Shape). However, the increase appeared to be similar for all three dimensions for the Non_Location group.

The mean number of trials with $WRF < .0825$ for each dimension for all three groups are plotted in Figure 2.4. Similar to the trend shown in Figure 2.3, mean numbers appeared to be lower for the Area dimension than for the Shape and Location dimensions for the Non_Area group and lower for the Shape dimension than for the Area and Location dimensions for the Non_Shape group; however, the number for the Area dimension appeared to be higher than that for the Shape and Location dimensions for the Non_Location group. Standard errors were similar across the three dimensions for the Non_Area and Non_Location group when it was greater for the Shape than for the other two dimensions for Non_Shape group.

Confirming results were found from one-way repeated measures ANOVAs for the three dimensions for each group. For the Non_Area group, the number of trials with $WRF < .0825$ was significantly different across the three dimensions, $F(2, 38) = 5.69, p = .007, \text{partial } \eta^2 = .23$. Pairwise comparisons show that the mean number was significantly lower for the Area dimension ($M = 143.55, SD = 45.48$) than for the Shape dimension ($M = 173.65, SD = 48.64; p = .012$) and for

the Location dimension ($M = 178.85$, $SD = 44.81$; $p = .009$); the number of trials with $WRF < .0825$ for the Shape and Location dimensions did not differ significantly ($p = .639$).

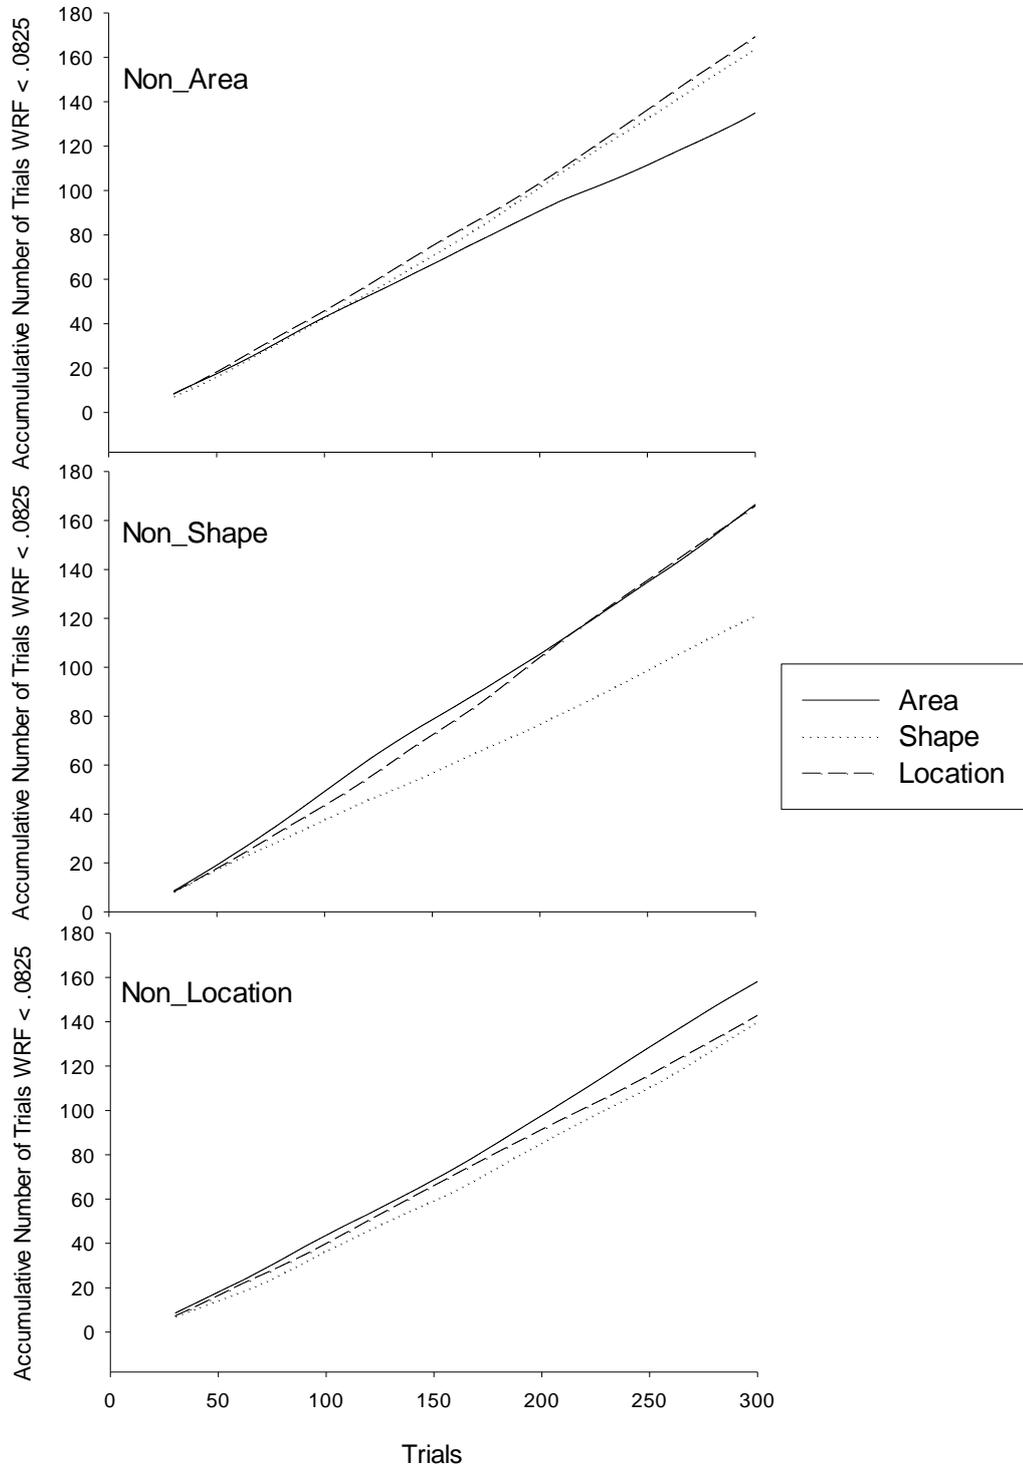


Figure 2.3. Moving average of the cumulative number of trials WRF less than .0825 for Area, Shape and Location dimensions for Non_Area, Non_Shape and Non_Location groups.

For the Non_Shape group, Mauchly's test indicated that the assumption of sphericity had been violated, $\chi^2(2) = 9.06, p = .011$, therefore degrees of freedom were corrected using Greenhouse-Geisser estimates of sphericity. Mean number of trials with $WRF < .0825$ significantly differed across the three dimensions, $F(1.43, 27.23) = 23.95, p < .001$, partial $\eta^2 = .56$. Pairwise comparisons indicated that the numbers of trials with WRF below the threshold value were lower for the Shape dimension ($M = 123.60, SD = 40.97$) than they were for the Area dimension ($M = 177.75, SD = 21.36; p < .001$) and for Location ($M = 172.35, SD = 38.69; p < .001$); the numbers for the Area and Location were not significantly different ($p = .584$). For the Non_Location group, the mean number of trials with $WRF < .0825$ were not significantly different across Area ($M = 166.30, SD = 53.04$), Shape ($M = 148.7, SD = 49.06$) and Location ($M = 151.00, SD = 51.93$), $F(2, 38) = 2.23, p = .122$.

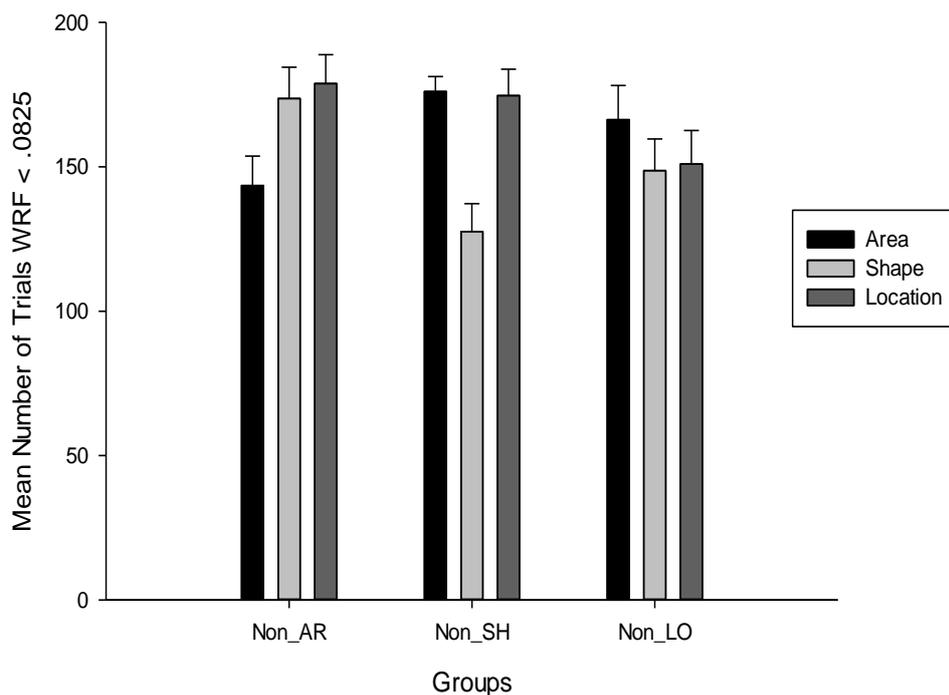


Figure 2.4. Mean number of trials $WRF < .0825$ for Area, Shape and Location dimensions for Non_Area, Non_Shape and Non_Location groups. Error bars represent standard errors.

Overall, when reinforcement was contingent on varying the Shape and Location dimensions (Non_Area), participants drew rectangles using more infrequently used shapes at more infrequently used locations than using infrequent sizes. When reinforcement was contingent on varying the Area and Location

dimensions, participants drew rectangles of relatively infrequent sizes and locations (Non_Shape) while the tendency to use infrequently used shapes was lower. However, when reinforcement was contingent on varying the Area and Shape dimensions, participants drew rectangles of similarly infrequently used sizes, shapes and locations.

U-values for Experiment 2 only

U-values were calculated independently for each dimensions for each of the three groups; and they were compared between the dimensions that were required to vary and the dimension that was not required to vary within each group and across groups.

Figure 2.5 shows mean U-values for the Area, Shape and Location dimensions for the Non_Area, Non_Shape and Non_Location groups. The differences in U-value across dimensions were similar to the differences in number of trials with $WRF < .0825$; U-values appeared to be lower for the dimension that was not required to vary than the two which were required to vary only for Non_Area and Non_Shape groups.

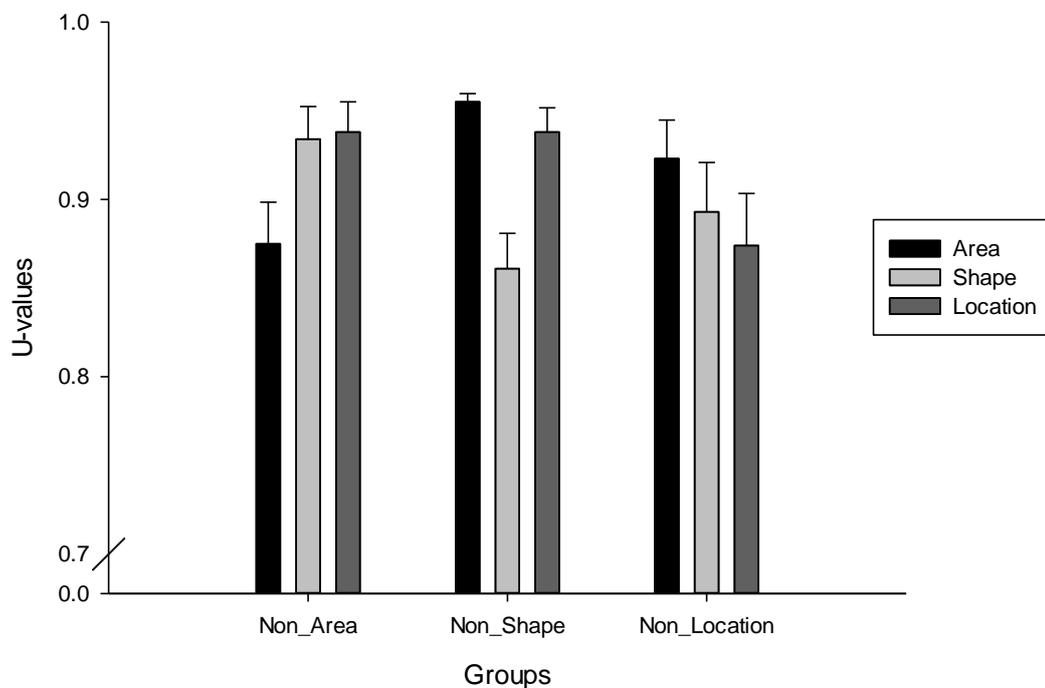


Figure 2.5. Mean U-values for Area, Shape and Location dimensions for Non_Area, Non_Shape and Non_Location groups. Error bars represent standard errors.

One-way repeated measures ANOVAs were carried out for each of the groups. For the Non_Area group, Mauchly's test indicated that the assumption of sphericity had been violated, $\chi^2(2) = 17.58, p < .001$; therefore the degrees of freedom were corrected using Greenhouse-Geisser estimates of sphericity. U-values were found to be significantly different across dimensions, $F(1.23, 23.41) = 4.87, p = .031$, partial $\eta^2 = .20$. Pairwise comparisons showed that U-value for the dimension that was not required to vary, Area ($M = .88, SD = .11$), was lower than that for both Shape ($M = .93, SD = .08; p = .04$) and Location ($M = .94, SD = .08; p = .029$), which were required to vary. For the dimensions that were required to vary, Shape and Location, no significant difference was found ($p = .702$).

For the Non_Shape group, Mauchly's test indicated that the assumption of sphericity had been violated, $\chi^2(2) = 13.64, p = .001$, therefore the degrees of freedom were corrected using Greenhouse-Geisser estimates of sphericity. U-values were found to be significantly different across dimensions, $F(1.31, 24.81) = 20.66, p < .001$, partial $\eta^2 = .52$. Pairwise comparisons showed that U-value for the dimension that was not required to vary, Shape ($M = .86, SD = .09$), was lower for both the Area ($M = .95, SD = .02; p < .001$) and the Location ($M = .94, SD = .07; p < .001$) dimensions, which were required to vary. For the dimensions that were required to vary, Shape and Location, no significant difference was found ($p = .257$).

For the Non_Location group, the difference in U-values across dimensions was found to be non-significant, $F(2, 38) = 2.88, p = .069$. However, pairwise comparisons showed that mean U-value for the Location dimension ($M = .87, SD = .13$) was significantly lower than that for the Area dimension ($M = .92, SD = .10; p = .016$) but not than the Shape dimension ($M = .89, SD = .12; p = .319$); no significant difference was found between the two dimensions required to vary – Area and Shape ($p = .229$). Again, results from this group is inconsistent with the other two groups.

As can be seen in Figure 2.5, the U-values for the Location dimension for the Non_Location group appeared to be similar to the U-values for Shape dimension for the Non_Shape group and for the Area dimension for the Non_Area group; all of these dimensions were not required to vary. A one-way ANOVA

across these dimensions on U-values showed no significant difference, $F(2, 57) = 0.095, p = .9$.

For all the dimensions that were required to vary, U-values for the Area and Shape dimensions for the Non_Location group appeared to be lowest of the six. A one-way ANOVA across these six U-values showed no main effects, $F(5, 48.39) = 1.629, p = .170$; Leven's test showed that the assumption of homogeneity was violated (Leven statistic $(5, 114) = 3.771, p = .003$) therefore, the Welch's test statistics is reported instead. Post Hoc analysis using LSD indicated that mean U-value for the Shape dimension for the Non_Location group was significantly lower than it was for the Area dimension for the Non_Shape group, $p = .022$; no significant difference in U-values was found across other dimensions.

Overall, U-values for the dimensions that were not required to vary were not significantly different across the three groups. For the dimensions that were required to vary, although the differences in U-values between these dimensions were generally not significant, U-values for the Area and Location dimensions for the Non_Location group were slightly lower than those for the Shape and Location dimensions for the Non_Area group and the Area and Location dimensions for the Non_Shape group. If the variability in the Area and Shape dimensions for the Non_Location group had been higher, the results for the Non_Location group would have been consistent with those of the other two groups. The lack of difference for this group because of the lower variability in the dimensions that were required to vary (as compared to dimensions required to vary in other groups) could be due to the fact that when Location was not required to vary, participants varied less on the other two dimensions that were required to vary too. The requirement to vary on Location dimension in the other two groups, Non_Area and Non_Location, may have made the reinforcement contingency more salient to the participants. This suggests that the non-orthogonality of the three dimensions could have affected the variability of individual dimensions.

Summary

Generalisation of increased variability to a third dimension was evident for two of three groups when variability was required on two other dimensions. When variability was not required, the variation was higher for the Shape and Location

dimensions when they occurred with the other dimensions where variability was required compared to when they occurred with the other dimensions for which variability was not required. This suggests that it was likely the learned variability had generalised to the Shape dimension when reinforcement was contingent on the variability of the Area and Location dimensions; and to the Location dimension when reinforcement was contingent on variability of the Area and Shape dimensions. However, the reinforced level of variability in the Shape and Location dimensions did not appear to have generalised to the Area dimension because variability in the Area dimension was similar when it occurred with other dimensions that were required to vary (Non_Area) and when it occurred with other dimensions that were not required to vary (YOKE). The lack of difference in the variability for the Area dimension between these two conditions would have been the result of relatively high variability obtained for the Area dimension for the YOKE group. Therefore, the failure to find generalisation of learned variability for Area dimension may be due to the lack of orthogonality of the three dimensions as suggested in the discussion in Experiment 1.

For the three groups in Experiment 2, although participants met the reinforcement contingency similarly well, results from both measures of variability, the number of trials with WRF less than threshold value and U-values, suggested that the group that received reinforcement for varying both Area and Shape dimensions responded differently from the other two groups. The difference in the dimensions required to vary may have confounded the level of variability in individual dimensions.

Results from both Experiment 1 and Experiment 2 identified that the non-orthogonality of the three dimensions could be confounding; however, exactly how it affected the results was only an interpretation. Also, how the orthogonality issue interacted with the reinforcement contingency (e.g., in Experiment 2) was not clear. The non-orthogonality of the three dimensions in the rectangle drawing task which was not reported to have affected the results in Ross and Neuringer (2002) was shown to cause concerns from the analyses in Experiment 1 and Experiment 2 in the current thesis. Because of this, it was premature to conclude that generalisation was found. The issue of non-orthogonality of dimensions in the rectangle drawing task turned out to be a much bigger problem than both Ross and

Neuringer (2002) and we originally thought. To continue with investigation of the generalisation of learned variability across dimensions, alternative methodologies with completely independent dimensions is required.

Experiment 3

Introduction

The extent to which learning in one context extends to another has long been of interest to behaviour analysts (e.g., Arnold-Saritepe et al., 2009; Stokes & Baer, 1977; Stokes & Osnes, 1989). Comparisons of results between Experiment 1 and Experiment 2 showed generalisation of learned variability was likely to have occurred; however, the non-orthogonality of the three dimensions of the rectangles confounded the results. To continue with the investigation of generalisation of learned variability across dimensions, a task with multiple dimensions that were completely independent of each other was needed. Studies in reinforcing behavioural variability related to enhancing creativity reviewed in general introduction in this thesis commonly study the diversity in colours and forms (e.g., Fallon & Goetz, 1975; Ryan & Winston, 1979). Generalisation of trained diversity across tasks had also been investigated in creativity-related studies (Holman, et. al., 1977; Goetz et al., 1973). The dimensions being commonly studied, such as colours and shapes in drawings and forms of block-building, were completely independent dimensions; and they were common in children's school learning environment and thus are worthwhile studying. However, those studies related to creativity relied on subjectivity in the reinforcement procedure which could have confound the results. Specifically, there was a judgement made by the experimenter as to whether the behaviour exceeded a subjective criteria for the delivery of reinforcement. Procedures that require an observer to judge immediately whether a response meets reinforcement criteria and then deliver the reinforcers immediately in a uniform manner can cause delays to reinforcement, and errors and bias in judgement can also occur.

Experiment 3 aimed to use a computerised task that incorporated the dimensions commonly used in creativity related studies to examine the impact of reinforcement on variability on multiple dimensions of a behaviour whilst simultaneously eliminating problems with accurate reinforcer delivery evident with alternative procedures reviewed above. Methodologies for reinforcement schedules and measures of variability used in the previous two experiments were employed. Success in obtaining similar results would form a foundation for studies of the generalisation of learned variability across dimensions that are more

common in everyday life and also that are completely independent from each other.

Methods

Participants

Participants were recruited through advertisement on noticeboards and on course forum and on social networking sites (e.g. Facebook). Data from 34 participants (20 aged 16-20 years; 11 aged 21-25 years; 2 aged 31 years and over; and of these 10 were male and 24 were female) were included for analysis. Data from six participants were excluded due to programming errors. Participants who were first year psychology students were given course credits for their participation. Ethics approval was obtained from the Research and Ethics Committee of School of Psychology at University of Waikato (12:14) before data collection started.

Apparatus

Sessions were conducted in either of two small offices at the School of Psychology or in a small office at the Tauranga campus of the University of Waikato. All sessions conducted at the School of Psychology had one participant in a room; and all used a desktop computer for their participation. Most of the sessions at the Tauranga campus had one participant in the room, with a few sessions having more than one participants in the same room. Some participants from the Tauranga campus used a provided laptop (a Sony VAIO SVE-151A11W) and some used their own laptop (details not recorded). The ratios of the display for the experiment were fixed so that using different screens would not affect the appearances of the elements presented on the screen.

Procedures

Creating Categories for the Dimensions

Three completely independent dimensions were used; these were shape, colour and pattern (texture). Each dimension included eight distinctively different categories; that is, there were eight shapes, colours and patterns (see Appendix I for a screen capture of the elements used in this experiment). Colours used were red (rgb 216, 24, 35), pink (rgb 231, 90, 148), green (rgb 121, 189, 119), orange (rgb 246, 129, 67), deep sky blue (rgb 87, 196, 228), yellow (rgb 255, 206, 75),

brown (186, 131, 10) and royal blue (rgb 124, 150, 205). The number of categories was eight for each of the three dimensions, which is half of that of Experiment 1. The main reason for this was to try to use colours that were relatively distinct from each other. Shapes and patterns were created using Adobe Illustrator Artwork 8.0. Efforts had been made to use simple two-dimensional shapes to provide distinctive shapes and patterns for the designs. Appendix I shows the categories used for the three dimensions as well as the instructions presented. Each category for each dimension was assigned a number for the convenience of data collection and analysis; the categories and their corresponding numbers were shown in Appendix I.a. The corresponding numbers for the categories were not visible to participants.

Threshold contingencies

The method for obtaining an appropriate threshold value for this task was the same as it was for the rectangle drawing task in the previous experiments. Highly variable responses are believed to be similar to random responses. If responses were randomly distributed over the eight categories for each dimension, the relative frequencies for each category for each dimension would be 0.125; each category would be responded to equally often. The computer was programmed to generate 300 random responses; a threshold of 0.165 was used instead because 75% of the trials from the random generator met reinforcement contingency.

Reinforcing variability

In order to encourage participants to vary their use of categories when creating the combinations of the three dimensions, points were given for combinations created using categories that were used infrequently for all three dimensions. Specifically, a response (combination) was reinforced only when the categories used for all three dimensions was lower than the threshold value 0.165. That is, participants had to vary all three dimensions: shapes, colours and patterns.

Reinforcement independent of variability

A second group of participants received points regardless of the categories from the three dimensions they used to create the combinations. The same yoking procedure used in Experiment 1 was also used in this experiment. A participant from the yoke group received a point whenever the participant he or she was

yoked to in the first group received a point. The yoking procedure was used to ensure the two groups received the same amount of reinforcement.

Experimental procedures

At the beginning of each session, the researcher briefly introduced the experiment and the approximate time needed for the session. Participants were then instructed to sit in front of the computer and to follow the instructions on the screen. The computer program presented brief information about the research project and collected consent before the experiment started.

Appendix I showed screen captures of the computer program used. Once the game started, the instructions (in black font) and the elements (categories) for Shapes (horizontally displayed on the top of the screen), Patterns (vertically displayed on the left edge of the screen) and Colours (vertically displayed on the right edge of the screen) appeared on a light grey background (Appendix I.b). The instructions were as follow:

“Please choose a shape (top), then choose a pattern (left) and then a colour (right). You won’t be able to change once you have clicked on an element. Whether you will get a point will be determined by the combination of the three elements. Your GOAL is to get as many points as possible. You will hear a tone (“beep”) when you earn a point. Click Start when you are ready. Good luck!”

After participants clicked on the button “Start” at the end of the instruction screen, the first trial started. The instructions disappeared and the shapes on the top of the screen were illuminated while elements for the other two dimensions were darkened (Appendix I.d). Participants had to move the mouse over the shape they decided to use. Once a shape was clicked on, that particular shape would appear in the centre of the screen. At the same time, the elements for shape were darkened while the elements for pattern on the left were illuminated; the elements for colour on the right remained darkened (Appendix I.e). Once a pattern was clicked on, the shape in the centre of the screen would be filled with the pattern chosen; and elements for shape went dark while the colours on the right were illuminated (Appendix I.f). When the participant clicked on a colour, the colour would fill the shape which had been filled with pattern in the centre; elements for all three

dimensions were darkened at this point, with only the combination in the centre illuminated (Appendix I.g). Ascending tones (100-ms 1500-Hz followed by 100-ms 2000-Hz) indicating the receipt of points were delivered by the computer through the built-in speakers immediately after each combination that met the reinforcement criteria. In sessions with more than one participant, participants used earphones for the auditory feedback. No other feedback was provided. Every combination remained on the screen for 500 ms and mouse clicks during this period had no effect; the screen cleared after 500 ms and the next trial started with the shapes on the top of the screen being illuminated. Participants were debriefed at the end of the experiment. For all trials, mouse clicks were only effective on the dimensions that were illuminated (to control for variability in the sequences of dimensions use). A dialogue box with instruction “*Please pick a (shape/pattern/colour) from the (top/left/right) by clicking on any of them*” (Appendix I. h & i) would appear if no responses were made after 10 s had elapsed since a dimension had been illuminated. The total number of points received will be displayed on the screen once participants finished creating 300 combinations.

Measures

Number of trials meeting reinforcement criteria

The number of trials meeting reinforcement for the VAR group equals the number of points participants obtained. For YOKE group, although the points were given independently of the categories used in the responses, the computer recorded whether those responses would have met the reinforcement criteria if it were in effect.

U-values

U-values were calculated using the same formula for shape, colour and pattern dimension independently.

$$U - value = - \sum_{i=1}^{\beta} \frac{\alpha_i \times \log(\alpha_i)}{\log(\beta)}$$

Same as Experiment 1 and 2, U-values were calculated based on the relative frequencies of each of the eight categories to examine the evenness of responses distributed to these categories. It was the main measure of variability.

Results

Number of trials meeting reinforcement criteria

Figure 3.1 shows the percentage of participants meeting the reinforcement criteria over trials on a moving average of 30 trials. As can be seen in the figure, both the VAR and YOKE group started with an average of 40 percent participants meet reinforcement criteria from the first 30 trials and the percentages did not differ between the two groups until after the 80th trial. The percentage of participants for the VAR group then increased slightly through to the end of the 300 trials while the percentage of participants for the YOKE group fluctuated. The percentage of participants meeting reinforcement criteria remained higher for the VAR group than it was for the YOKE group from approximately the 100th trial till the end of the 300 trials. Paired samples t-test shown that, overall, the number of trials meeting reinforcement criteria was significantly higher for the VAR group ($M = 151.35, SD = 52.30$) than it was for the YOKE group ($M = 117.94, SD = 56.68$), $t(16) = 2.79, p = .013, r = .57$.

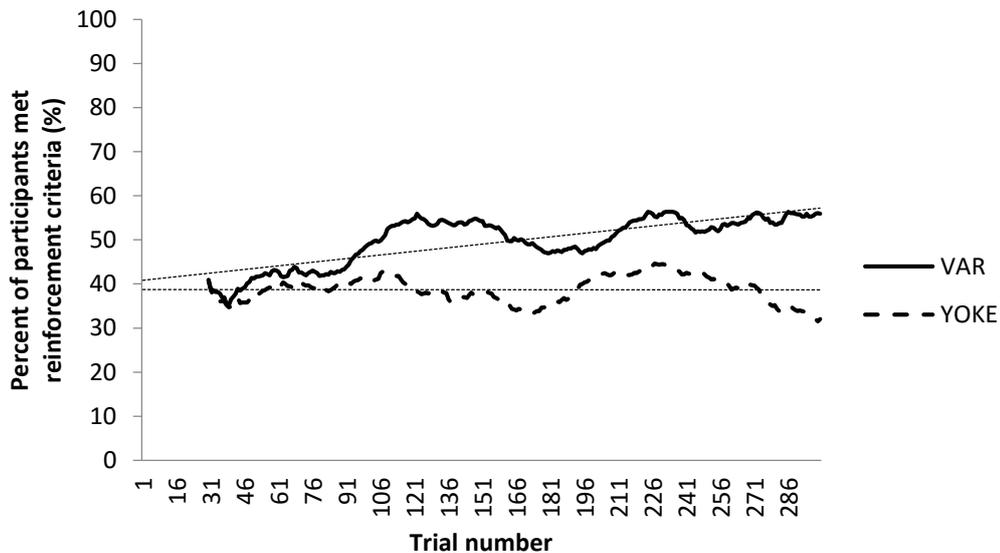


Figure 3.1. Percentage of participants meeting reinforcement criteria over 300 trials for both VAR (solid line) and YOKE (dotted line) group. Each data point represents the moving average of 30 trials.

U-values

Mean U-values for the three dimensions for both VAR and YOKE groups are shown in Figure 3.2. U-values were the main measure of variability and were compared between the group which received reinforcement contingent on variability (VAR) and the group which received reinforcement independent of

variability (YOKE). U-values were calculated independently for each dimension for both VAR and YOKE groups. The left three sets of bars in Figure 3.2 (8 categories) show mean U-values for the three dimensions for both VAR and YOKE group. As can be seen, all mean U-values were greater than .9; and although the mean U-values appeared to be higher for the VAR group for Colour and Pattern dimensions, the standard errors for these two dimensions were greater than those for the YOKE group. Results from repeated measures ANOVA show that the U-values calculated based on eight categories were significantly affected by the reinforcement contingencies, $F(1, 32) = 5.30, p = .028$, partial $\eta^2 = .14$. However, paired-samples t-tests for each dimension between VAR and YOKE groups found no significant differences. Specifically, as indicated by U-values, variability in the use of shapes was not significantly different between the VAR group ($M = .98, SD = .03$) and the YOKE group ($M = .96, SD = .05$), $t(16) = .95, p = .357$.

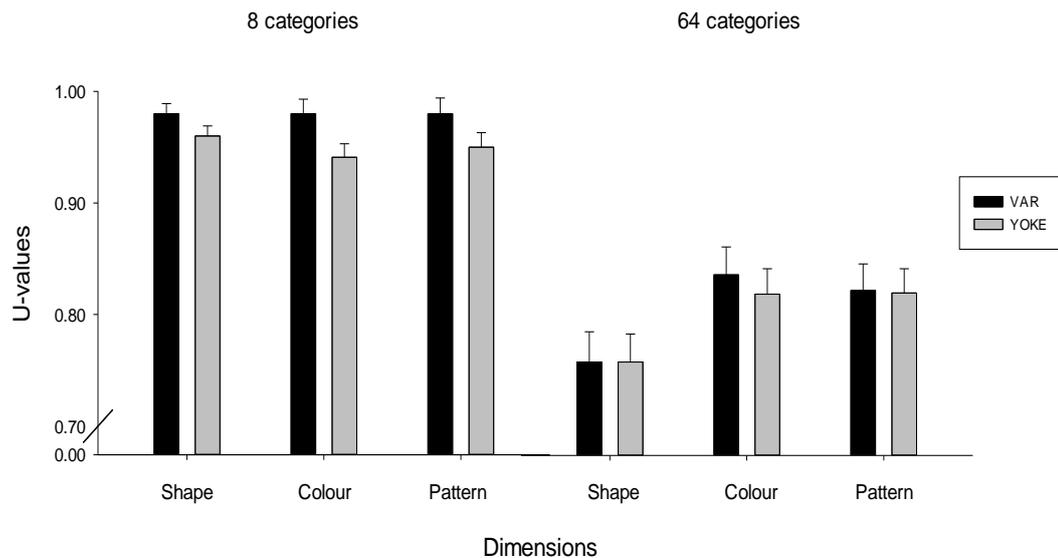


Figure 3.2. Mean U-values for the three dimensions for VAR and YOKE groups. The three sets of bars on the left were for U-values calculated based on 8 categories; and the three sets on the right were calculated based on 64 categories (Lag 1). Error bars represent standard errors.

This was also true for variability in the use of colours for VAR ($M = .97, SD = .03$) and YOKE ($M = .93, SD = .10$), $t(16) = 1.67, p = .114$. Also, no significant difference in the variability in patterns used was found between the VAR group ($M = .98, SD = .03$) and the YOKE group ($M = .93, SD = .10$), $t(19.50) = 1.65, p = .115$. These results were not consistent with the results from

Experiment 1 not only in that there was no significant difference in U-values between the VAR and YOKE groups, but also in that the U-values for the YOKE group were very high. Judging from the high U-values, variability in the categories used for all three dimensions appeared to be high for both the VAR and YOKE group.

Note because of the number and complexity of the figures used to help display the data they are presented in Appendix II, which shows graphs of the categories used for each dimension in each trial over the 300 trials (x-axis) for all participants. The categories used (y-axis) are represented by the corresponding numbers assigned to them (as shown in Appendix I.a). Figures for responses generated randomly from the computer (panels r1, r2, r3 in Appendix II) and responses from the VAR group (panels a1, a2, b1, b2, c1 and c2 in Appendix II) and the YOKE group (d1, d2, e1, e2, f1 and f2 in Appendix II) were included. The participant ID and U-values are shown next to each of the response graph. The same numbers were assigned to the pair of participants who were yoked, for example, participant YOKE1 was yoked to participant VAR1 and so forth.

Looking at panels r 1, 2, 3 in Appendix II, it shows that there is no systematic patterns in the responses generated randomly by the computer (panel r 1, 2, 3); there were a few repetitions of the same category but the number of repetitions was small and non-systematic. The responses from the participants (panels a1 to f2) show that most of the response patterns from both groups deviate from those of the random responses in that they appear to be more systematic. Three main patterns can be observed; one is like steep stairs (e.g., VAR9 in panel a1, VAR17 in panel b2, YOKE2 and YOKE14 in panel d1, and YOKE13 in panel d2) which would be a result of participants cycling through the options. Another pattern looks like wider stairs (e.g., VAR14, VAR1 and VAR17 in panel a2, VAR13 in panel b2) which would result from participants repeating the same category a few times before they moved onto the next adjacent category. A third pattern is a mix of the previous two (e.g., VAR14 in panel a2, YOKE 5 in panel f1 and YOKE13 in panel f2). These responses, although being stereotypic and deviating greatly from the pattern of random responses, obtained similarly high U-values as random responses did. Therefore, the stereotypy in the responding from the participants mentioned above did not appear to have been captured by the high

U-values previously calculated. A modified version of U-values were then calculated using the same sets of responses with consideration of response sequences to better capture the difference in responding between the random generator and the participants.

U-values calculated based on 64 categories (Lag 1)

The new set of U-values was calculated taking into account the sequences of the categories used in two consecutive trials; that is, the category used in the current trial and the category used in the trial immediately prior to the current trial (referred to as U-64 thereafter). Although it might be sensible to calculate U-values considering sequences over more than two trials, calculating U-values required a large number of responses so that there are enough opportunities to respond to each category. Calculating U-values based on sequences of more than two consecutive trials would not be meaningful as there would be a large number of categories not used and this would result in very low U-values, and would not help in understanding the variability of the responses emitted. In order to calculate U-values based on the two-consecutive-trials sequences, new sets of categories were created for each dimension. Sixty four new categories including all possible combination of any two original categories were created. If you take the dimension of Colour for an example, the sequence of colours appearing on the screen was red (1), pink (2), green (3), orange (4), deep sky blue (5), yellow (6), brown (7) and royal blue (8), the new categories would be red-pink (12), pink-green (23), green-royal blue (38) etc. etc.. If responses tended to be stereotypical in ways observed from the response graphs in Appendix II, either cycling from one category through to the last one or repeatedly using the same categories before moving onto using the next adjacent category (original), some categories would have been used more frequently than others, which would result in lower U-values. To illustrate this using colour dimension again, cycling through categories 1 to 8 (e.g., 1, 2, 3, 4, 5, 6, 7, 8, 1, 2, 3, 4, 5, 6, 7, 8, 1, 2, 3, 4, 5, 6, 7, 8, 1, 2, 3, 4, 5, 6, 7, 8...) would result in higher frequencies for categories of 12, 23, 34, 45, 56, 67, 78, 81 (12, 3; 23, 3; 34, 3; ...) than any other possible sequences (e.g., 16, 0; 35, 0; 76, 0; 83, 0; ...). On the other hand, repeatedly using the same categories until no reinforcement could be obtained then moving on to the next category (e.g., 1111222233334444555566667777888811112222...) would result

in higher frequencies for categories 11, 22, 33, 44, 55, 66, 77, and 88 (e.g., 11, 3; 22, 3; 33, 3; ...) than any other possible sequences (e.g., 24, 0; 78, 1; 73, 0; ...). Either of these two response patterns would result in lower U-values calculated based on 8 categories than if these responses were random. Based on the difference in the response patterns observed between the random responses and those from the participants, we would expect U-64 to be lower for both the VAR and the YOKE group compared to the random generator. Moreover, by calculating U-64, we would be able to find out whether participants from the VAR group responded more or less stereotypical than those from the YOKE group did.

U-64 was calculated independently for each of the three dimensions. The three sets of bars on the right in Figure 3.2 (64 categories) show the mean U-values calculated based on 64 categories for all three dimensions for the VAR and the YOKE groups. There did not appear to be much difference between the VAR and the YOKE for the Shape and Pattern dimensions; and when U-64 appeared to be lower for the YOKE group than for the VAR group, the standard errors for YOKE group were much greater than they were for the VAR group. Repeated measures ANOVA showed that when the sequences of two consecutive trials were taken into account, U-values were not significantly impacted by the difference in reinforcement contingencies between VAR and YOKE group, $F(1, 32) = .157, p = .695, \text{partial } \eta^2 = .005$. Paired samples t-tests between the VAR and YOKE groups for all three dimensions showed that there were no significant differences ($p = .918$ for Shape, $p = .638$ for Colour, and $p = .976$ for Pattern).

To summarise, the overall level of variability of responses, with consideration of the sequences of responding, was similar for the two groups receiving different reinforcement contingencies.

Discussion

One of the aims of the current experiment was to compare the variability of responses on the use of shapes, colours and patterns between two groups of participants receiving different reinforcement contingencies. Specifically, higher variability was expected to be found for the group who received reinforcement dependent on the variability of all three dimensions (VAR) compared to the group who received reinforcement independent of variability in their responses (YOKE). Overall, participants from the VAR group met reinforcement criteria on more

trials than participants from the YOKE group did; and more participants met reinforcement contingency from the VAR group than from the YOKE group over trials. This was consistent with results from Experiment 1 using a rectangle drawing task in which dimensions were continuous and the category boundaries were not visually available to participants. As can be seen in Figures 1.1 and 3.1, there were higher percentages of participants meeting reinforcement criteria for both VAR and YOKE group in the current experiment than there were in Experiment 1. This could be due to the task being easier with higher threshold value required by an eight-category task. It could also be that the categories were visible in the current experiment which made it easier for participants to work out ways of obtaining reinforcement.

U-values as the main measure of variability were initially calculated as they normally are, based on the relative frequencies of the 8 categories (U-8). However, high U-8 ($> .93$) were obtained for all three dimensions for both the VAR and YOKE group; and there were no significant differences in U-8 between the two groups on any of the three dimensions. That is, participants from both the VAR and YOKE groups, regardless whether the reinforcement were delivered dependent on variability or not, used all of the eight categories for all three dimensions equally often. Recall that up to the current experiment, including Experiment 1 and Experiment 2, there was a same underlying assumption that highly variable responses would be similar to random responses. Therefore, three sets of random responses (random between 1 and 8) were generated using the computer in order to calculate U-8 to compare with the ones obtained from the participants. The three U-values calculated based on 8 categories for the random responses were .998, .997, .989; U-values for the three dimensions for VAR and YOKE were similarly high, with those of VAR (.978, .976, .976) closer to that of the random responses than those of YOKE (.964, .933, .933). As indicated by the high U-8 for both groups, responses had been allocated to each of the 8 categories evenly as random responses would be. Inspection of the response graphs in Appendix II revealed that there were significant deviations in responses produced by the participants from both the VAR and YOKE group from those generated randomly; responses from most of the participants were more systematic than random responses.

New sets of U-values (U-64) were calculated using the same responses but taking into account the sequences of the responses to better capture the stereotypy observed in the graphs. Figure 3.3 shows U-values calculated based on 8 (U-8; filled circles) and 64 (U-64; empty circles) categories for Shape, Colour and Pattern dimensions for both VAR and YOKE groups and for random responses generated by the computer. As can be seen, U-8 and U-64 were similar for random responses (.998, .997, .989 and .966, .976, .970, respectively) while U-64 for the three dimensions for both VAR and YOKE groups were generally lower than U-8. U-64 appeared to have captured at least some of the stereotypy in the responding from both groups of participants; and responses from the two groups were less variable than the random responses despite the similar high U-8. Both U-8 and U-64 show that there was no significant difference between the VAR and YOKE groups; however, participants from the VAR group had more trials meeting reinforcement criteria. If the two groups were similarly variable in their responses on all three dimensions, there should be no difference in the number of trials meeting criteria when the criteria were based on the relative frequencies of the categories. This suggested that the two groups could have come to similar U-8 and U-64 values by different means. As shown in Figure 3.3, U-8 and U-64 were very close for random responses; and there were more similar U-8 and U-64 observations for the YOKE group compared to the VAR group. This means that participants from the YOKE group might have produced more random like responses than participants from the VAR group did.

To better show how the changes in U-64 reflect the changes in responding, figures in Appendix II were arranged in descending order based on U-64; for the same dimension, the first panel showed figures with higher U-64 when the second panel showed figures with lower U-64. Thus, responses with the highest U-64 would appear first on the first panel for a dimension; and responses with the lowest U-64 would appear at the bottom of the second panel of that dimension. Looking at the first panels of each dimension for both groups (a₁, b₁, c₁, d₁, e₁ and f₁), only a few figures (YOKE11, YOKE 8 and YOKE9 from panel e₁ and YOKE11, YOKE12 and YOKE14 from panel f₁) appeared to have similar pattern to the random responses; and all were from the YOKE group with none from the VAR group. These observations were consistent with what were found in Figure

3.3 in which more instances of U-64 being close to U-8 were found for the YOKE group than for the VAR group.

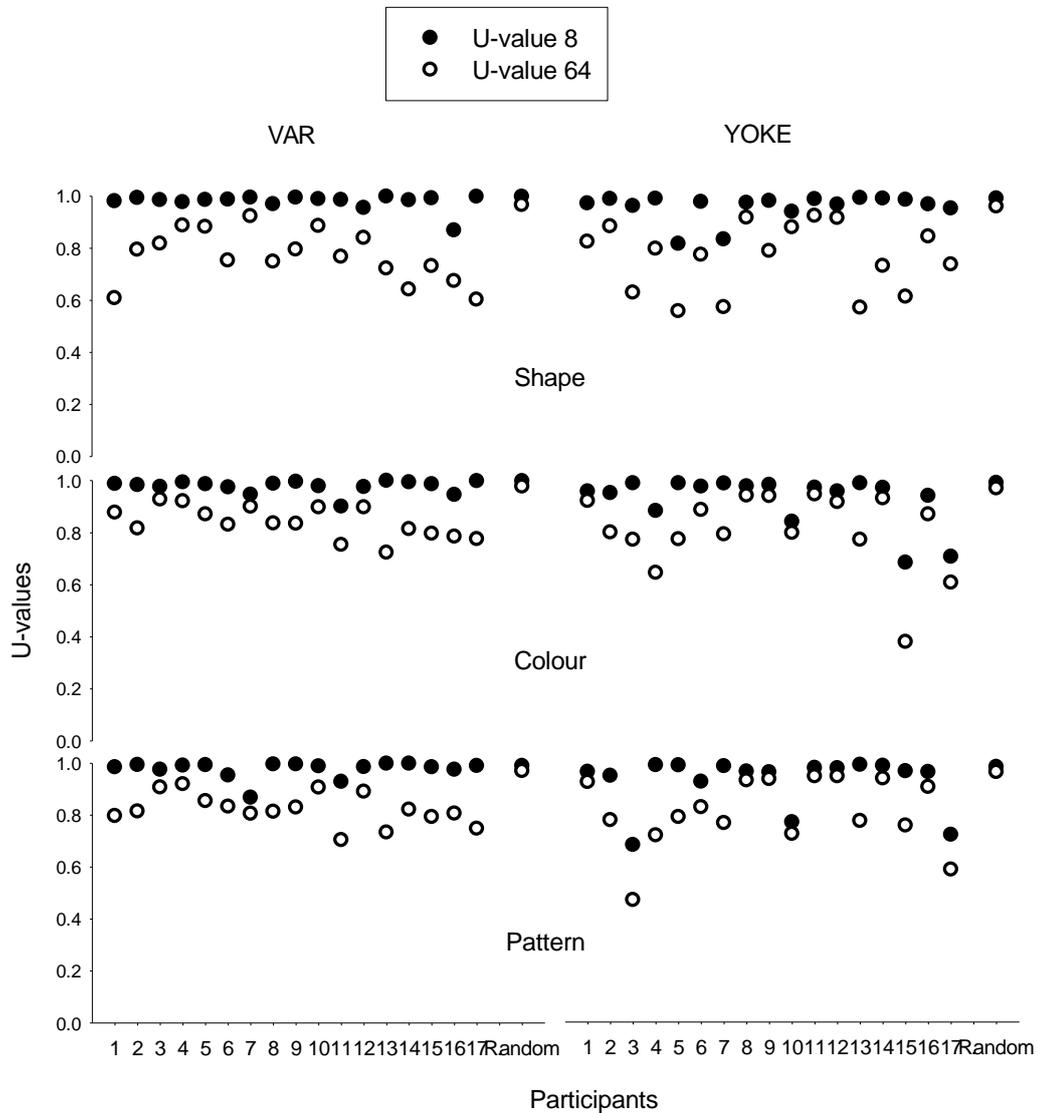


Figure 3.3. U-values calculated based on 8 (filled circles) and 64 (empty circles) categories for Shape, Colour and Pattern dimensions for 17 pairs of VAR-YOKE participants. Data points corresponding to a participant in a “Random” condition were U-values calculated from 300 computer generated random responses (between 1 and 8).

Summary

The analysis of U-8 and U-64 as measures of variability indicated that the threshold reinforcement contingency based on the relative frequencies of the options used for each dimension failed to produce more variable (random like) responses in the variability contingent group (VAR) than the control group (YOKE). Response graphs suggested that responses from the VAR group were actually less variable than those from the YOKE group. Therefore, the current

experiment failed to replicate results from Experiment 1 and other studies where higher variability usually resulted from a condition where reinforcement was contingent on variability over multiple dimensions compared to conditions where variability was not required for reinforcement.

One possible reason for the inconsistent results obtained from the current experiment was in the difference in the nature of stimuli used and the way they were presented. The options for each dimension in the current experiment were visible to the participants; this enabled participants to formulate stereotypic responses to obtain reinforcers. Whereas in the rectangle drawing task and sequence generation tasks, usually the total available options are not available to the participants/animals, neither were these options visible; thus formulating a strategic response pattern to meet reinforcement contingency was almost impossible.

Another possible reason for the present finding was that the threshold value used in the current experiment was too permissive; a lower threshold value which required higher variability for reinforcement could have resulted in more variable responses in the VAR group. However, it was observed that some participants simply cycled through the options; increasing the variability requirement by using a lower threshold value would not have prevented this type of stereotypical responding to obtain high number of reinforcers. Using other schedules of reinforcement that have been used in the literature to increase variability such as the lag procedures, percentile schedules or the frequency dependent schedules also would not have prevented the cycling through options pattern. There might be methods to prevent all the strategic responses by imposing more rules on the reinforcement criteria to make random responses more likely; however, this could end up in participants making random responses based on the order of the options rather than the intended features of the options (e.g., colours, shapes and patterns). More importantly, choosing the options randomly might not be necessary for the dimensions of Colour, Shape and Pattern. When the dimensions under investigation changes, the methodologies will have to change to adapt to the nature of the new dimensions. The definition of variability used in the current experiment as well as Experiment 1 and Experiment 2, which was to

distribute responses evenly in a random manner, appeared to be inappropriate for defining variability in dimensions of Colours, Shape and Patterns.

U-values calculated based on the relative frequencies of available options (U-8) as a measure of variability have been shown to be limited as they did not capture the highly stereotypical responses observed in the data. Although the limitation in identifying higher order stereotypy has been noted (Machado, 1992; 1993; Page & Neuringer, 1985; Stokes, 1995), U-values based on relative frequencies of available options have continued to be used as a main measure of variability. U-values calculated based on the newly created categories consisting two consecutive trials (U-64) was shown to better represent the variability in the responses from our data. Because stereotypical responding would result in some sequences being more likely than others, U-64 will capture any stereotypy when it occurs, not restricting the types of stereotypy – the patterns of cycling through options or repeating the same option for a few times before moving to another will result in lower U-64 than if the responses were made randomly.

In order to capture the stereotypy in the responses observed in the current experiment, U-64 was created as an additional measure of variability that better captured the stereotypy. The method of calculating U-64 could also be utilised to create a modified threshold contingency for reinforcement that could be appropriate for the current task or for future investigations using similar tasks. The threshold contingency could be based on the relative frequencies of the newly created 64 categories, with each category consisting of two consecutive trials. If the relative frequency of this category is lower than the threshold value, the response could be reinforced. This new threshold contingency would encourage equal use of other sequences of two consecutive trials that are not from the cycling pattern or the pattern of repeatedly using one option before moving onto using another option.

Experiment 4

Introduction

Experiment 3 attempted to use the same methodologies as used in the rectangle drawing task to increase variability in a new task using completely independent dimensions (colour, shape and pattern). However, the results showed that the threshold contingency which successfully produced higher variability in the three dimensions of rectangles in Experiment 1 failed to produce the difference in variability in dimensions of colour, shape and pattern in that the variability as measured by U-values found no significant difference between the group receiving reinforcement for varying and the group receiving reinforcement independent of the variability in the options used for these dimensions. Moreover, it was likely because of the onscreen presentations of all the options available for each dimension, the threshold contingency produced stereotypical responding that resulted in high number of reinforcers.

One possible way to overcome the unintended stereotypy resulted from threshold contingency proposed previously was to use threshold contingency based on pairs of consecutive responses so that stereotypical responding patterns would not be reinforced more often than non-stereotypical responding. However, a question arose: Is choosing the options randomly among the dimensions of shapes, colours and pattern necessary? The task in Experiment 3 was selecting a shape and filling the shapes with colours and patterns. This differs from sequence generation tasks where usually 2 or 3 keys/keyboards (operanda) are used and the keys themselves have no distinguishing quality from each other except for the physical location. However, for the dimensions of shape, colour and pattern, each option is actually qualitatively different; the characteristics of the options matters.

Reinforcement schedules that emphasise choosing the options from a dimension randomly may result in participants simply moving from options to options without regard to the characteristics of each option, which is also important for the task. Therefore, randomness in the choices of shapes, colours and patterns appeared to be unnecessary. Take the dimension of Colour for example, if a child draws a picture using 12 colours, the variability in colours use will be considered to be higher than when a child draws a picture using 4 colours; whether the child used the different colours randomly did not matter. Alternative

reinforcement schedules that can increase the variability in the use of options in these dimensions is needed.

Reinforcing novel responses had been shown to increase response diversity in colours and shapes in children's drawings (e.g., Fallon & Goetz, 1975; Holman et al., 1977; Kratochwill et al., 1979) as well as in the forms constructed in block play in children (e.g., Goetz & Baer, 1973); it seems appropriate that it could also be used for the dimensions of Shape, Colour and Pattern that were used in Experiment 3. Although generalisation of learned variability across dimensions was the original goal for this thesis, it is considered important to first test the methodologies using only one dimension. The Colour dimension has many options readily available therefore it is suited to the investigation. Reinforcing the use of colours that have never been used by an individual may increase the range of colours used by the individual. This next experiment aimed to examine this.

The number of colours available for participants to choose from in Experiment 3 was only 8; and the results showed all colours were used equally often over 300 trials. The number of options available appeared to be too limited to allow for variability. On the other hand, however, having too many colour options (e.g., 256) would inflate variability in colours used simply because it is difficult to pick the same colours again. An intermediate number of 135 colours that allows for variability but does not artificially inflate variability was used here. Also, to show whether some colours are preferred over others, the number of options available has to be greater than the number of opportunities for choosing (the number of trials). For example, if there is a preference for certain colours, the proportion of trials using a new colour will be smaller than if there is no preference (e.g., using new colours in almost all the trials).

The present experiment therefore aimed at testing whether reinforcing the use of novel colours expanded the range of colours used by individuals in a computer task.

Methods

Participants

Ninety-three students (Male = 41, Female = 52) participated in this experiment; 21 were students from the University of Waikato, 25 from a local

high school and 47 from Ren Min University from China. Participants from first year psychology courses were given course credits for their participation. Ethics approval was obtained from the Research and Ethics Committee of School of Psychology at University of Waikato.

Apparatus

Experimental sessions were carried out with only one individual present at a time in a small office or with a small group of participants present at the same time in computer labs. Participants used a computer mouse to select colours on the computer screen and a sound was provided when the colour selected met the criterion in effect. Participants in the group sessions used headphones to prevent interference.

Procedures

Computer program

A computer program was created which provided 135 colour options (on the left of the screen) for participants to select from; these would be used to colour an on-screen white t-shirt (on the right of the screen). The number of t-shirts coloured and the number to be coloured for each phase was shown on the top left corner above the t-shirt. For Phase 2, a sticker was shown on the top right corner above the t-shirt, along with the number of stickers obtained next to the sticker. Appendix III shows the screen captures of the experiment.

Baseline – Phase 1

During Phase 1, participants were asked to colour blank t-shirts; no feedback was provided. The colours used were recorded by the computer and used as reference for reinforcement in Phase 2.

Reinforcing novel responses – Phase 2

During Phase 2, participants were given positive feedback in the form of simulated stickers (indicated by the increase in the sticker counter and an ascending tone) every time they used a colour which they had not used previously, in Phase 1 or in trials preceding the current trial in Phase 2.

Post-check – Phase 3

During Phase 3, as in Phase 1, no feedback was provided. The colours used were recorded by the computer.

Experimental procedures

On arrival, participants were given information sheets about the experiment and questions were answered. Consent forms were collected before the experiment started. The experiment started with the instructions being displayed on a light grey screen (see Appendix III for screen captures of the experiment):

“Scenario: Suppose you volunteer to help a charity organization create t-shirts to give people who donate to the organization.

Your task: You will help them create 220 t-shirts in total in three phases.

Overview of duties involved: First phase: create 60 t-shirts for back up (in storage). Second phase: create 100 for direct distribution – people will give you a sticker if they like your design. Final phase: create another 60 t-shirts for backup (in storage).”

Once the “next” button was clicked on the Phase 1 screen, a 15*9 colour palette of 135 colours appeared on the left side of the screen and a white box showing top half body of a model in a white t-shirt appeared on the right. On top of the colour palette there was a box showing the dimension “colour”; and on top and aligned to the left of the white box there was a counter showing the number of t-shirts completed and the total number of t-shirts needed for Phase 1 (0/60). Under the colour palette, there was a box “ok”, participants had to click on “ok” to indicate they had decided on a colour. When a colour was clicked on, the colour on the t-shirt in the white box will change to the selected colour accordingly and would remained on the t-shirt until another colour or “ok” was clicked on. Participants could change their choice of colour before they clicked on “ok” for each trial. Once “ok” was clicked on, the number on the left in the counter increased and the last colour before “ok” was clicked stayed on the t-shirt for 500ms. The colour of the t-shirt returned to white after 500ms elapsed and the next trial began. When all 60 t-shirts had been created in Phase 1, the instruction for Phase 2 appeared on the computer screen,

“Create 100 t-shirts for direct distribution. Create one t-shirt at a time. Feel free to change your selections, click ‘ok’ when done. Each t-shirt created will be produced immediately and handed to the person who had just donated. If the person likes the t-shirt, s/he will give you a sticker. You will hear a sound (“ding”) when someone likes the t-shirt you just created. The number next to the sticker will show you how many of your t-shirts have been liked. ”

The same screen with colour palette and white box with the model in white t-shirt appeared, the number of t-shirts to be completed in the counter box increased to 100 and there was a sticker and a sticker counter “00” on top of the t-shirt box. Each time a participant used a colour – confirmed with the “ok” button – that s/he had not used before, the counter to the right of the sticker would increase. Once the participants had completed 100 t-shirts, they would be reminded on the screen with the number of stickers they received,

“Well done! You’ve received n stickers! End of phase 2. Please click next when you are ready for phase 3.”

Instructions and the number of t-shirts to be completed for Phase 3 were the same as Phase 1. No feedback was provided.

The experiment finished when participants completed all three phases.

Measures

The number of different colours used in a phase: the number of different colours used in a phase alone.

The number of new colours used over all phases: the number of different colours used throughout the three phases.

The number of colours used in Phase 3 which had not been used in Phase 1: the number of different colours used in Phase 3 that had not been used in Phase 1. It did not matter whether these colours have been used in Phase 2.

Results and Discussions

The question we are attempting to answer in this experiment is whether reinforcing the use of different colours will increase the use of colours after reinforcement is withheld. That is, whether the colours participants used in Phase 3 (post-check) differed from the ones they used in Phase 1 (baseline). Before we look into that, we first looked at whether the reinforcement contingency had successfully increased the number of colours used in Phase 2. Figure 4.1 shows the number of different colours used in each phase (dark bars), the number of new colours (light grey bars) and the number of colours used in Phase 3 that was new to those used in Phase 1 (dark grey bar).

Baseline – Phase 1

In Phase 1, participants completed the t-shirts without receiving any feedback on the colours used. The colours and the number of colours used were recorded as reference for reinforcement in Phase 2 and as a baseline. On average, participants used 35.2 colours (as shown by the first bar in Figure 4.1; $SD = 16.04$). The number of colours used ranged from 1 to 59; therefore, there appeared to be great differences across participants (also indicated by the large standard deviation). Overall, participants used a new colour on only half of the trials; this should reflect a stable colour preference for colouring the t-shirts.

Reinforcing novel responses – Phase 2

In Phase 2, participants received reinforcers when they used colours that they had not used before in any previous trials including those in baseline (Phase 1). As shown by the two bars for Phase 2 in Figure 4.1, on average, participants used 66.12 (out of 135 possibilities) different colours within the 100 trials ($SD = 16.94$), ranging from 9 to 97. Among these colours, on average 46.11 ($SD = 17.64$) were new to Phase 1. This suggests that overall, the reinforcement contingency successfully expanded the choices of colours participants used for colouring the t-shirts. It should be noted however, there was great differences among participants as indicated by the large standard deviation and the range.

Post-check – Phase 3

In Phase 3, participants again completed colouring 60 t-shirts without receiving any feedbacks on the colours they used. The three bars for Phase 3 in Figure 4.1 show the mean numbers of colours used. On average, participants used

40.11 different colours ($SD = 15.36$); 26.68 of these colours were new to Phase 1 ($SD = 13.12$). There were on average 10.19 ($SD = 6.19$) colours new to Phase 1 and Phase 2. When compared to the number of colours used in Phase 1 ($M = 35.20$, $SD = 16.04$), on average, participants used more colours in Phase 3, $t(92) = -3.30$, $p = .001$, $r = .32$. That is, participants' choices of colours for t-shirts were wider after they had been reinforced to use different colours.

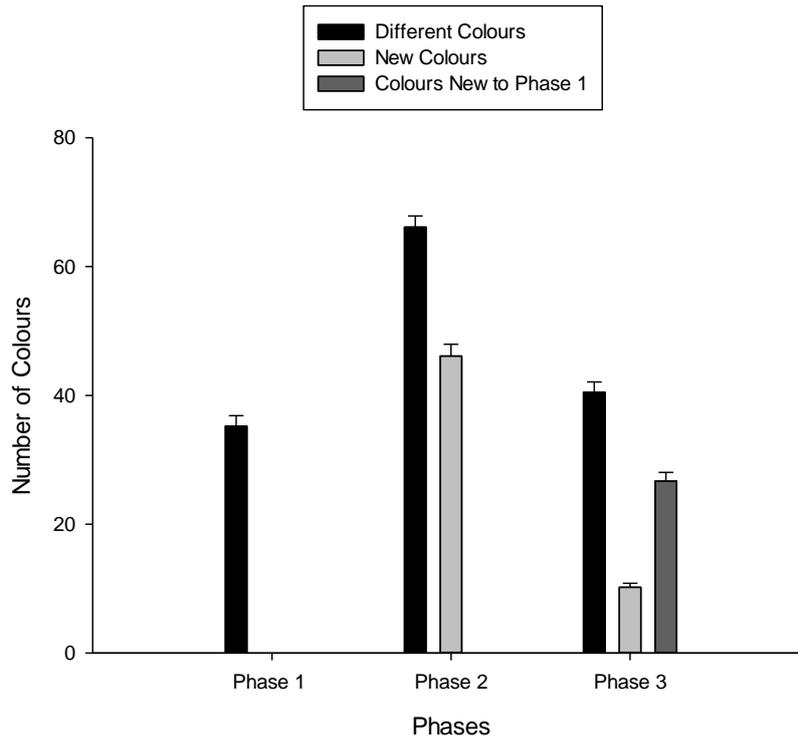


Figure 4.1. Mean number of colours used in Phase 1, Phase 2 and Phase 3. Number of different colours were represented by dark bars, number of new colours represented by light grey bars and the number of colours in Phase 3 that were new to Phase 1 was in dark grey bar. Error bars represent standard errors.

Post-check – Phase 3

In Phase 3, participants again completed colouring 60 t-shirts without receiving any feedback on the colours they used. The three bars for Phase 3 in Figure 4.1 show the mean numbers of colours used. On average, participants used 40.11 different colours ($SD = 15.36$); 26.68 of these colours were new to Phase 1 ($SD = 13.12$). There were on average 10.19 ($SD = 6.19$) colours new to Phase 1 and Phase 2. When compared to the number of colours used in Phase 1 ($M = 35.20$, $SD = 16.04$), on average, participants used more colours in Phase 3, $t(92) =$

-3.30, $p = .001$, $r = .32$. That is, participants' choices of colours for t-shirts were wider after they had been reinforced to use different colours.

Throughout the analysis of the number of colours used for the three phases, it became apparent that there were substantial differences between participants.

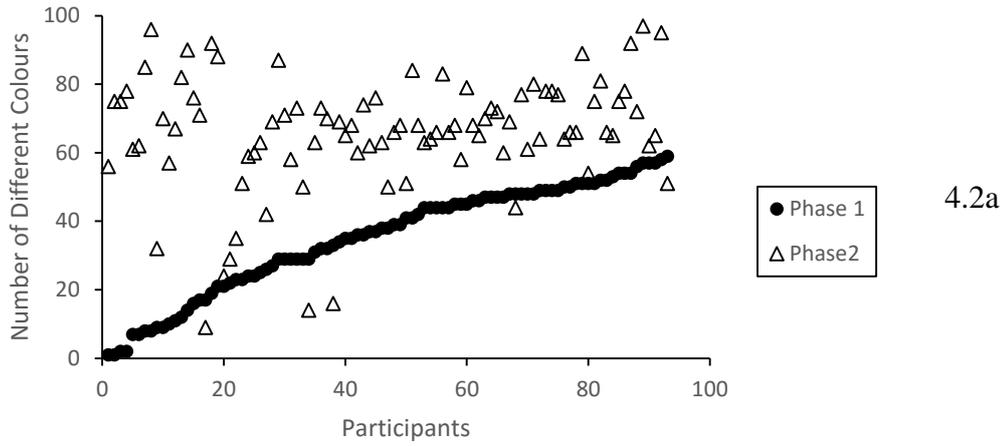
For example, the minimum number of colours used in Phase 1 was 1 while the maximum was 59; one participant used only one colour while one used different colours for almost all 60 trials. This was true for Phase 2 and Phase 3. Figure 4.2 shows the number of colours used in three phases by all participants. The number of colours used in baseline (Phase 1) were rearranged in an ascending order; the number of colours used in other phases corresponded to these orders accordingly. As shown in Figure 4.2, participants who used very few colours in Phase 1 used as many in Phase 2 and Phase 3 as those who used many colours in Phase 1 (panel 4.2a and 4.2b). Although it was found that, on average, participants used more colours in Phase 3 than in Phase 1, data shown in Figure 4.2 suggested that some participants started low and increased largely during (Phase 2) and after (Phase 3) reinforcement phase while some participants used the maximum number of colours in both baseline and after reinforcement phase. To clarify this, further analyses were carried out.

Groupings of data and re-analysis

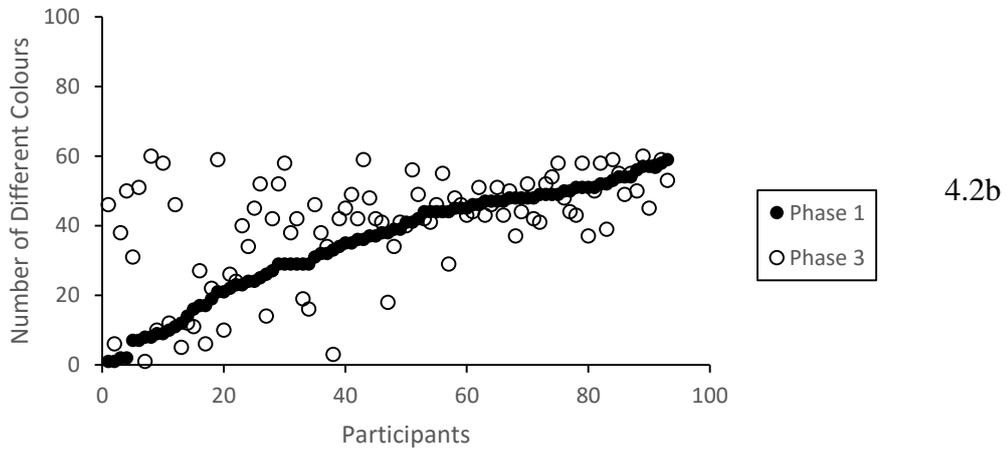
In order to examine more closely the effect that reinforcement had on the change in colours used among participants differing in the number of colours used over the baseline, the data from all participants were re-arranged to create three groups. The number of colours used in Phase 1 were arranged in an ascending order; and data for all other phases changed accordingly. The first 31 participants with the lowest numbers formed the first group – referred to as group “Bottom”, the second 31 with medium numbers formed the second group – referred to as group “Middle” and the last 31 with the highest numbers formed the third group – referred to as group “Top”. Data were then reanalysed based on these three groups.

Figure 4.3 shows the mean number of colours used in all three phases for each of the newly created groups. Overall, for the number of different colours used, there appeared to be little change between Phase 1 and Phase 3 for groups Middle and Top while there was an increase for group Bottom – participants who

The number of colours used in Phase 1 and Phase 2



The number of colours used in Phase 1 and Phase 3



The number of colours used in Phase 3 but not in Phase 1

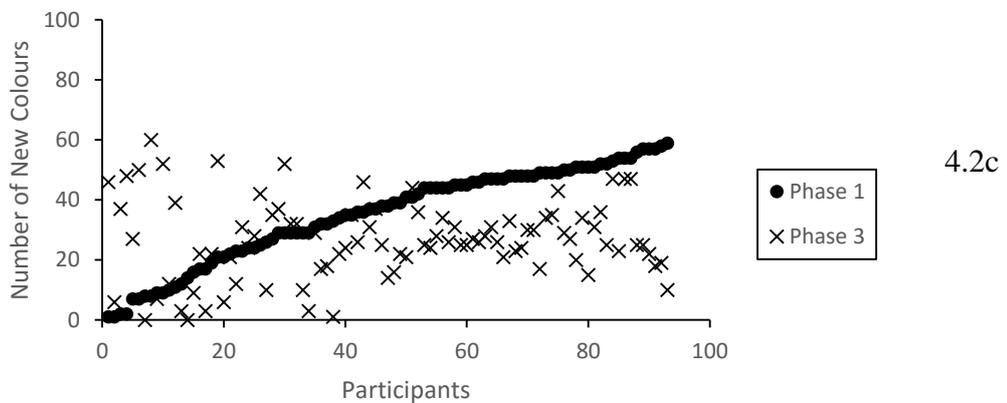


Figure 4.2. The number of colours used in Phase 1 and Phase 2 (panel 4.2a), in Phase 1 and Phase 3 (panel 4.2b.) and the number of colours used in Phase 1 and the number of colours used in Phase 3 but not in Phase 1; panel 4.2c). Plotted in order of increasing use of the number of colours in Phase 1.

used the lowest number of colours during baseline. The number of colours used in Phase 2 appeared to be similar across groups.

In Phase 1, as can be seen in the first set of bars in Figure 4.3, group Bottom used the lowest number of colours, group Middle used a much higher number of colours and group Top used the most. The variability in the numbers of colours used in Phase 1 appeared to be bigger for the Bottom group than the other two groups as indicated by the standard errors. Analysis of variance showed a main effect of groups on the number of colours used in Phase 1, $F(2, 54.26) = 225.31, p < .001$, partial $\eta^2 = .81$; Leven's test showed that the assumption of homogeneity was violated (Leven statistic $(2, 90) = 18.92, p < .001$) therefore, the Welch's test statistics is reported instead. Post Hoc analysis using LSD indicated that the number of colours used in Phase 1 was lower for the Bottom group ($M = 15.90, SD = 9.02$) than both the Middle group ($M = 38.39, SD = 5.54; p < .001$) and the Top group ($M = 51.32, SD = 3.64; p < .001$); the number was lower for the Middle group than it was for the Top group ($p < .001$).

At baseline, the Bottom group used a significantly smaller number of colours than the Middle and Top groups; and the Middle group used significantly fewer colours than the Top group. These differences were expected as participants were grouped based on the number of colours used in Phase 1.

In Phase 2, the number of different colours used (as shown by the middle sets of bars in Figure 2) did not appear to be different between the Bottom and the Middle group; and it was highest for the Top group. Results from an ANOVA showed that the difference across the groups was only marginally significant, $F(2, 57.42) = 3.33, p = .043$, partial $\eta^2 = .05$. Leven's test indicated that the assumption of homogeneity had been violated, therefore, the Welch test statistics is reported. Post hoc analysis found no significant difference in the number of colours used across groups Bottom ($M = 63.55, SD = 21.37$), Middle ($M = 63.32, SD = 15.23$) and Top ($M = 71.48, SD = 12.14$). These results show that under the reinforcement contingency, participants who used a small number of colours in baseline (Phase 1) performed similarly as participants who had used many colours in baseline.

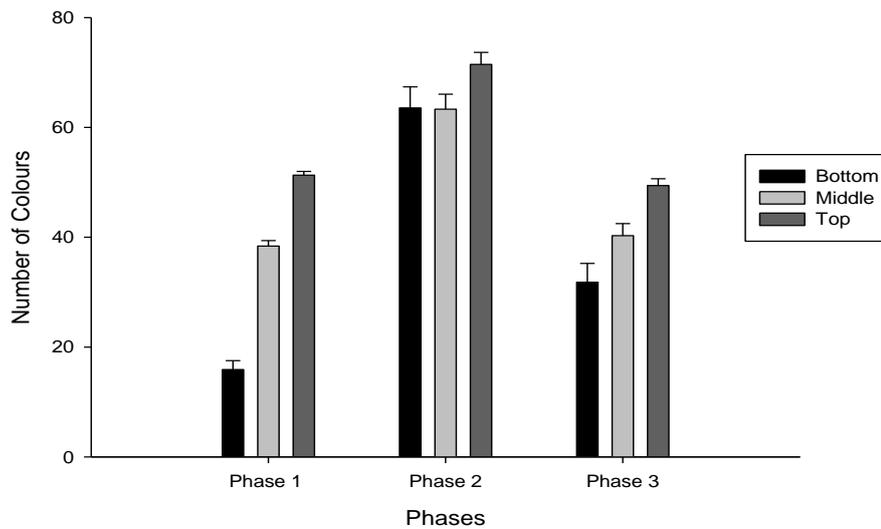


Figure 4.3. Mean number of different colours used in the three phases for groups of participants using least number of colours (Bottom), intermediate number of colours (Middle) and highest number of colours (Top) in Phase 1. Error bars represent standard errors.

In Phase 3 (the last set of bars in Figure 4.3), the number of colours used by the three groups shows a trend which is similar to that in Phase 1: the Bottom group used the least number of colours, the Middle group used a medium number and the Top group used the most colours. However, the difference between the Bottom group and the other two groups decreased dramatically compared to Phase 1. Results from an ANOVA showed that there was a significant main effect of groups on the number of colours used in Phase 3, $F(2, 52.27) = 15.47, p < .001$, partial $\eta^2 = .22$. Leven's test indicated that the assumption of homogeneity had been violated (L statistic (2, 90) = 18.62, $p < .001$), therefore, the Welch test statistics was reported. Post hoc analysis by LSD indicated that the number of colours used was significantly lower for the Bottom group ($M = 31.81, SD = 19.11$) than either the Middle group ($M = 40.29, SD = 12.25; p = .044$) or the Top group ($M = 49.42, SD = 6.94; p < .001$); and the number was significantly lower for the Middle group than it was for the Top group ($p = .027$). This shows that participants who used the least number of colours in Phase 1 also used the least number of colours in Phase 3; this was also true for the other two groups. Recall that in Figure 4.2 there appeared to be an increase in the number of colours used for the Bottom group between Phase 1 and Phase 3, so although participants started and finished using lowest number of colours, their performance appeared to have improved.

Figure 4.4 re-plotted some values from Figure 4.3 to better show the differences in the number of colours used between Phase 1 and Phase 3 for each of the three groups. As shown in Figure 4.4, the number of colours used increased from Phase 1 to Phase 3 only for the Bottom group; and the values varied greatest among participants in the Bottom group. Paired-samples t-tests were carried out for the number of colours used between Phase 1 and Phase 3 for each of the Bottom, Middle and Top groups. For the Bottom group, on average, participants used significantly less different colours in Phase 1 than in Phase 3, $t(30) = -4.43$, $p < .001$, $r = 0.63$. No significant difference was found for the Middle ($t(30) = -.96$, $p = .35$, $r = 0.17$) or the Top group ($t(30) = 1.71$, $p = .098$, $r = 0.30$). These results suggest that the increase in colour choices for t-shirts in this experiment was significant only for participants who started using small number of colours but not for those who started with using medium to high numbers of colours.

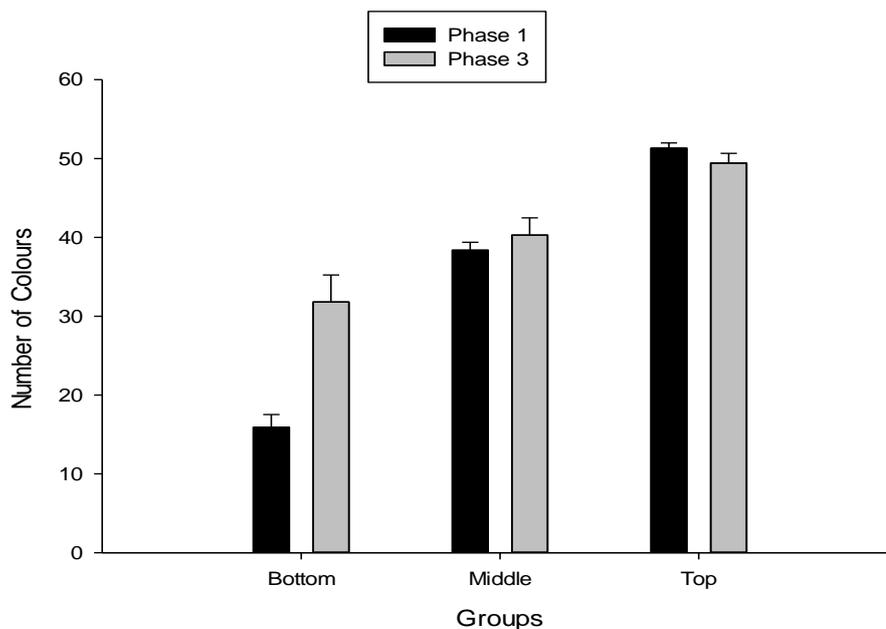


Figure 4.4. Mean number of colours used in Phase 1 (dark bars) and in Phase 3 (light grey bars) for Bottom, Middle and Top groups. Error bars represent standard errors.

While overall the results suggest there was an impact of reinforcement on variability, however, the impact was shown to be evident for only one group of participants after the data was broken into three groups with different levels of initial variability. This analysis shows that, for participants who already had maximum variability, it was impossible for reinforcement to increase variability. There were at least two possible reasons for the high number of colours used for

the Top group in Phase 1. One was that preferences for colours for T-shirts was wide for this group of participants. Another possibility was that participants from this group was designing for a diverse market so they used many different colours; however these colours did not reflect their personal preferences for colours. Either way, in practice, the test for control by reinforcement on variability is demonstrated only for participants who had low level of initial variability (the Bottom group). For the Top group, reinforcement obviously cannot have an impact because there was no opportunity for the variability to increase because it was already maximum. While there's an outcome overall, the more important outcome is the analysis of the difference in the Bottom group where the measure can vary and this shows control by reinforcement.

Summary

The present experiment investigated whether the newly created computer task that reinforced the use of novel colours increased variability in colour use on that task. Results showed that the diversity in colour choices increased only for people who initially used a relatively limited number of colours; however such increases were not observed for those who used a wide range of colours initially. Analyses that combine results from individuals with different initial levels of variability can result in false representation of the impact of reinforcement on variability. Therefore, for experimental designs like the present experiment, it is important to assess whether the impact of reinforcement is different for individuals with different level of initial variability.

An Analysis of U as a Measure of Variability

Introduction

U-value is a measure of uncertainty (Miller & Frick, 1949) which shows the overall variability of all emitted responses. The U-value as described earlier has been most commonly used in studies to assess the level of variability in responses. U-value is a molar measure of variability – it shows only the overall response distributions of responses, but does not reflect the order of the sequences of those responses at a molecular level (Maes, 2003; Page & Neuringer, 1985; Stokes, 1995). Results from Experiment 3 in the current thesis clearly demonstrated this limitation. Although the limitation of U-value being not able to capture the sequences of responses had been recognized and additional measures had been used (Maes, 2003; Stokes, 1995), little analysis has been done to understand U-values and the impacts of changes in the number of categories/options used on the changes in the values of U. It is also unclear what level of variability is represented by the different values of U (e.g., is variability in responses of the same set of options always the same for a value of .8?).

As described earlier, the U-value is calculated according to the following formula:

$$U - value = - \sum_{i=1}^{\beta} \frac{\alpha_i \times \log(\alpha_i)}{\log(\beta)}$$

In the formula, β equals the number of possible categories and α equals the relative frequency of category i . In essence, greater U-values would reflect greater variability in the responses as compared to lower U-values.

Although the formula for calculating the U-value used across the different studies are slightly different (e.g., using log or log2), the calculation is always based on the relative frequencies of all available options; and the results from using the different formula are algebraically equivalent. The U-values show whether the available options have been used equally often. Most research using U-values as a measure of variability has reported the U-values and has compared variability in responses between groups or conditions based on the U-values obtained. However, there is no detailed information about how much variability a particular value of U reflects, put differently, is there a linear relation between the

U-values and amount of variability (e.g., how variable responses are with a U-value of 0.8 or 0.9)? Also, U provides no rational to distinguish “high” from “low” variability. For example, we might conclude that the variability in the distributions of responses across categories or options with a U-value of .9 to be higher than that with a U-value of .8 based on the difference in the values. However, it is not clear whether the variability in responses that resulted in a lower value of U is necessarily lower than variability in responses that resulted in a higher value of U.

The current section aimed at examining U-values as a measure of variability in more details by using two sets of simulated data. The first simulation, which consisted of 100 responses distributed over 4 categories, shows possible stereotypical response patterns that would result in high U-values. The second simulation presents a more detailed analysis of the relationship between U-value and response patterns and consisted of 300 responses distributed evenly across a number of categories.

Simulation 1 – Stereotypical Responding

In Experiment 3, highly repetitive or stereotypical responding resulted in high U-values; the first set of simulation was designed to explore ways that frequent stereotypical responding could result in high U-values.

Four types of response distributions were simulated. For the first type, 100 responses were distributed randomly over four categories (Random); for the second, 100 responses were distributed over four categories evenly in a way that 25 responses were made to one category before moving to another category (e.g., 111111...222222...333333...444444...; Repeat); for the third type, 100 responses were again evenly distributed over four categories, but this time in a cycling pattern (e.g., 1234123412341234...; Cycle); and for the last type, all 100 responses were distributed to only one category and zero to the other three (One).

U-values shown in Figure 5.1 were calculated based on these four types of simulations. Also shown in Figure 5.1 are the frequencies of responses to each category for the four types of simulation. As can be seen in Figure 5.1, U-values were maximum for the Repeat and Cycle types of simulation and minimum for the One type, where only one category was used; although all of these three types were highly stereotypical. For the Random type, frequencies of responses to each

category were not strictly even and therefore a U-value of .98 rather than 1 was obtained. Another 10 sets of 100 responses were generated randomly over 4 categories and the resulting U-values ranged from .95 to .99; random responses never reached the highest possible U-value.

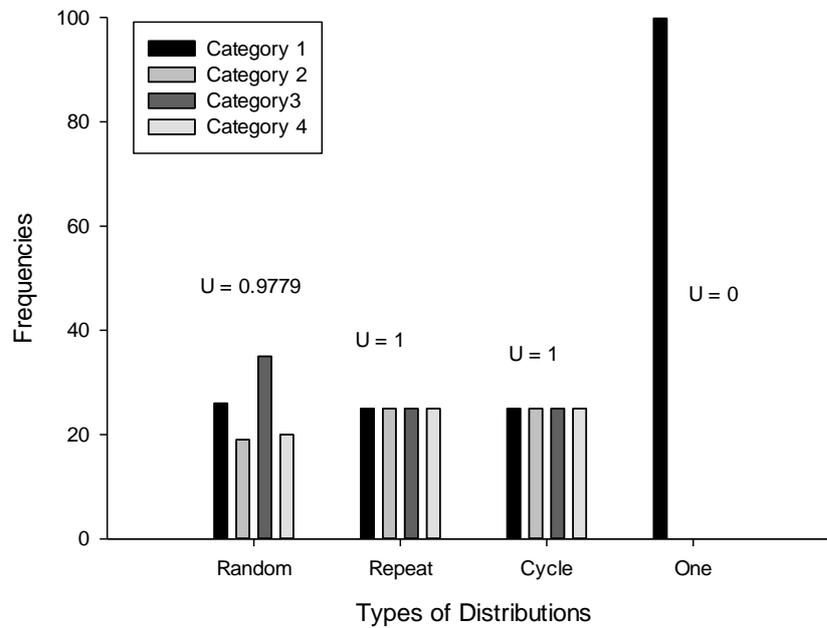


Figure 5.1. U-values from four types of simulation of 100 responses over 4 categories and Frequencies of responses to each category from each type of simulation.

Both of the response patterns, Repeat and Cycle, ensure equal number of responses to all 4 categories which result in a high U-value of 1. However, the response patterns are highly stereotypical, which is not what one might intuit from a high U-value. For some tasks, for example, rectangle drawing (Ross & Neuringer, 2002), it is rather difficult to generate such high level of stereotypical response patterns because the categories were continuous and unknown to participants; however, this could be a problem for tasks in which categories are discrete and visible to the participants as in Experiment 3. Therefore, using the U-value as an index of variability has a potential limit: such as when U-values show us variability of one aspect of responses, such as the evenness of responses over categories used, it is not sufficient as an indication of response variability because it does not capture the sequences of the responses.

Reinforcement of variability in an attempt to rule out the possibility of higher order stereotypy (e.g., repeating one option before moving onto another

one and cycling through all options) can be difficult. For example for response options of only four categories, Lag procedures and the frequency dependent methods by themselves cannot prevent the Cycle type response pattern; and threshold contingency used by itself will allow both Repeat and Cycle types of stereotypy. Using a combination of at least two types of schedules, for example threshold contingency and lag procedures reduces the chance of the simplest stereotypical responding, but is likely to create just another type of higher order stereotypy that meets the reinforcement criteria, especially when experimenting with humans where elements of verbal regulation can occur. U-value would be a useful measure when the response options are categorical or discrete (not continuous); however scheduling contingencies have to be arranged so that only truly random responding is reinforced.

Simulation 2 – U-values from using different number of the available options

The second simulation consisted of different sets of 300 responses over 16 possible categories; each set differed in the number of categories used out of the 16. The number of categories being used ranged from one to 16. For example, if 4 categories were used, each of these four categories were allocated 75 responses and zero response was allocated to all other 12 categories. For categories that had not been used, the lag of 0 normally used is replaced by 0 only to enable calculation. Note that responses to the used categories were equally distributed so that the maximum U- values could be obtained for the number of categories used. Seven additional sets of 300 responses were simulated but instead of having zero response to the unused categories, 1, 2, and 3 responses were allocated to each of the unused categories (others = 1, others = 2, others = 3). The purpose of having the additional sets of simulations was to explore the different impacts on U-value when one or more categories were completely unused and when all categories have been used only a small number of times.

U-values were calculated based on the different sets of simulations and plotted in Figure 5.2 and Figure 5.3.

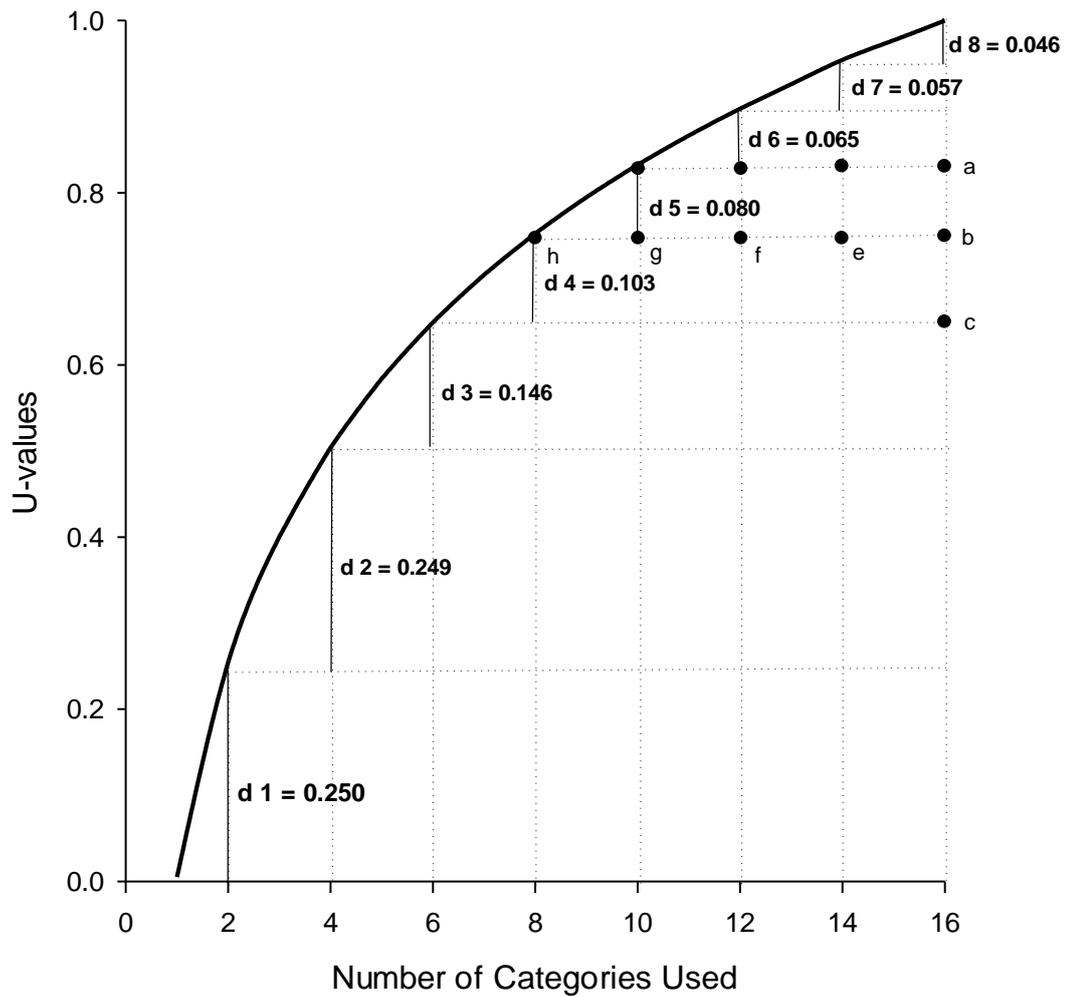


Figure 5.2. U-values resulting from 300 responses over different number of categories out of 16 used.

Figure 5.2 shows the relationship between U-values (y-axis) and the number of categories used (x-axis) when no responses were made to the unused categories (others = 0). Values on the curve corresponding to the number of categories used represented the maximum U-value one can get from evenly distributing the 300 responses over all the used categories, with no responses made at all to the unused categories.

When all categories were used, 16 in this case, and when the responses were evenly distributed, the maximum U-value is 0.999 which reflects the highest variability possible. The U-value from random responses is 0.993, very close to the highest value. U-values decreased when the distributions of responses over these 16 categories were less even. For example, judging by the values of U for

points *a*, *b*, and *c* in Figure 5.2, one could conclude that responses were less evenly distributed over the 16 categories; and responses resulted in U-value corresponding to point *a* were more evenly distributed than those corresponding to point *b* and point *c*, with those corresponding to point *c* being the least evenly distributed. This would be true when comparisons were made for situations when the same number of categories were used; the higher the U-value, the more evenly distributed responses will be.

When only half of the available categories were used, 8 in this case, even if the responses were evenly distributed (maximum variability for 8 categories being used), the highest value U can reach is 0.753 (point *h* in Figure 5.2). A U-value of 0.753 can also be resulting from using more than 8 categories (e.g. 9, 10, 11... 15 and 16). While U-value of 0.753 is maximum when 8 categories are used, it is not the maximum value for using 10 (point *g*), 12 (point *f*), 14 (point *e*) and 16 (point *b*) categories. This makes comparing variability of responses difficult as it appears that with the same value of U, the level of variability can be different, depending on the number of categories being used. Therefore, it is worth noting that before one makes comparisons between groups or subjects based on U-values, one should take into account the number of categories used in addition to reporting only U-values.

One possible way of comparing the evenness of distributions over different number of categories is to look at the distance between the obtained U-value to the maximum U-value one can obtain for using particular number of categories. For example, the U-value for point *g* (using 10 categories) is 0.753 and the maximum U-value for 10 categories is 0.832, thus the distance for point *g* to the maximum U-value is 0.080. By using the same calculation, the distance to maximum U-values for point *f* (12 categories) is 0.145, 0.202 for point *e* (14 categories) and 0.247 for point *b* (16 categories). We already discussed that for the same number of categories used, the greater distance from the maximum U-value, the less even the distributions will be. Thus we can assume that the evenness for responses will be greatest for point *h* and least for point *b*, even though the number of categories used for point *h* is 8 and but 16 for point *b*.

As we can see from Figure 5.2, the maximum U-values increased as the number of categories used increased; however, the increase was non-linear. For

example, with the increase in the number of categories used, the differences in the U-values decreases, even though the difference in the number of categories used remains the same.

Data from Experiment 1 fitted into Figure 5.2 (Simulation 2)

Figure 5.3 added U-values from both the VAR and YOKE groups on Area, Shape and Location dimensions from Experiment 1 to show the impacts of using different number of categories on U-values, as well as the different response patterns resulted from the differential reinforcement schedules. As can be seen in Figure 5.3, participants from the VAR group tended to use all 16 categories for all three dimensions; and the distances between these obtained U-values and the maximum U-values which can be obtained using 16 categories were relatively small. Participants from the YOKE group (empty circles) tended not to use all 16 categories; and the number of categories used varied greatly among participants, ranging from 6 to 16. The distances between the obtained U-values and the maximum U-values using the same number of categories were relatively large. This indicates that participants from the YOKE group used fewer categories; and responses over these relatively small number of categories were uneven. Given the information on the number of categories used and how evenly responses were distributed over these used categories, U-value is more informative when comparing the variability in responses between these two groups.

SUMMARY

Simulation 1 showed stereotypical ways of responding that resulted in an extremely high U-value; Simulation 2 showed comparisons of U-value have to take into account the number of available options used. When U-value might provide useful information about overall variability in responses, analyses of U-value using the two simulations in this section revealed that it is also a limited measure of variability. For one, it does not capture stereotypical patterns in responses; for the second, it is ambiguous as there is no specific values distinguishing between high and low variability; and for the third, it is ambiguous because same values of U could result from different level of evenness in response distributions, depending on the number of options used. For these

reasons, alternative measures of variability should be considered when studying the effect of reinforcement on response variability.

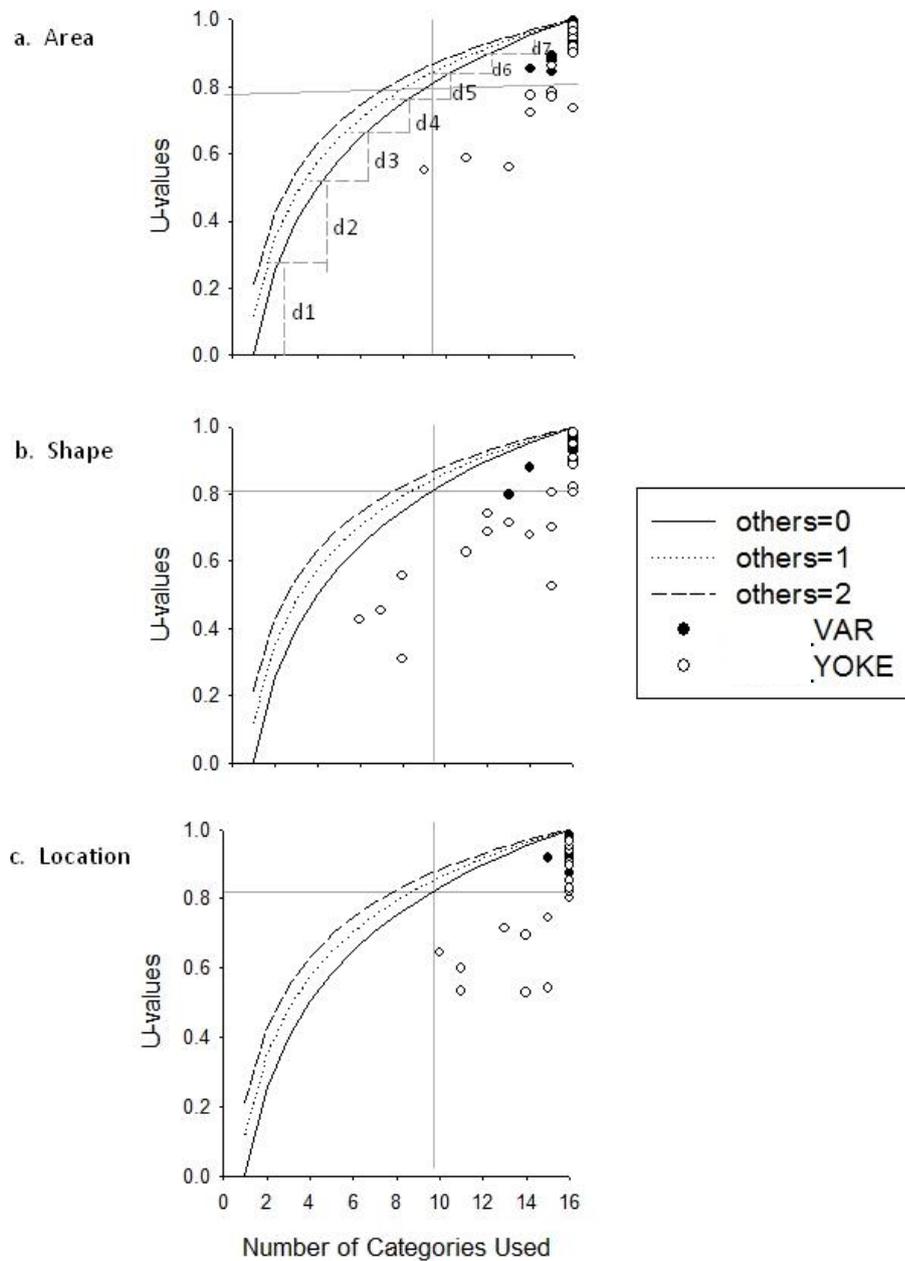


Figure 5.3. U-value from Experiment 1 fitted in the figure with simulation data of 300 responses over 16 categories. Filled circles represent data from VAR group and empty circles represent those from the YOKE group.

General Discussion

Overall aim of thesis

A series of experiments were carried out with the initial aim of examining the effects of variability in responses on one or more dimensions of that response on the variability of other dimensions of the same response.

Summary of the experiment findings

In the first experiment, it was demonstrated that reinforcement can generate variability in multiple dimensions when participants were reinforced to draw rectangles of different sizes, shapes and at different locations on the computer screen. In the second experiment, it was demonstrated that the unreinforced dimension had lower variability than the two reinforced dimensions using the rectangle drawing tasks with different participants. When data from Experiment 1 and Experiment 2 were compared, there appeared to be some generalisation of learned variability from the reinforced to the unreinforced dimensions. However, close inspection of Experiment 2 suggested that there was some confound between the impact of reinforcing variability on two dimensions having an impact on variability of a third dimension and the non-orthogonal nature of the three dimensions. That is, when the reinforced dimension was made highly variable, it inevitably led to variability in the unreinforced dimension. Thus it was unclear whether it was the non-orthogonality of dimensions or the positive contingency itself which created the higher variability in the unreinforced dimension.

Experiment 3 was designed to investigate the non-orthogonality confound found in the nature of the task used in Experiment 2; and to this end, the task was changed to selecting shapes and filling the shapes with colours and patterns. However, there was a failure to observe higher variability in all three dimensions in the group receiving reinforcement for variability; the variability was high in all three dimensions regardless whether the individuals received reinforcers for varying or not. It was suggested that the obvious and visible nature of the variations in the dimensions of the task might have been responsible for this as they allowed easy categorisation. The final experiment was a more naturalistic investigation carried out in response to the limitation of using too few options for

a dimension in Experiment 3; more options were used to reflect variability in responses. The task was now to vary colour use over many (135) shades. The reinforcement schedule and measure of variability were also changed to suit the new task. It was found that reinforcing the use of novel colours increased the number of colours used by individuals who initially used only a small number of colours.

Reinforcing behavioural variability

Experiment 1 successfully replicated Ross and Neuringer's (2002) finding of the reinforcement control of the variability over multiple dimensions, and results from Experiment 1, 2 and 4 were consistent with findings from other research on reinforcing variability in general (e.g., Cherot et al., 1996; Denney & Neuringer, 1998; Goetz & Baer, 1973; Machado, 1992; 1993; Page & Neuringer, 1985). Despite the disagreements among researchers on whether variability should be considered an operant dimension (Barba, 2012; Holth, 2012; Machado & Tonneau, 2012; Marr, 2012; Neuringer, 2012), results from Experiment 1, 2 and 4 in the current thesis, using threshold contingencies for reinforcement and reinforcing novel responses, showed that variability in the dimensions where variability was directly reinforced was higher than when it was not. The results of the present experiments give support to position that the level of variability in responses is under the control of the contingency.

Variability vs. stereotypy

Experiment 3 attempted to move the dimensions in the rectangle drawing tasks where the dimensions were non-orthogonal to more independent dimensions (shape, colour and pattern) while using the same reinforcement schedule used in Experiment 1 and 2. However, the reinforcement contingency that aimed at reinforcing highly variable responses actually resulted in highly stereotypical response patterns. Responses from participants receiving reinforcement contingent on variability had responded more stereotypically than those who received reinforcement independent of it. It was possible that the task used in Experiment 3 was too easy, with a small number of options and the options being visible, so that participants easily found rules that lead to reinforcement for every response but not to responding variably. The ability to develop stereotypical patterns of

responding to meet variability requirement is not unique to humans. For example, Holth (2012) reported that rats learned to cycle through four response options to meet a Lag 3 criteria which led to reinforcement on every response. However, the stereotypical pattern was broken when five response options were used and a Lag 4 criteria was imposed. The number of options used in Experiment 3 was 8; and for humans this might be like the four response options enabling easy formulation of stereotypical response patterns observed in rats in Holth's (2012) study.

One way to prevent such stereotypy as found in Experiment 3 would be to impose more stringent reinforcement criteria (e.g., lowering the threshold value, incorporating more requirements or using criteria with multiple requirements) to prevent stereotypical responding being reinforced. However, because of the nature of the dimensions used in this experiment, as well as the tendency of humans to formulate rules in obtaining reinforcement, training participants to respond randomly might lead them to focus on the locations of each option rather than the physical characteristics of them (e.g., colours, shapes, patterns). Research in reinforcing behavioural variability in sequence generation with rats and pigeons usually assume highly variable responses to be similar to random responses. For sequence generation, locations of the keys are important as the generation of sequences involves moving between locations or staying at the same location; subjects do not have to pay attention to the characteristics (e.g., colours, textures) of the keys or levers in order to get reinforcement. However, for a task like the one we used in Experiment 3 where characteristics of the dimensions actually mattered more than locations – the changes in location should be the result of choosing different options (e.g., colours). Therefore, for dimensions such as colours, patterns and shapes, randomness in the locations of the elements of these dimensions might not be necessary when defining highly variable responses.

Generalisation

Generalisation was examined by comparing the variability in dimensions where variability was not reinforced contingently between Experiment 1 and Experiment 2 using the rectangle drawing task. For two out of three dimensions, reinforced variability in two dimensions appeared to have generalised to the third dimension (the Shape and Location dimensions). This is consistent with findings from Maes (2003) and da Silva Souza and Abreu-Rodrigue (2010) where they

found that the variability in sequences generated by students to be higher if variability has been previously reinforced. Results from Maes (2003) and da Silva Souza and Abreu-Rodrigue (2010) showed that learned variability of the same task has likely generalised from one point of time to a later time; results from the first two experiments in the current thesis showed that such generalisation may also generalise across dimensions of a same response.

The possible generalisation of reinforced variability, however, was not found for the Area dimension; the variability of Area dimension was at a similarly intermediate level regardless whether the other dimensions it occurred with were reinforced to vary or not. One possible explanation was that when the reinforcement contingency was unclear, participants responded to the three dimensions differently; they naturally varied more on the Area dimension than the other two. As a result, even if there was enhanced variability in the Area dimension as the result of generalisation from reinforcing variability in responses on the other two dimensions, the difference did not look large. Another possible explanation was that varying on some of the dimensions may result in inevitable increase in the variability of responses in another dimension (non-orthogonality of the dimensions). Therefore, the increased level of variability in the unreinforced dimension may not be the result of generalisation but simply the result of dependency among these three dimensions. Both of the explanations point to the fact that these three dimensions are non-orthogonal; and such non-orthogonality prevented a solid conclusion of generalisation of learned variability across dimensions.

Interdependency across dimensions of a response is a major difficulty faced in the current project when examining the generalisation of learned variability. This is likely to be faced also by other researchers intending to investigate the generalisation of reinforced variability across dimensions of the same response because the dimensions will often correlate naturally. One example is that when reinforcing a higher speed of running (speed measured by dividing the distance by time), it will inevitably result in longer strides and higher lifts – the length of strides and the height of the leg lift are part of the running process and could change not because of the reinforcement contingency but because of the attempt to increase the speed of running. One line of research examines the

distances between dimensions in psychological space so that one can examine how close or distant the dimensions of interest are (e.g., Soto & Wasserman, 2010). However, the distances between dimensions in psychological space would differ greatly from individual to individual based on their personal history; it is therefore difficult to make conclusions based on group data. When the interdependency of different dimensions of a same response or behaviour cannot be controlled for as a confounding variable, investigations on the generalisation of learned variability might be more appropriate when examined across independent responses.

U-value as a measure of variability

This thesis examined the extent to which U-values, a measure of variability that is commonly used in the literature, represented response variability. U-value has been commonly used to show overall variability of responses; however, little analysis has been done to see how the values of U are affected by factors other than the evenness of the distributions of responses. The limitation of U-value for showing only overall response distributions but not the sequences of responses has long been recognised (e.g., Maes, 2003; Neuringer, 2012; Page & Neuringer, 1985; Stokes, 1995); simulation 1 of the analysis of U in the current thesis demonstrated two highly stereotypical patterns that could result in the highest U-value ($U = 1$).

Another concern for using U-value as a measure of response variability is that the values of U are ambiguous; a same value can be the results of response distributions with different level of evenness. Barba (2012) pointed out that the calculation of U-values that is based on the complete set of possible response options might not be appropriate when only some of the options are used. The second simulation of the analysis of U confirmed his concern; responses resulted in the same value of U, for example, .75, can be the result of distributing responses completely evenly across 8 categories or less evenly across any number of categories that is greater than 8 – the more categories used, the less even responses will be. Therefore, comparing the variability across response sets with only the values of U and without information regarding the number of options each used is meaningless. Even if the information on the number of options used was provided, however, which is more or less variable is difficult to determine.

For example, is it more variable to use more options but less evenly, or is it more variable to use fewer options but more evenly? In addition, when there are large number of options unused (e.g., only 8 categories out of 16 are used), the maximum U-value can appear to be quite low (e.g., the U-value for distributing responses evenly over 8 categories out of 16 was below .8). This becomes a concern when there is a large number of possible options without enough opportunities for responding, particularly when the number of options is greater than the number of total response required. For example, a sequence of 8 responses over three keys on a computer keyboard has 6561 possible sequences, if only 3000 responses, which is large enough for human participants, even if all 3000 sequences are different, the U-value will still be quite low.

Future research in reinforcing variability intending to use U-value as a measure of variability should use additional measures that can capture possible stereotypy in responding. Also, when comparing the variability between individuals based on U-values, information in regard to the number of options used should be included. Finally, for tasks that has a large number of possible options but not a large enough number of response opportunities, U-value is not an appropriate measure.

Implications

One area where increased variability in responses is desired is in individuals with autism. Individuals on the autistic spectrum typically have repetitive response patterns and display rigid responding in social interactions; these characteristics are believed to have prevented them from getting into contact with reinforcement in the social contexts where variability in responding is reinforced (Leekham, Prior, & Uljarevic, 2011; Rodriguez & Thompson, 2015). There has been a growing number of research in applying reinforcement schedules that increase variability in functional responding in individuals with autism; lag procedures has been the most commonly employed method (e.g., Esch et al., 2009; Lee et al., 2002; Murray & Healy, 2013; Susa & Schlinger, 2012). Studies using lag schedules to increase variability typically used low lag requirement to sustain responding; however it has been shown that lag schedules with a low requirement result in stereotypical responding (e.g., Lee et al., 2002; Lee & Sturmey, 2006; Page & Neuringer, 1985; Susa & Schlinger, 2012).

One of the concerns raised from the results from Experiment 3 in the current thesis was that schedules intending to reinforce variable responses resulted in highly stereotypical responding that was not reflected by the measure used. It was likely due to the small number of options available and human's tendency to form rules to obtain reinforcers. This result translates to the autism research is that measures used in these studies might not pick up stereotypical responding. Although higher lag requirements in Susa and Schlinger (2012) has been shown to prevent alternating responding in the individual, it was not clear whether the individual had cycled through the four responses. For research in areas where only small number of response options are available, additional measures which can detect stereotypy in responding are needed in addition to counting the number of different responses.

Another area where results from the current thesis can be extended to is investigations in increasing variability in responding in typically developing individuals so that creative outcomes are more likely (Epstein, 1999; Goetz, 1989; Neuringer, 2003; Runco, 1993; Stokes, 2012).

One of the defining characteristics of creativity is novelty. A creation or behaviour usually has to be novel within a certain period of time or within a certain society to be considered creative. However, it is also important that responses or creations to be novel within an individual's history before novelty to the wider society to occur. Experiment 4 showed that for people who initially had a limited colour selections, their choices of colours widened after being reinforced to use colours they had never used previously in the experiment. This result is consistent with research that reinforced the use of novel shapes and colours in paintings and drawings (Fallon & Goetz, 1975; Holman et al., 1977; Kratochwill et al., 1979; Ryan & Winston, 1978) as well as those that reinforced variability in block-building play (Goetz & Baer, 1973). Experiment 4 used the method of reinforcing novel responses to increase variability in responses, which is an approach that has not been used recently. The success in using this task to increase the number of colours use by individuals who initially used only few colours provides a potentially viable methodology for future investigation in increasing creative behaviours. Difficulties involving subjectivity and the immediacy and accuracy of reinforcement faced by early studies of reinforcing

variability that is related to creativity can be overcome by using a computerised task like the one used in Experiment 4.

Conclusion

Results from the current thesis support to the notion that variability in responses can be reinforced. The attempt to examine whether such trained levels of variability generalise was not achieved because the confounding non-orthogonality of the three dimensions of the rectangle drawing task could not be ruled out. Although the current thesis did not answer the question raised at the beginning, results from our four experiments point out key aspects future investigation into generalisation should pay attention to. These are: to have a clear functional definition of variability and use appropriate reinforcement schedules accordingly; to choose appropriate measure(s) for the specific dimensions chosen; to examine whether higher order stereotypical responding patterns occurred for certain dimensions or tasks when required; and finally, effort should be made to identify independent dimensions when examining the impact of generalisation on the level of variability in dimensions or responses that are not trained. Current thesis attempted to study whether trained level of variability in responses would generalise across different dimensions of a same behaviour. However, it was realised that the test for generalisation across dimension will be hindered because of the interdependent nature of simultaneous dimensions of a response. Although not being answered by the current thesis, research aimed at understanding whether trained level of variability generalise across behaviours, tasks and contexts has practical significance.

References

- Arnold-Saritepe, A. M., Phillips, K. J., Mudford, O. C., De Rozario, K. A., & Taylor, S. A. (2009). Generalization and maintenance. In J. L. Matson (Ed.), *Applied Behavior Analysis for Children with Autism Spectrum Disorders*. New York: Springer.
- Arnesen, E. M. (2000). *Reinforcement of object manipulation increases discovery*. Reed College, Unpublished Thesis.
- Baldwin, J. D., & Baldwin, J. I. (1978). Reinforcement theories of exploration, play, creativity, and psychological growth. In E. O. Smith (Ed.), *Social play in primates* (pp. 231-258). New York: Academic Press.
- Barba, L. S. (2012). Operant variability: a conceptual analysis. *The Behavior Analyst, 35*, 213-227.
- Baruni, R. R., Rapp, J. T., Lipe, S. L., & Novotny, M. A. (2014). Using lag schedules to increase toy play variability for children with intellectual disabilities. *Behavioral Interventions, 29*, 21-35.
- Betz, A. M., Higbee, T. S., Kelley, K. N., Sellers, T. P., & Pollard, J. S. (2011). Increasing response variability of mand frames with script training and extinction. *Journal of Applied Behavior Analysis, 44*, 357-362.
- Blough, D. S. (1966). The reinforcement of least frequent interresponse times. *Journal of the Experimental Analysis of Behavior, 9*, 581-591.
- Brown, J. S., Bilodeau, E. A., & Baron, M. R. (1951). Bidirectional gradients in the strength of a generalized voluntary response to stimuli on a visuospatial dimension. *Journal of Experimental Psychology, 41*, 52-61.
- Bryant, D., & Church, R. M. (1974). The determinants of random choice. *Animal Learning & Behavior, 22*, 245-248.
- Campbell, D. T. (1960). Blind variation and selective retention in creative thought as in other knowledge processes. *Psychological Review, 67*, 380-400.
- Cherot, C., Jones, A., & Neuringer, A. (1996). Reinforced variability decreases with approach to reinforcers. *Journal of Experimental Psychology: Animal Behavior Processes, 22*, 497-508.
- Cohen, L., Neuringer, A., & Rhodes, D. (1991). Effects of ethanol on reinforced variations and repetitions by rats under a multiple schedule. *Journal of the Experimental Analysis of Behavior, 54*, 1-12.
- Da Silver Souza, A., Abreu-Rodrigues, J., & Baumann, A. A. (2010). History effects on induced and operant variability. *Learning & Behavior, 38*, 426-437.
- Denney, J., & Neuringer, A. (1998). Behavioral variability is controlled by discriminative stimuli. *Animal Learning & Behavior, 26*, 154-162.
- Doolan, K. E., & Bizo, L. A. (2013). Reinforced behavioral variability in humans. *The Psychological Record, 63*, 725-734.
- Doughty, A. H., Giorno, K. G., & Miller, H. (2013). Effects of reinforcer magnitude on reinforced behavioral variability. *Journal of the Experimental Analysis of Behavior, 100*, 355-369.
- Doughty, A. H., & Lattal, K. A. (2001). Resistance to change of operant variation and repetition. *Journal of the Experimental Analysis of Behavior, 76*, 195-215.
- Duker, P. C., & van Lent, C. (1991). Inducing variability in communicative gestures used by severely retarded individuals. *Journal of Applied Behavior Analysis, 24*, 379-386.

- Edelstein, B. A. (1989). Generalization: terminological, methodological and conceptual issues. *Behavior Therapy*, 20, 311-324.
- Epstein, R. (1996). *Cognition, creativity, and behavior: Selected essays*. Westport, CT: Praeger.
- Epstein, R. (1999). Generativity theory. In M. A. Runco & S. R. Pritzker (Eds.), *Encyclopedia of Creativity* (Vol. 1, pp. 759-766). California, USA: Academic Press.
- Esch, J., Esch, B., & Love, J. R. (2009). Increasing vocal variability in children with autism using a lag schedule of reinforcement. *The Analysis of Verbal Behavior*, 25, 73-78.
- Fallon, M. P., & Goetz, E. M. (1975). The creative teacher: effects of descriptive social reinforcement upon the drawing behavior of three preschool children. *School Applications of Learning Theory*, 7, 27-42.
- Gharib, A., Derby, S., & Roberts, S. (2001). Timing and the control of variation. *Journal of Experimental Psychology: Animal Behavior Processes*, 27, 165-178.
- Gharib, A., Gade, C., & Roberts, S. (2004). Control of variation by reward probability. *Journal of Experimental Psychology: Animal Behavior Processes*, 30, 271-282.
- Goetz, E. M. (1989). The teaching of creativity to preschool children: the behavior analysis approach. In J. A. Glover, R. R. Ronning & C. R. Reynolds (Eds.), *Handbook of Creativity*. New York: Plenum Press.
- Goetz, E. M., & Baer, D. M. (1973). Social control of form diversity and the emergence of new forms in children's blockingbuilding. *Journal of Applied Behavior Analysis*, 6, 209-217.
- Goetz, E. M., Jones, K., & Weamer, K. (1973). *The generalization of creativity "training" in easel painting to blockbuilding*. Paper presented at the Annual Convention of the American Psychological Association, Montreal, Canada
- Goode, M. K., Geraci, L., & Roediger III, H. L. (2008). Superiority of variable to repeated practice in transfer on anagram solution. *Psychonomic Bulletin & Review*, 15, 662-666.
- Grunow, A., & Neuringer, A. (2002). Learning to vary and varying to learn. *Psychonomic Bulletin & Review*, 9, 250-258.
- Holman, J., Goetz, E. M., & Baer, D. M. (1977). The training of creativity as an operant and an examination of its generalization characteristics. In B. C. Etzel, J. LeBlanc & D. M. Baer (Eds.), *New developments in behavioral research: Theory, method and application*. Hillsdale, N.J.: Lawrence Erlbaum Associates.
- Holth, P. (2012). Variability as an operant? *The Behavior Analyst*, 35, 243-248.
- Hopkinson, J., & Neuringer, A. (2003). Modifying behavioral variability in moderately depressed students. *Behavior Modification*, 27, 251-264.
- Koehler-Platten, K., Grow, L., Schulze, K. A., & Bertone, T. (2013). Using a lag reinforcement schedule to increase phonemic variability in children with autism spectrum disorders. *The Analysis of Verbal Behavior*, 29, 71-83.
- Kratochwill, C. E., Rush, J. C., & Kratochwill, T. R. (1979). The effects of descriptive social reinforcement on creative responses in children's easel painting. *Studies in Art Education*, 20, 29-39.

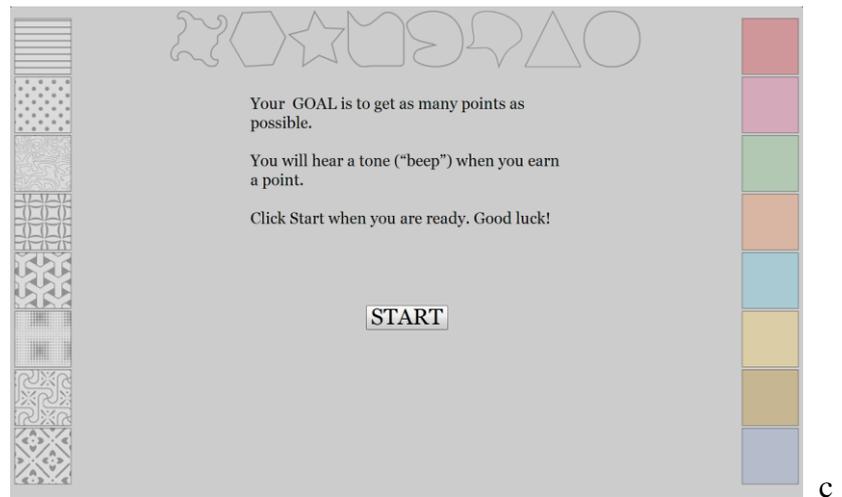
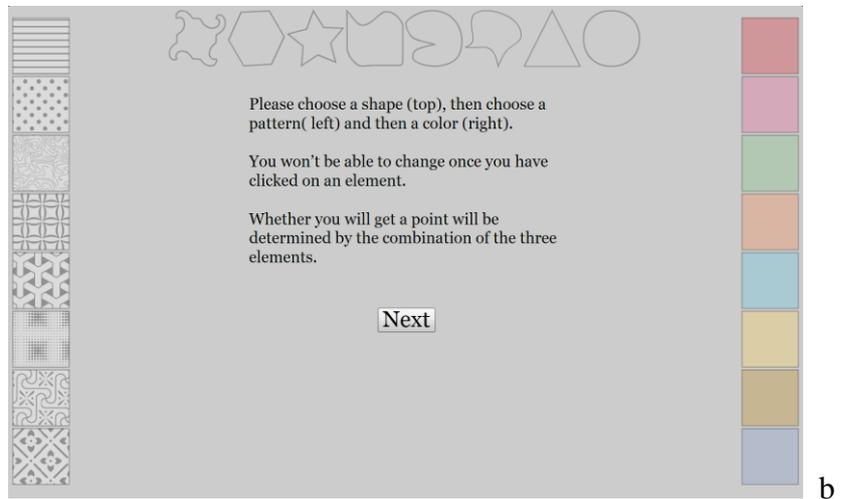
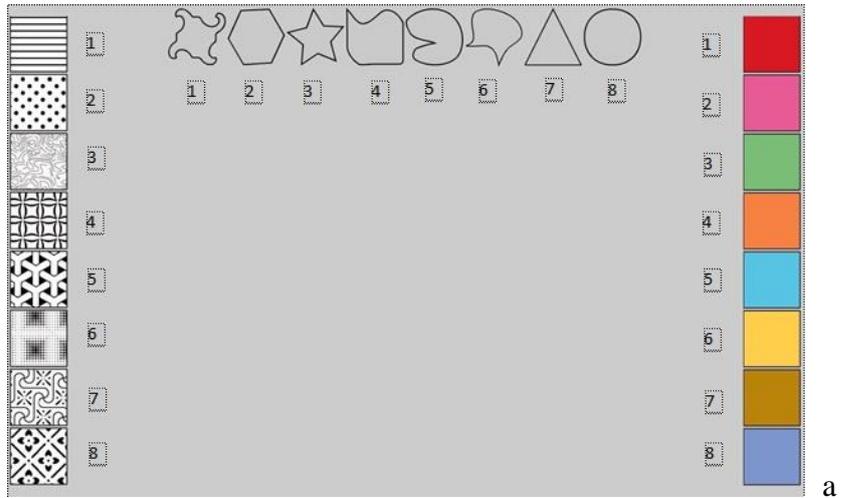
- Lee, R., McComas, J. J., & Jawor, J. (2002). The effects of differential and lag reinforcement schedules on varied verbal responding by individuals with autism *Journal of the Applied Behavior Analysis*, *35*, 391-402.
- Lee, R., Sturmey, P., & Fields, L. (2007). Schedule-induced and operant mechanisms that influence response variability: A review and implications for future investigations. *The Psychological Record*, *57*, 429-455.
- Leekham, S. R., Prior, M. R., & Uljarevic, M. (2011). Restricted and repetitive behaviors in autism spectrum disorders: A review of research in the last decade. *Psychological Bulletin*, *137*, 562-593.
- Leising, K. J., Ruprecht, C. M., & Stahlman, W. D. (2014). Modulation of variation by response-reward spatial proximity. *International Journal of Comparative Psychology*, *27*, 326-337.
- Machado, A. (1989). Operant conditioning of behavioral variability using a percentile reinforcement schedule. *Journal of the Experimental Analysis of Behavior*, *52*(155-166)
- Machado, A. (1992). Behavioral variability and frequency-dependent selection. *Journal of the Experimental Analysis of Behavior*, *58*, 241-263.
- Machado, A. (1993). Learning variable and stereotypical sequences of responses: Some data and a new model. *Behavioural Processes*, *30*, 103-130.
- Machado, A. (1997). Increasing the variability of response sequences in pigeons by adjusting the frequency of switching between two keys. *Journal of the Experimental Analysis of Behavior*, *68*, 1-25.
- Machado, A., & Tonneau, F. (2012). Operant variability: procedures and processes. *The Behavior Analyst*, *35*, 249-255.
- Maes, J. H. R. (2003). Response stability and variability induced in humans by different feedback contingencies. *Learning & Behavior*, *31*, 332-348.
- Maes, J. H. R., & van der Goot, M. (2006). Human operant learning under concurrent reinforcement of response variability. *Learning and Motivation*, *37*, 79-92.
- Manabe, K., Staddon, J. E. R., & Cleaveland, J. M. (1997). Control of vocal repertoire by reward in budgerigars (*Melopsittacus undulatus*). *Journal of Comparative Psychology*, *111*, 50-62.
- Marr, M. J. (2012). Operant variability: some random thoughts. *The Behavior Analyst*, *35*, 237-241.
- Miller, G. A., & Frick, F. C. (1949). Statistical behavioristics and sequences of responses. *Psychological Review*, *56*, 311-324.
- Miller, N., & Neuringer, A. (2000). Reinforcing variability in adolescents with autism. *Journal of Applied Behavior Analysis*, *33*, 151-165.
- Mook, D. M., Jeffrey, J., & Neuringer, A. (1993). Spontaneously Hypertensive Rats (SHR) readily learn to vary but not repeat instrumental responses. *Behavioral and Neural Biology*, *59*, 126-135.
- Morris, C. J. (1989). The effects of lag value on the operant control of response variability under free-operant and discrete-response procedures. *Psychological Record*, *39*, 263-270.
- Murray, C., & Healy, O. (2013). Increasing response variability in children with autism spectrum disorder using lag schedules of reinforcement. *Research in Autism Spectrum Disorders*, *7*, 1481-1488.
- Napolitano, D. A., Smith, T., Zarcone, J. R., Goodkin, K., & McAdam, D. B. (2010). Increasing response diversity in children with autism. *Journal of Applied Behavior Analysis*, *43*, 265-271.

- Neuringer, A. (1986). Can people behave "randomly"? The role of feedback. *Journal of Experimental Psychology: General*, *115*, 62-75.
- Neuringer, A. (1993). Reinforced variation and selection. *Animal Learning & Behavior*, *21*, 83-91.
- Neuringer, A. (2002). Operant variability: Evidence, functions, and theory. *Psychonomic Bulletin & Review*, *9*, 672-705.
- Neuringer, A. (2003). Creativity and reinforced variability. In K. A. Lattal & P. N. Chase. (Eds.), *Behavior theory and philosophy* (pp. 323-338). New York: Plenum Publishing.
- Neuringer, A. (2004). Reinforced variability in animals and people: Implications for adaptive action. *American Psychologist*, *59*, 891-906.
- Neuringer, A. (2012). Reinforcement and induction of operant variability. *The Behavior Analyst*, *35*, 229-235.
- Neuringer, A., Deiss, C., & Olson, G. (2000). Reinforced variability and operant learning. *Journal of Experimental Psychology: Animal Behavior Processes*, *11*, 429-452.
- Neuringer, A., & Huntley, R. W. (1991). Reinforced variability in rats: effects of gender, age and contingency. *Physiology & Behavior*, *51*, 145-149.
- Neuringer, A., & Jensen, G. (2012). The predictably unpredictable operant. *Comparative Cognition & Behavior Reviews*, *7*, 55-84.
- Neuringer, A., & Jensen, G. (2013). Operant variability. In G. J. Madden (Ed.), *APA Handbook of Behavior Analysis: Methods and Principles* (Vol. 1).
- Newman, B., Reinecke, D. R., & Meinberg, D. L. (2000). Self-management of varied responding in three students with autism. *Behavioral Interventions*, *15*, 145-151.
- Odum, A. L., Ward, R. D., Barnes, C. A., & Burke, K. A. (2006). The effects of delayed reinforcement on variability and repetition of response sequences. *Journal of the Experimental Analysis of Behavior*, *86*, 159-179.
- Page, S., & Neuringer, A. (1985). Variability is an operant. *Journal of Experimental Psychology: Animal Behavior Processes*, *11*, 429-452.
- Parsonson, B. S., & Baer, D. M. (1978). Training generalized improvisation of tools by preschool children. *Journal of Applied Behavior Analysis*, *11*, 363-380.
- Pryor, K. W., Haag, R., & O'Reilly, J. (1969). The creative porpoise: Training for novel behavior. *Journal of the Experimental Analysis of Behavior*, *12*, 653-661.
- Rodriguez, N. M., & Thompson, R. H. (2015). Behavioral variability and autism spectrum disorders. *Journal of Applied Behavior Analysis*, *48*, 1-21.
- Ross, C., & Neuringer, A. (2002). Reinforcement of variations and repetitions along three independent response dimensions. *Behavioural Processes*, *57*, 199-209.
- Runco, M. A. (1993). Operant theories of insight, originality, and creativity. *The American Behavioral Scientist*, *37*, 54-67.
- Ryan, B. A., & Winston, A. S. (1978). Dimensions of creativity in children's drawings: A social-validation study. *Journal of Educational Psychology*, *70*, 651-656.
- Schmidt, R. A., & Bjork, R. A. (1992). New conceptualizations of practice: Common principles in three paradigms suggest new concepts for training. *Psychological Science*, *3*, 207-217.

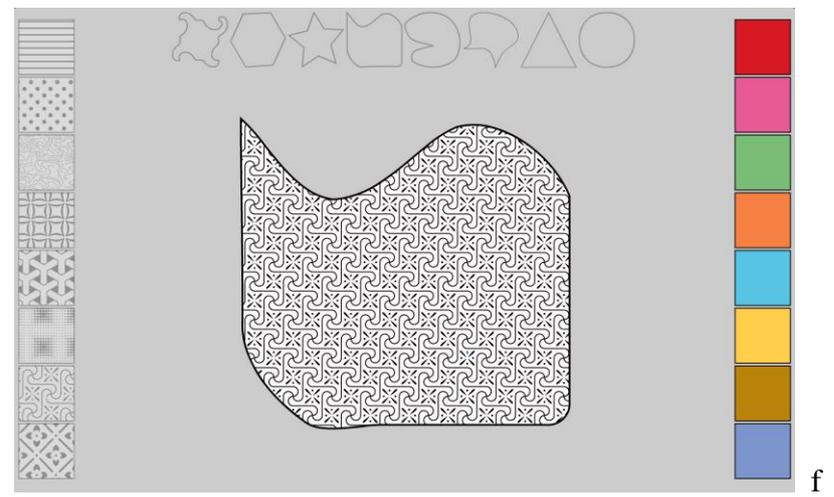
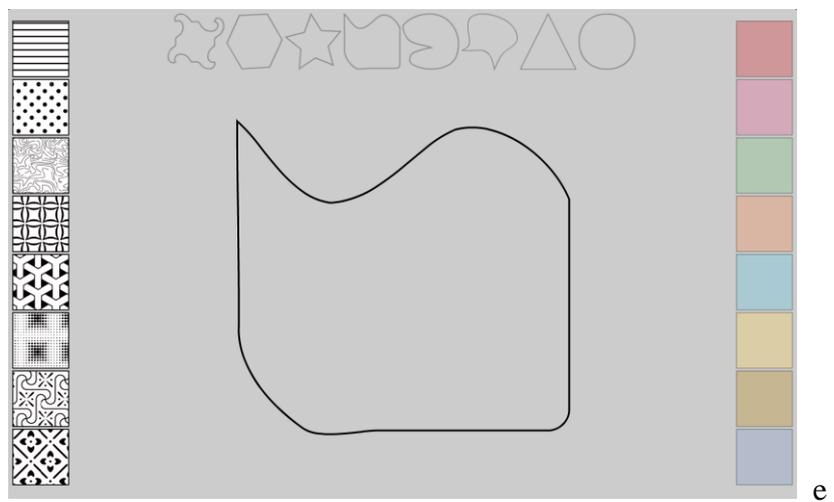
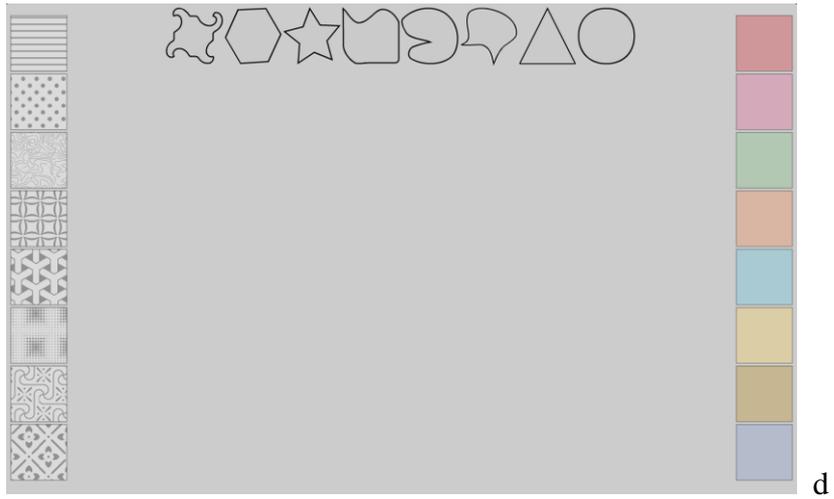
- Schoenfeld, W. N., Harris, A. H., & Farmer, J. (1966). Conditioning response variability. *Psychological Reports, 19*, 551-557.
- Schusterman, R. J., & Reichmuth, C. (2008). Novel sound production through contingency learning in the pacific walrus (*Odobenus rosmarus divergens*). *Animal Cognition, 11*, 319-327.
- Schwartz, B. (1982). Failure to produce response variability with reinforcement. *Journal of the Experimental Analysis of Behavior, 37*, 171-181.
- Shimp, C. P. (1967). Reinforcement of least-frequent sequences of choices. *Journal of the Experimental Analysis of Behavior, 10*, 57-65.
- Skinner, B. F. (1938). *The behavior of organisms: an experimental analysis*. Oxford, England: Appleton-Century.
- Skinner, B. F. (1953). *Science and human behavior*. New York: Macmillan.
- Soto, F. A., & Wasserman E. A. (2010). Integrality/separability of stimulus dimensions and multidimensional generalisation in pigeons. *Journal of Experimental Psychology: Animal Behavior Processes, 36*, 194-205.
- Stahlman, D., Leising, K. J., Garlick, D., & Blaisdell, A. P. (2013). There is room for conditioning in the creative process: associative learning and the control of behavioral variability. In O. Vartanian, A. S. Bristol & J. C. Kaufman (Eds.), *Neuroscience of creativity*. Cambridge, MA: The MIT Press.
- Stahlman, W. D., Roberts, S., & Blaisdel, A. P. (2010). Effect of reward probability on spatial and temporal variability. *Journal of Experimental Psychology: Animal Behavior Processes, 36*, 77-91.
- Stahlman, W. D., Young, M. E., & Blaisdel, A. P. (2010). Response variability in pigeons in a Pavlovian task. *Learning & Behavior, 38*, 111-118.
- Stokes, P. D. (1995). Learned variability. *Animal Learning & Behavior, 23*, 164-176.
- Stokes, P. D. (2012). Creativity and operant research: Selection and reorganization of responses. In M. A. Runco (Ed.), *The Creativity Research Handbook* (Vol. II, pp. 147-171). New York: Hampton Press.
- Stokes, P. D., Holtz, T. M., Carlis, A., & Eisenberg, J. (2008). Sources of variability in children's problem solving. *The International Journal of Creativity & Problem Solving, 18*, 49-67.
- Stokes, P. D., Mechner, F., & Balsam, P. D. (1999). Effects of different acquisition procedures on response variability. *Animal Learning & Behavior, 27*, 28-41.
- Stokes, T. F., & Baer, D. M. (1977). An implicit technology generalization. *Journal of Applied Behavior Analysis, 10*, 349-367.
- Stokes, T. F., & Osnes, P. G. (1989). An operant pursuit of generalization. *Behavior Therapy, 20*, 337-355.
- Susa, C., & Schlinger, H. D. (2012). Using a lag schedule to increase variability of verbal responding in an individual with autism. *The Analysis of Verbal Behavior, 28*, 125-130.
- Ward, R. D., Bailey, E. M., & Odum, A. (2006). Effects of d-amphetamine and ethanol on variable and repetitive key-peck sequences in pigeons. *Journal of the Experimental Analysis of Behavior, 86*, 285-305.
- Ward, R. D., Kynaston, A. D., Bailey, E. M., & Odum, A. (2008). Discriminative control of variability: Effects of successive stimulus reversals. *Behavioural Processes, 78*, 17-24.

- Weiss, A., & Neuringer, A. (2012). Reinforced variability enhances object exploration in shy and bold rats. *Physiology & Behavior, 107*, 451-457.
- Wolfe, K., Slocum, T. A., & Kunnavatana, S. S. (2014). Promoting behavioral variability in individuals with autism spectrum disorders: A literature review. *Focus on Autism and Other Developmental Disabilities, 29*, 180-190.

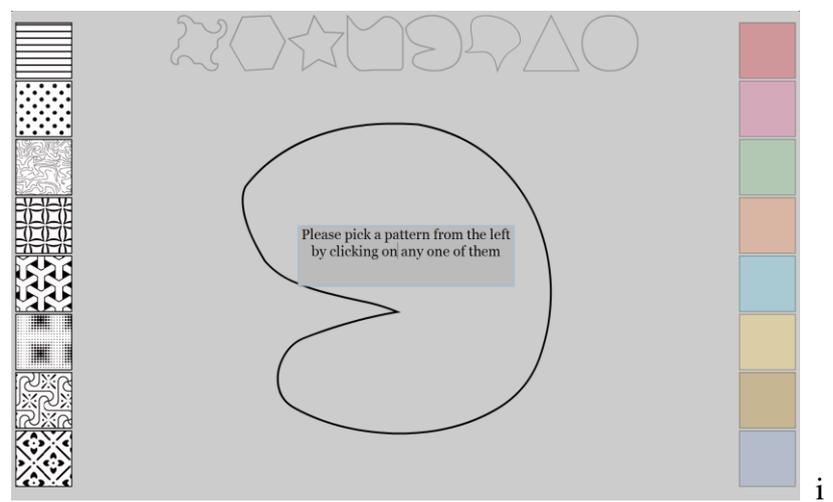
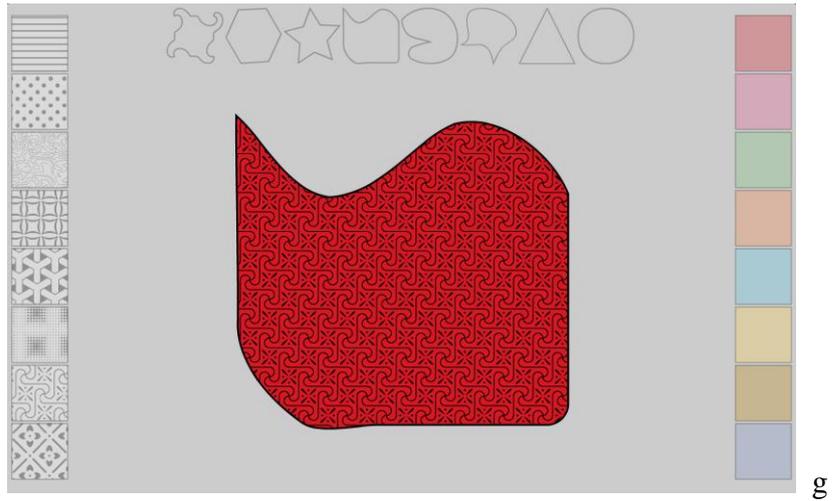
Appendix I



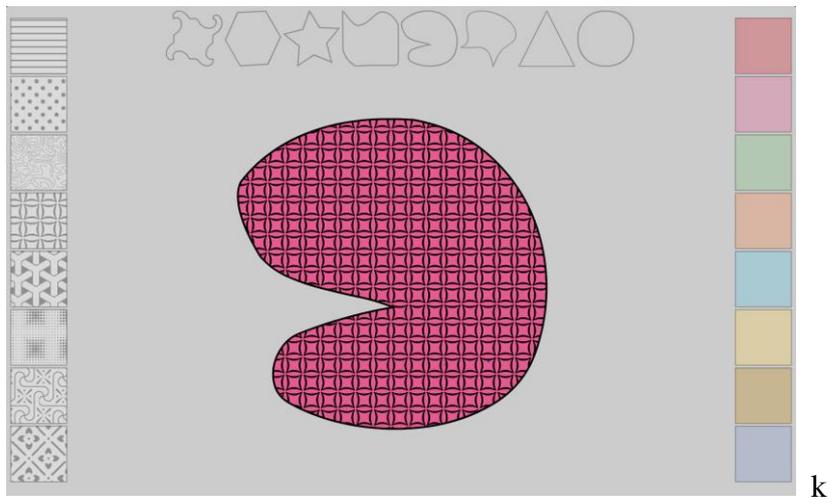
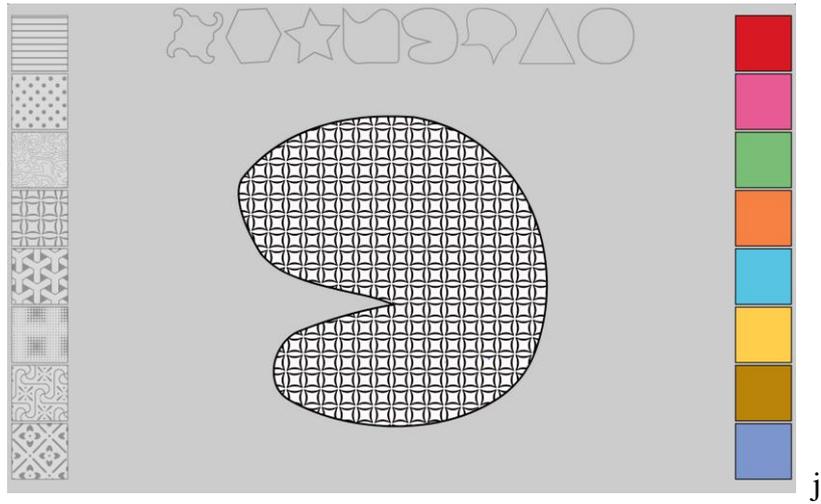
Appendix I. Screen captures of computer program for Experiment 3.



Appendix I continued. Screen captures of computer program for Experiment 3.



Appendix I continued. Screen captures of computer program for Experiment 3.



Appendix I continued. Screen captures of computer program for Experiment 3.

Appendix II

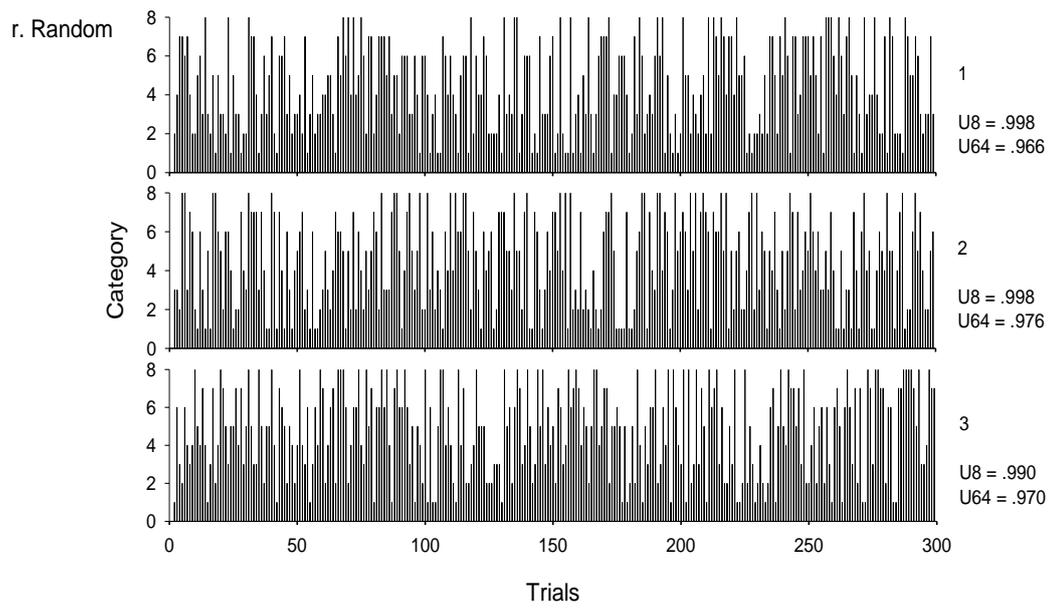
Bar graphs of categories used for each dimension for each participant in both Variability (VAR) and the yoking (YOKE) group and graphs for random responses.

Y-axis shows the number corresponding to the categories used; for number and category reference please refer to Appendix I.a. X-axis shows the trial numbers.

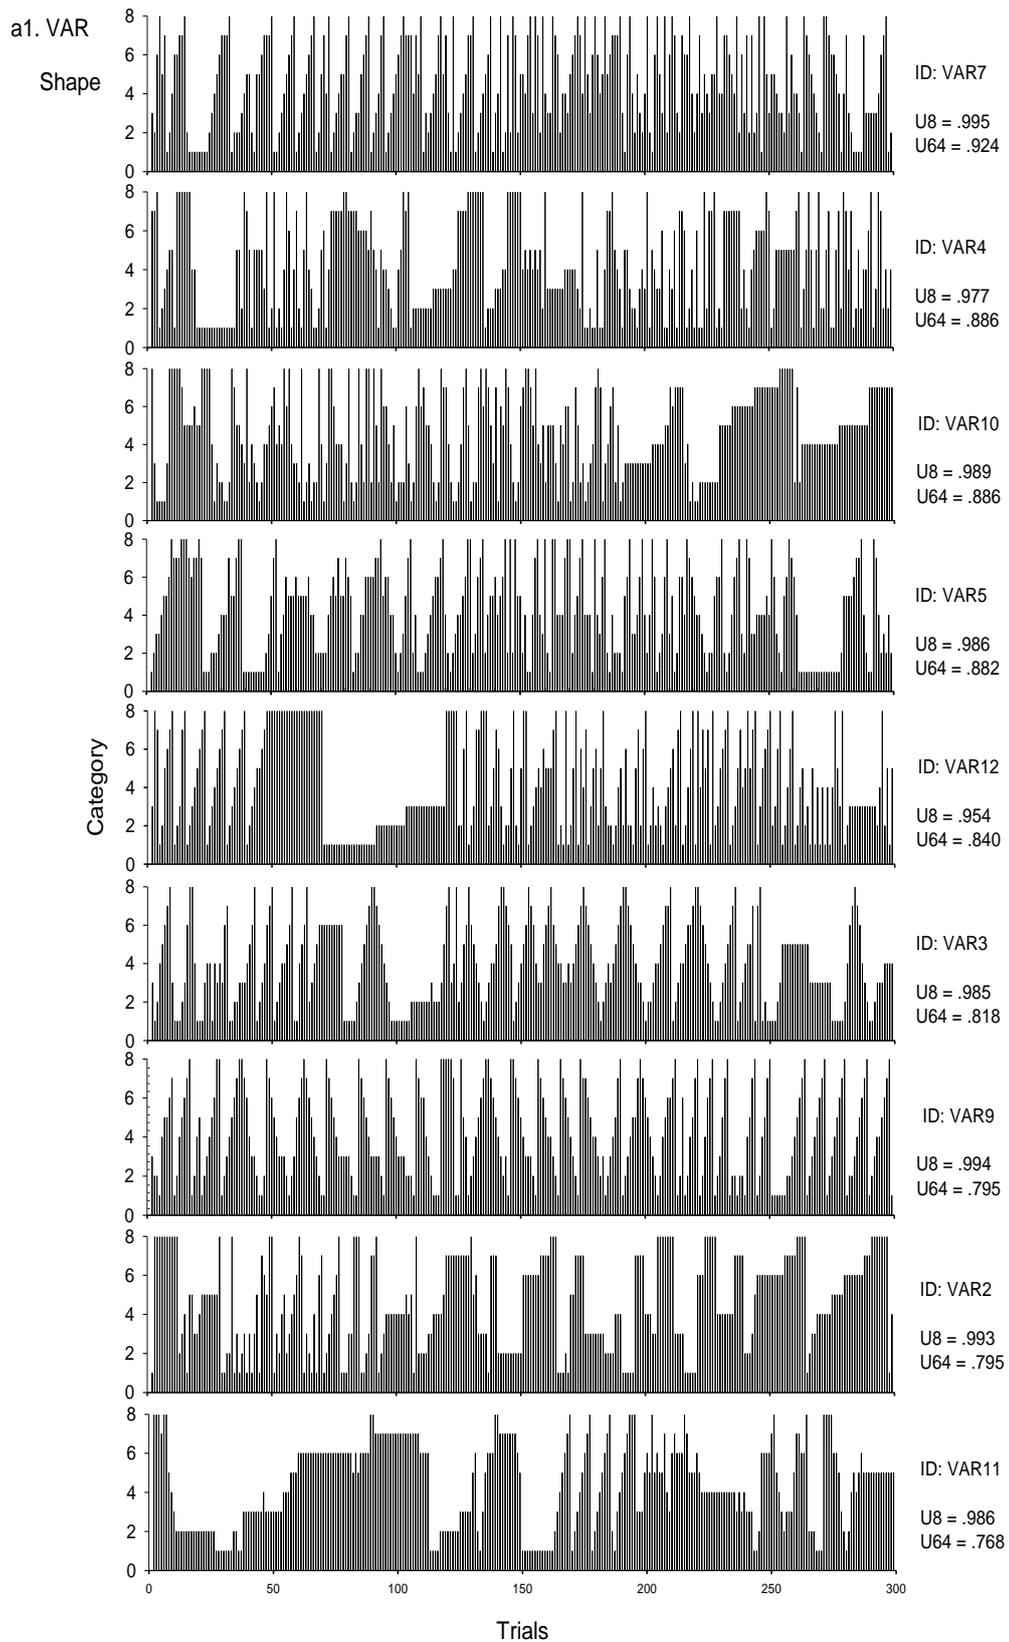
Participant identification numbers and the U-values (U8 and U64) calculated based on responses depicted in the bar graphs were shown next to the graph for each participant. U8 was the U-values calculated based on the original 8 categories; U64 was the U-values calculated based on the newly created 64 categories with consideration of the sequence of two consecutive trials.

The graphs for each dimension for each group were rearranged in descending order according to the values of U64.

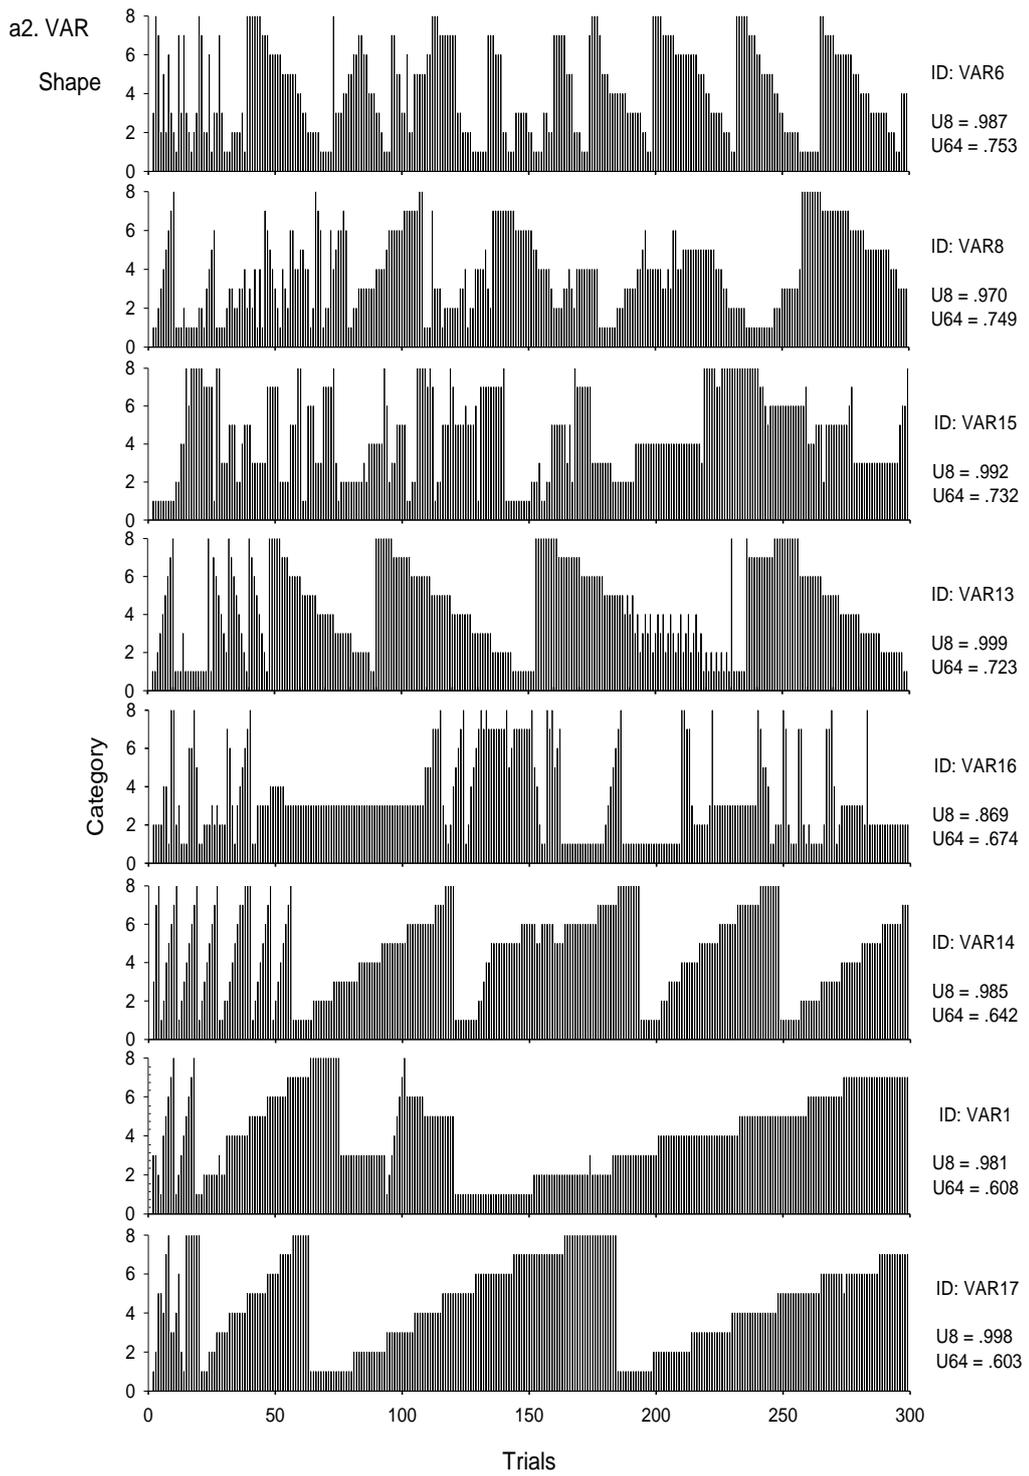
Panel r showed response graphs for the random responses generated by the computer. For the VAR group, panels a₁ and a₂ showed responses for shape, panels b₁ and b₂ showed responses for colour; and panels c₁ and c₂ showed responses for pattern. For the YOKE group, panels d₁ and d₂ showed responses for shape, panels e₁ and e₂ showed responses for colour and panels f₁ and f₂ showed responses for pattern.



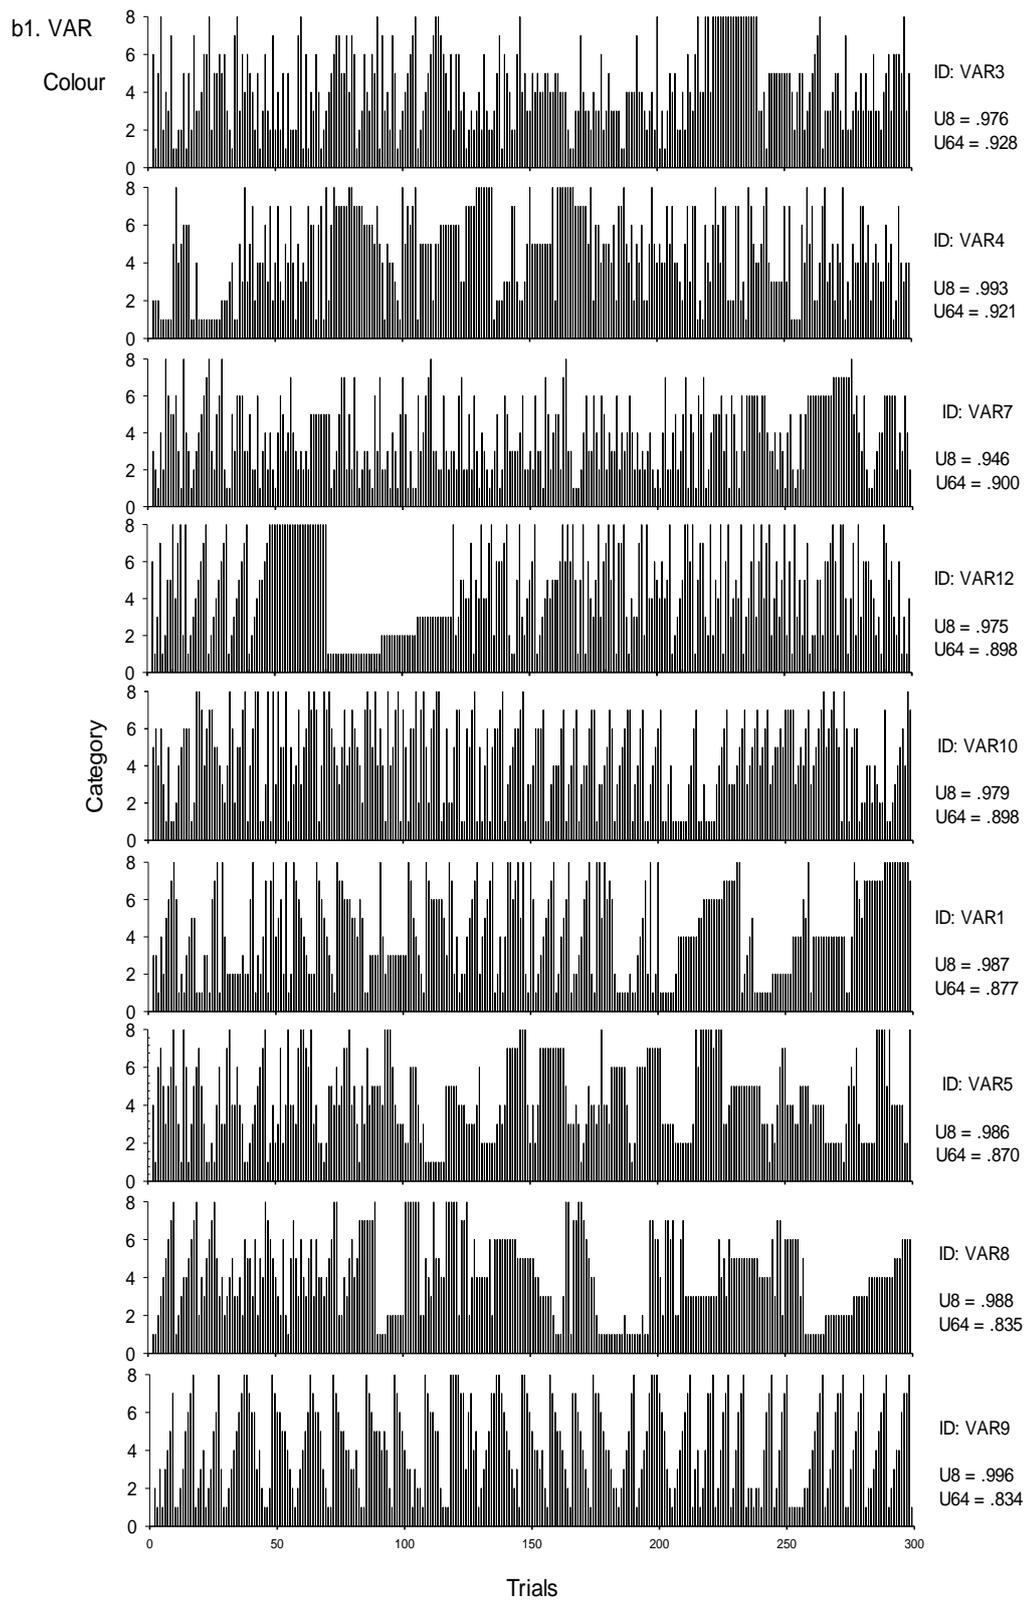
Appendix II.r. Responses generated by random generator and U-values calculated based on these responses.



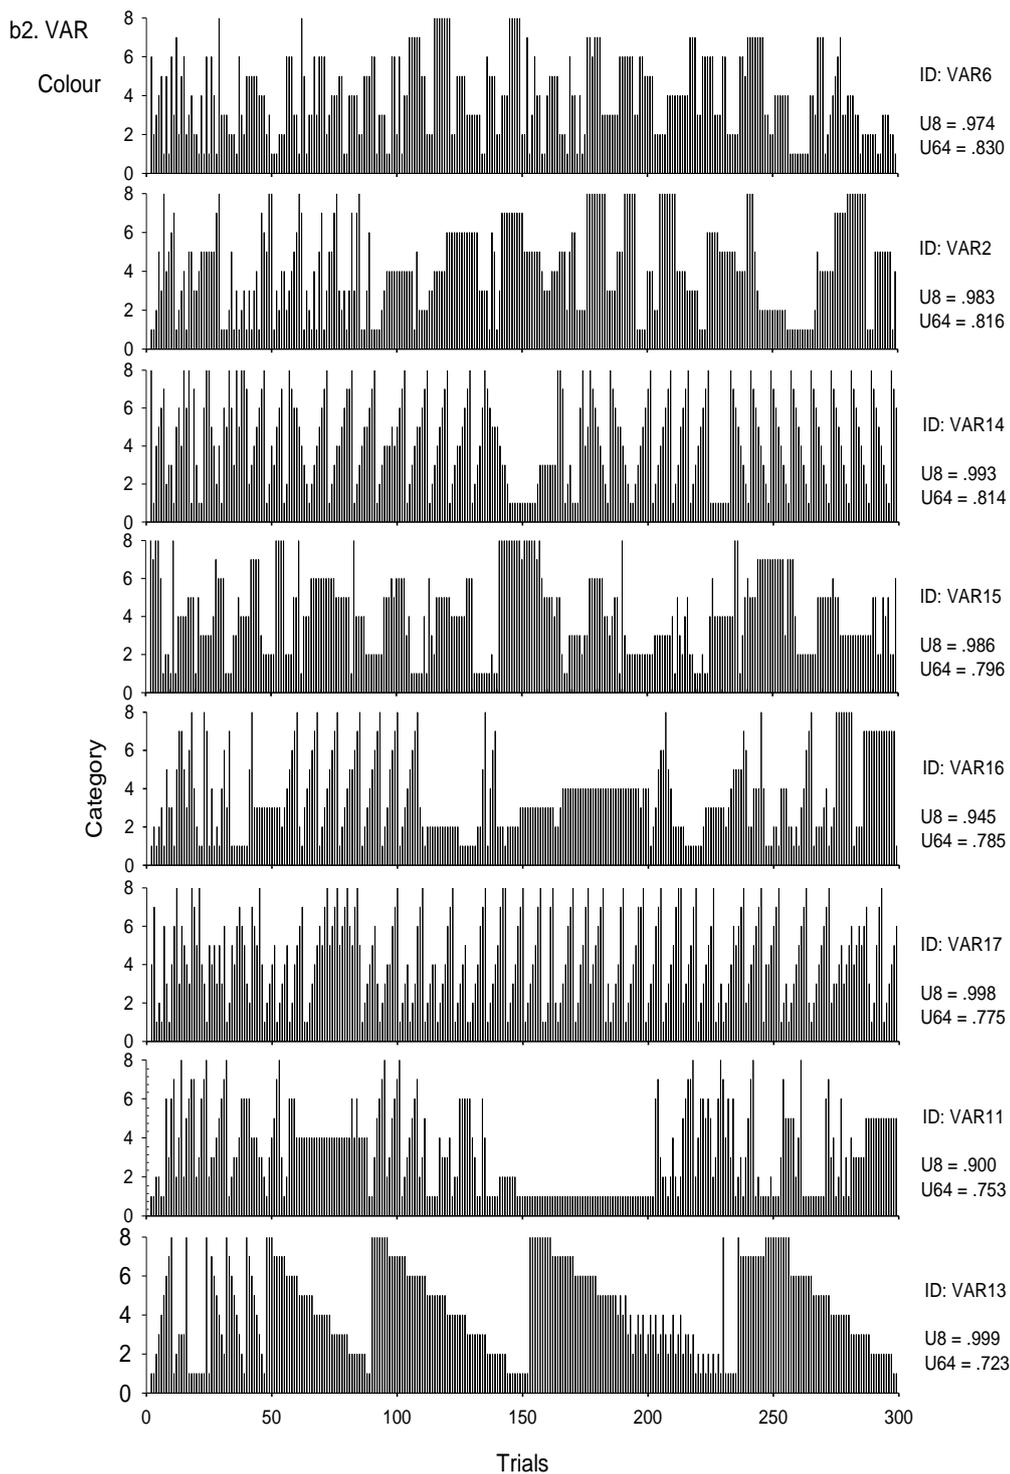
Appendix II. a. 1. Response graphs for VAR group for Shape dimension and U-values calculated based on these responses.



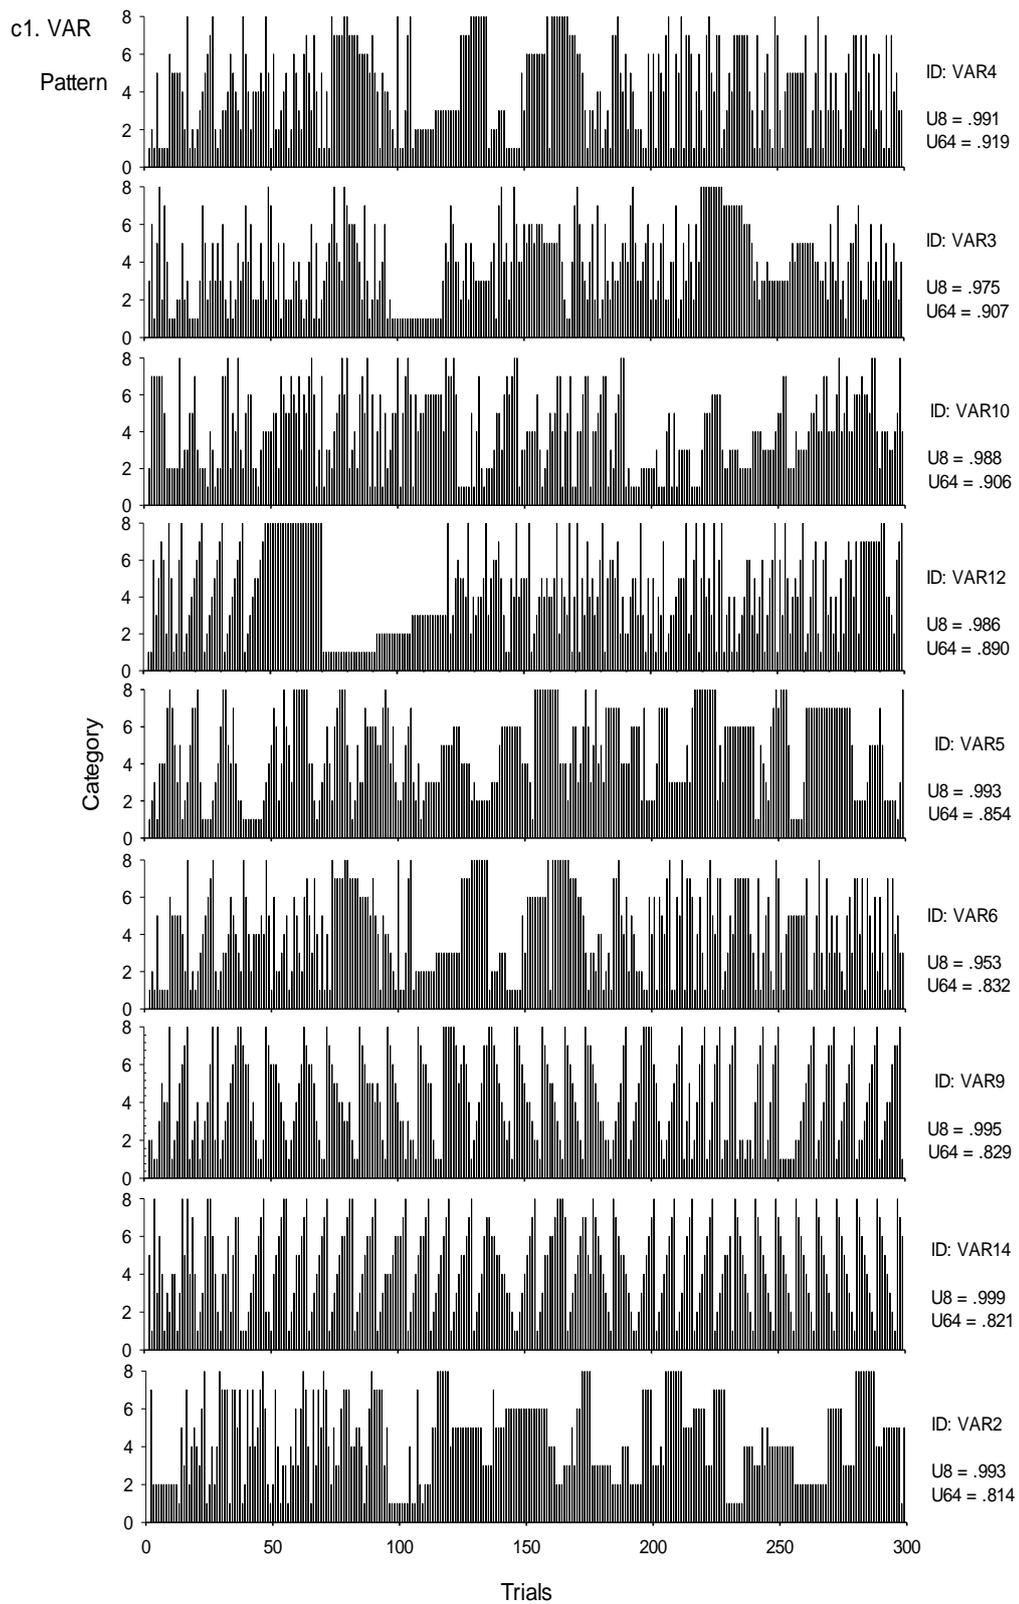
Appendix II. a. 2. Response graphs for VAR group for Shape dimension and U-values calculated based on these responses.



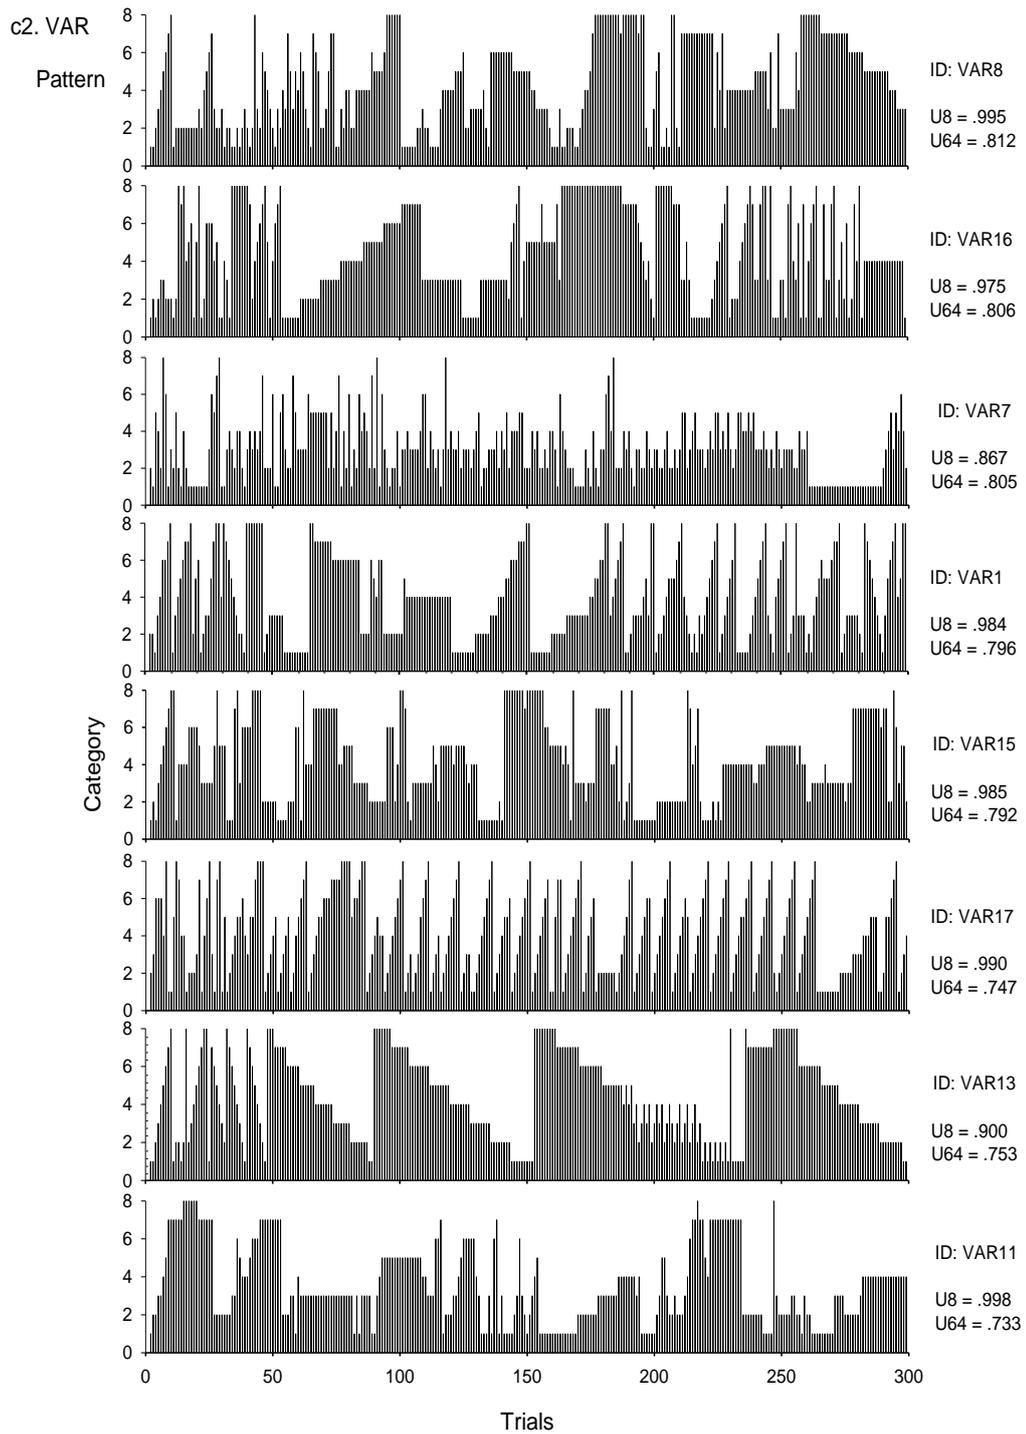
Appendix II. b. 1. Response graphs for VAR group for Colour dimension and U-values calculated based on these responses.



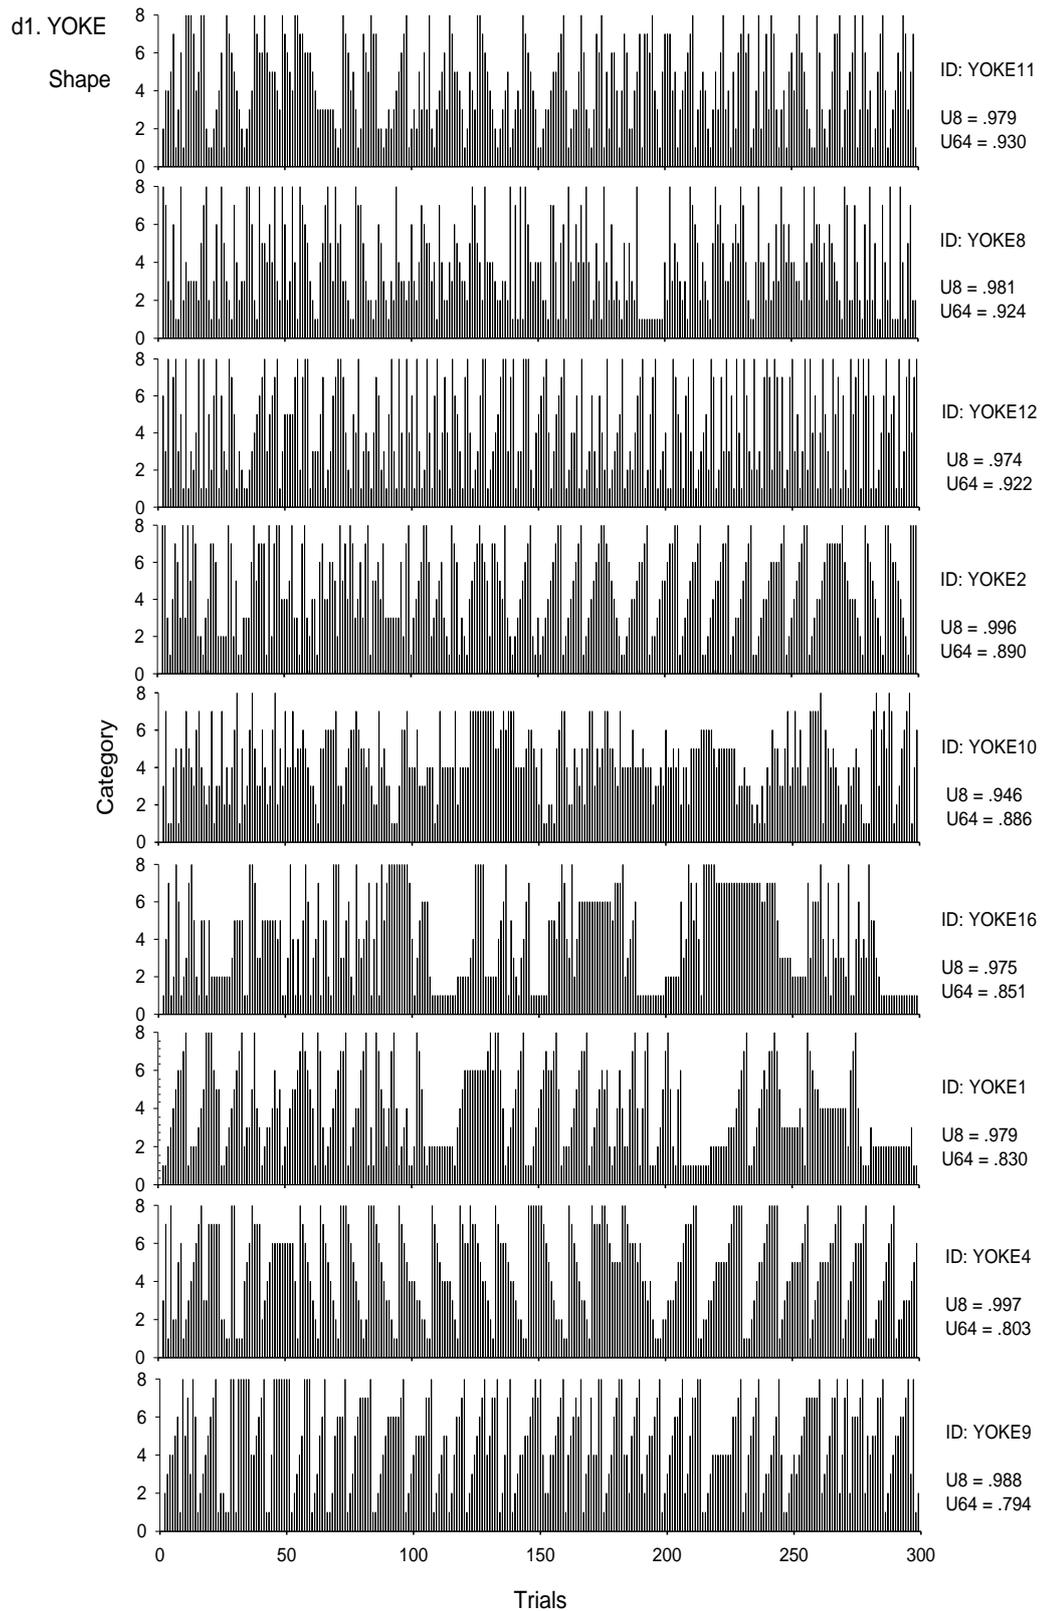
Appendix II. b. 2. Response graphs for VAR group for Colour dimension and U-values calculated based on these responses.



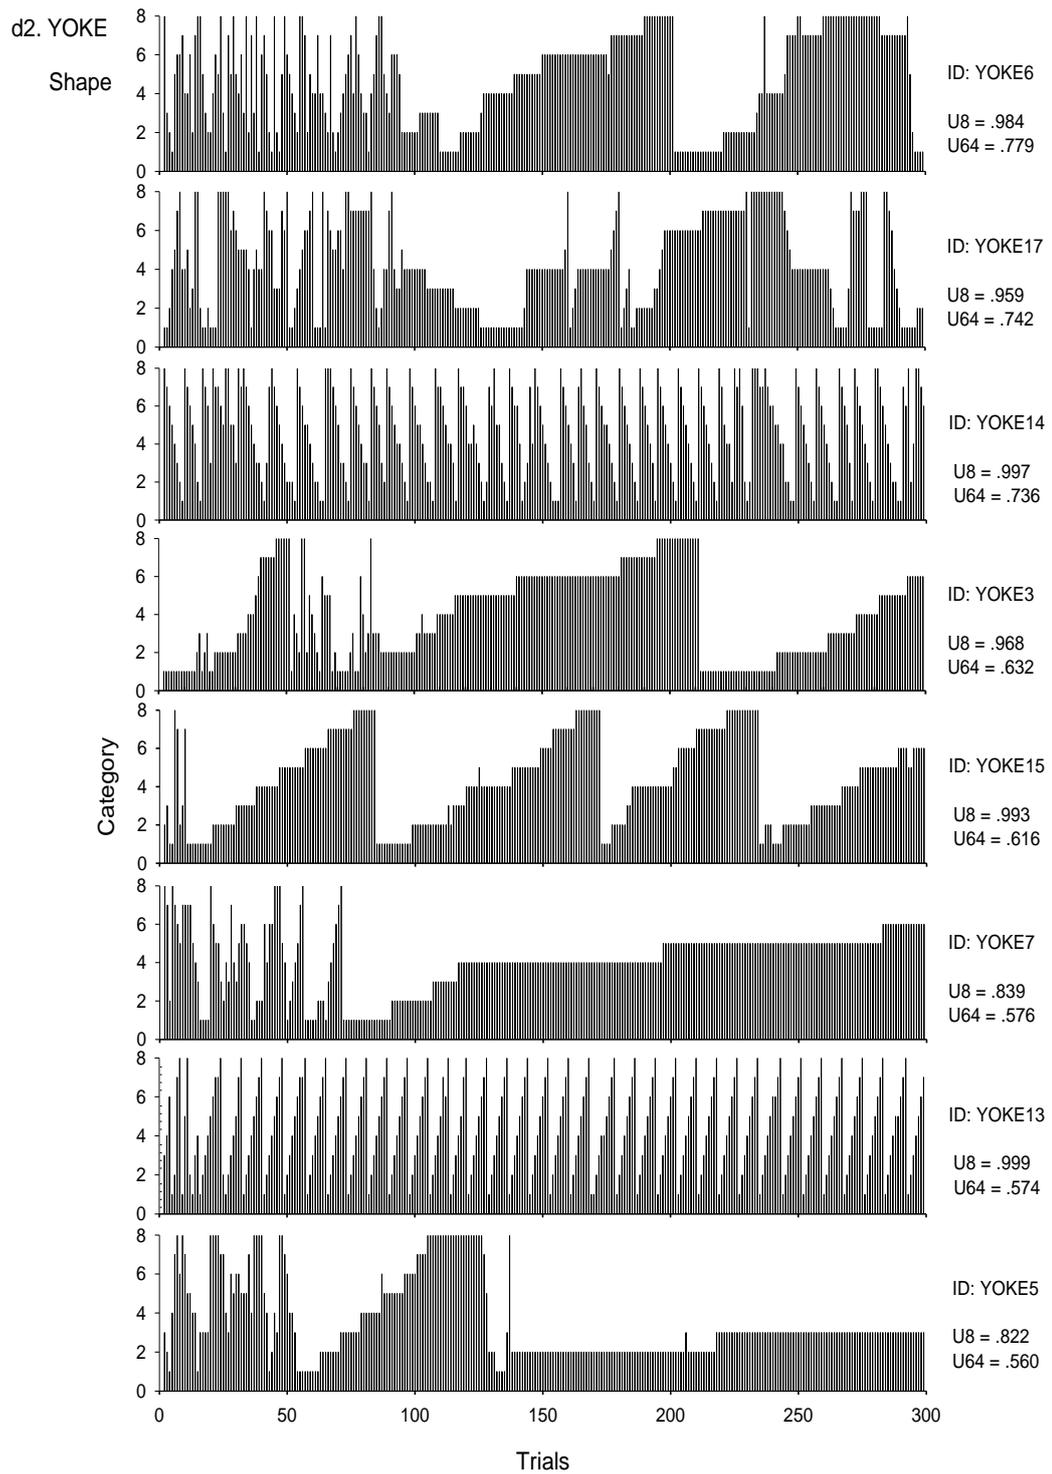
Appendix II. c. 1. Response graphs for VAR group for Pattern dimension and U-values calculated based on these responses.



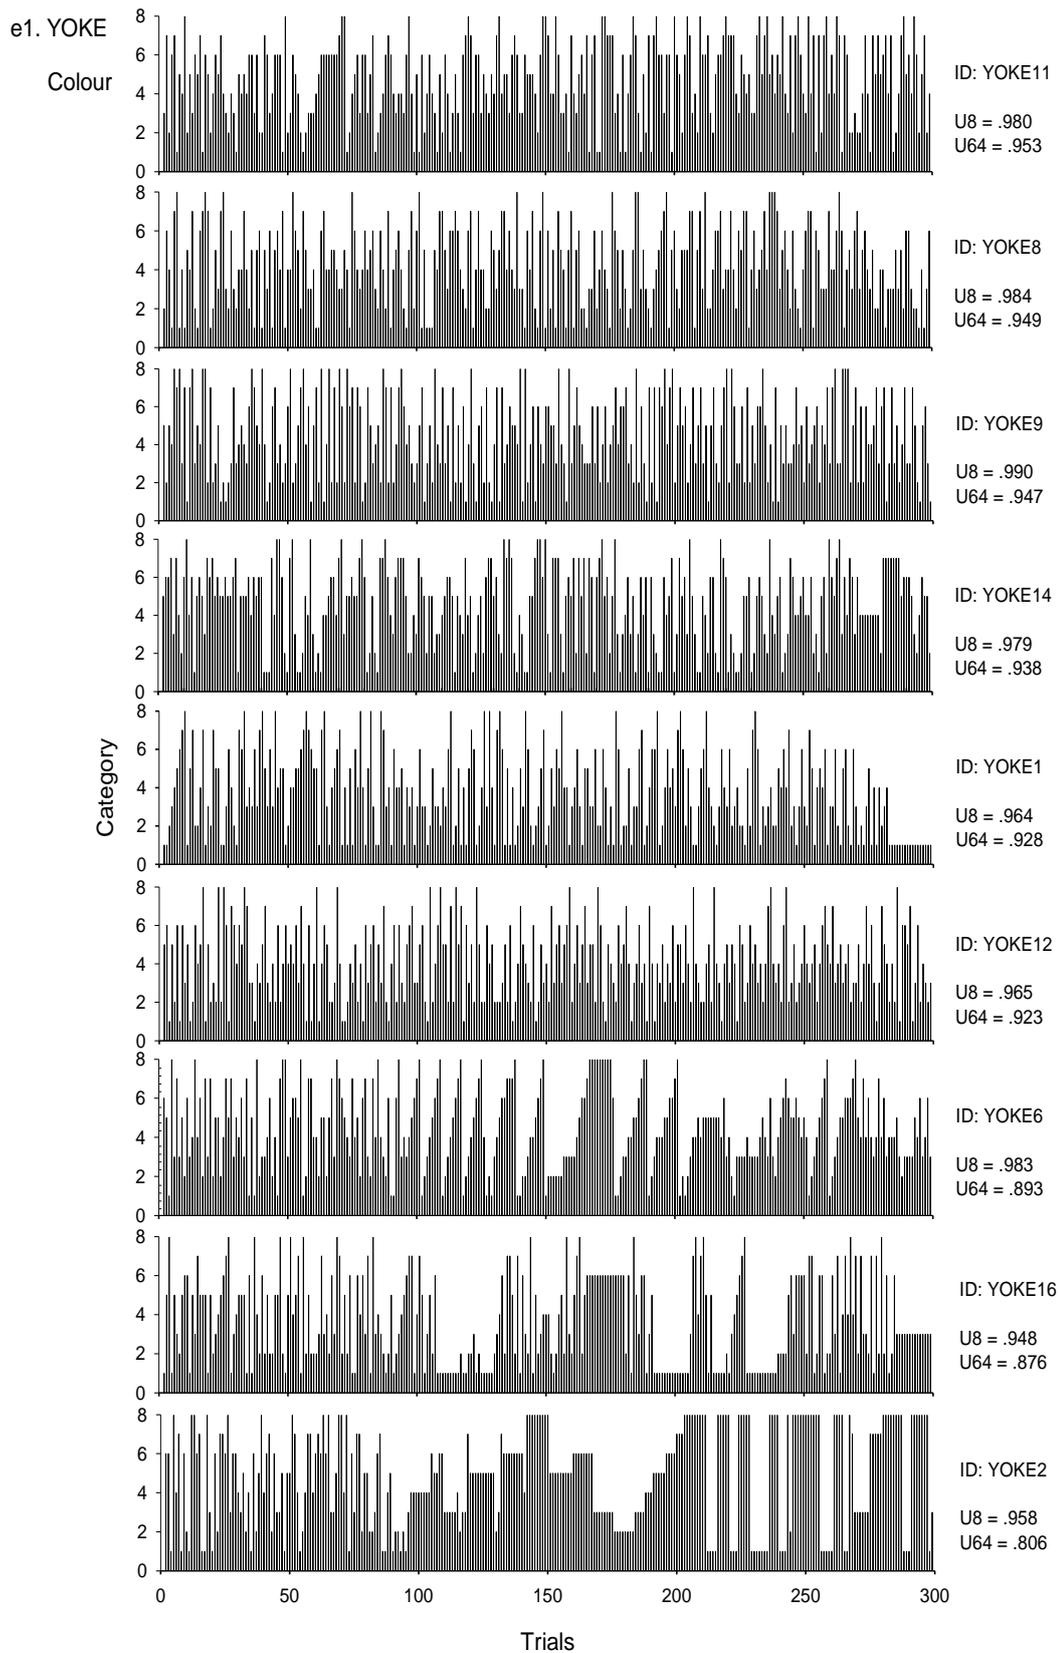
Appendix II. c. 2. Response graphs for VAR group for Pattern dimension and U-values calculated based on these responses.



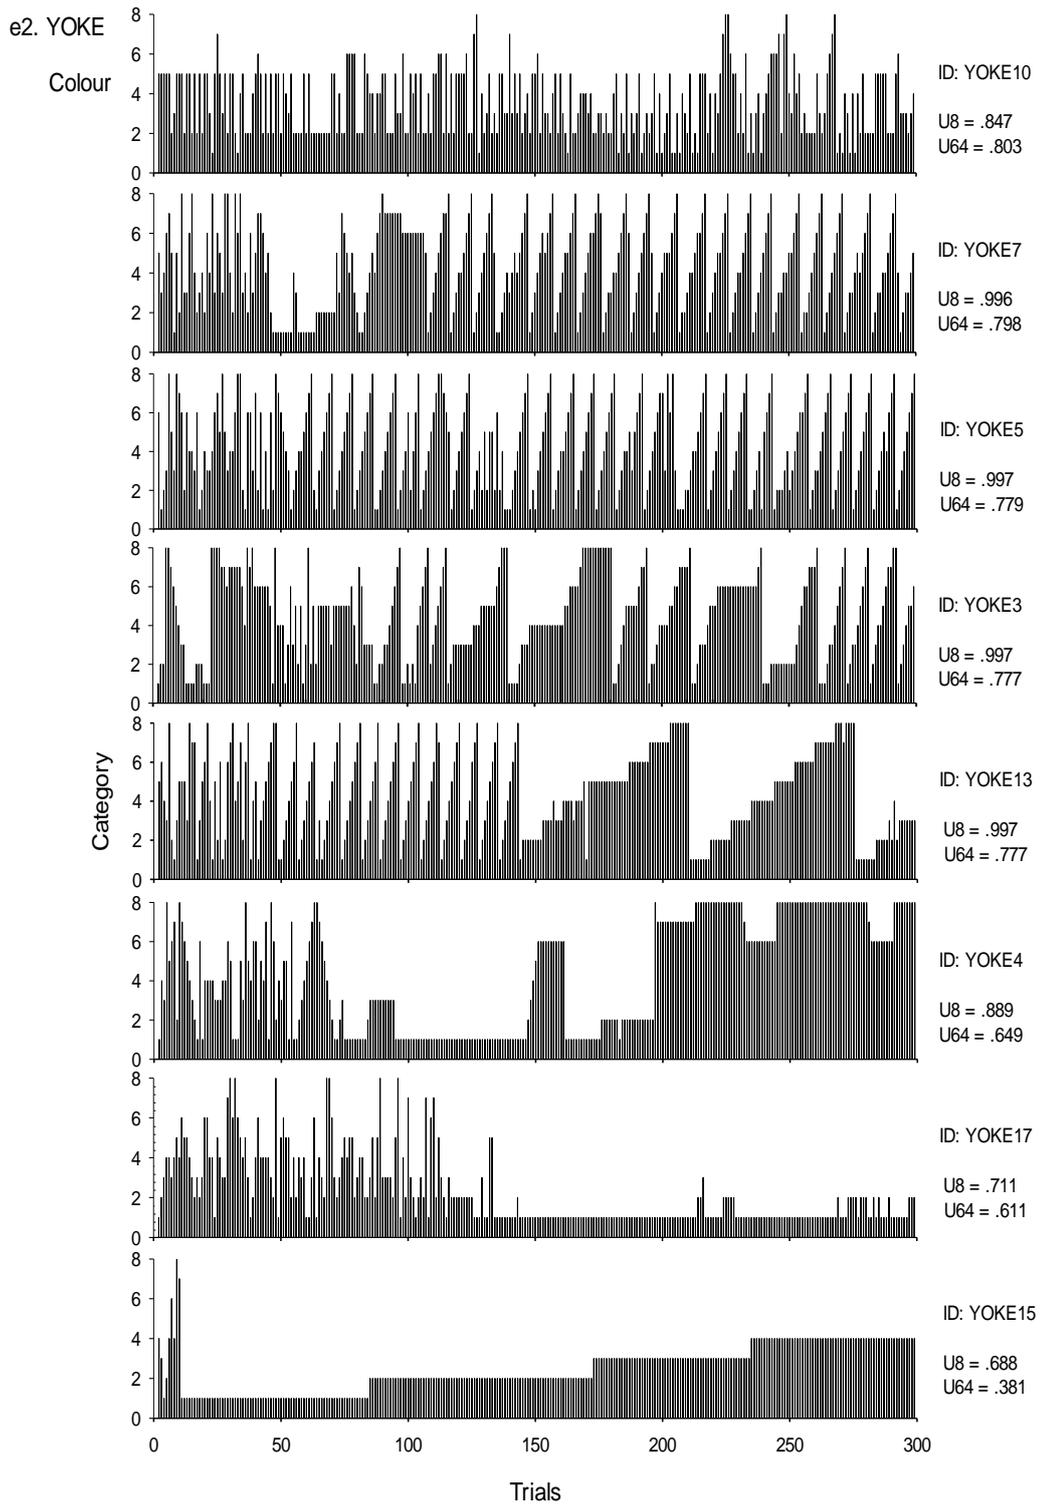
Appendix II. d. 1. Response graphs for YOKE group for Shape dimension and U-values calculated based on these responses.



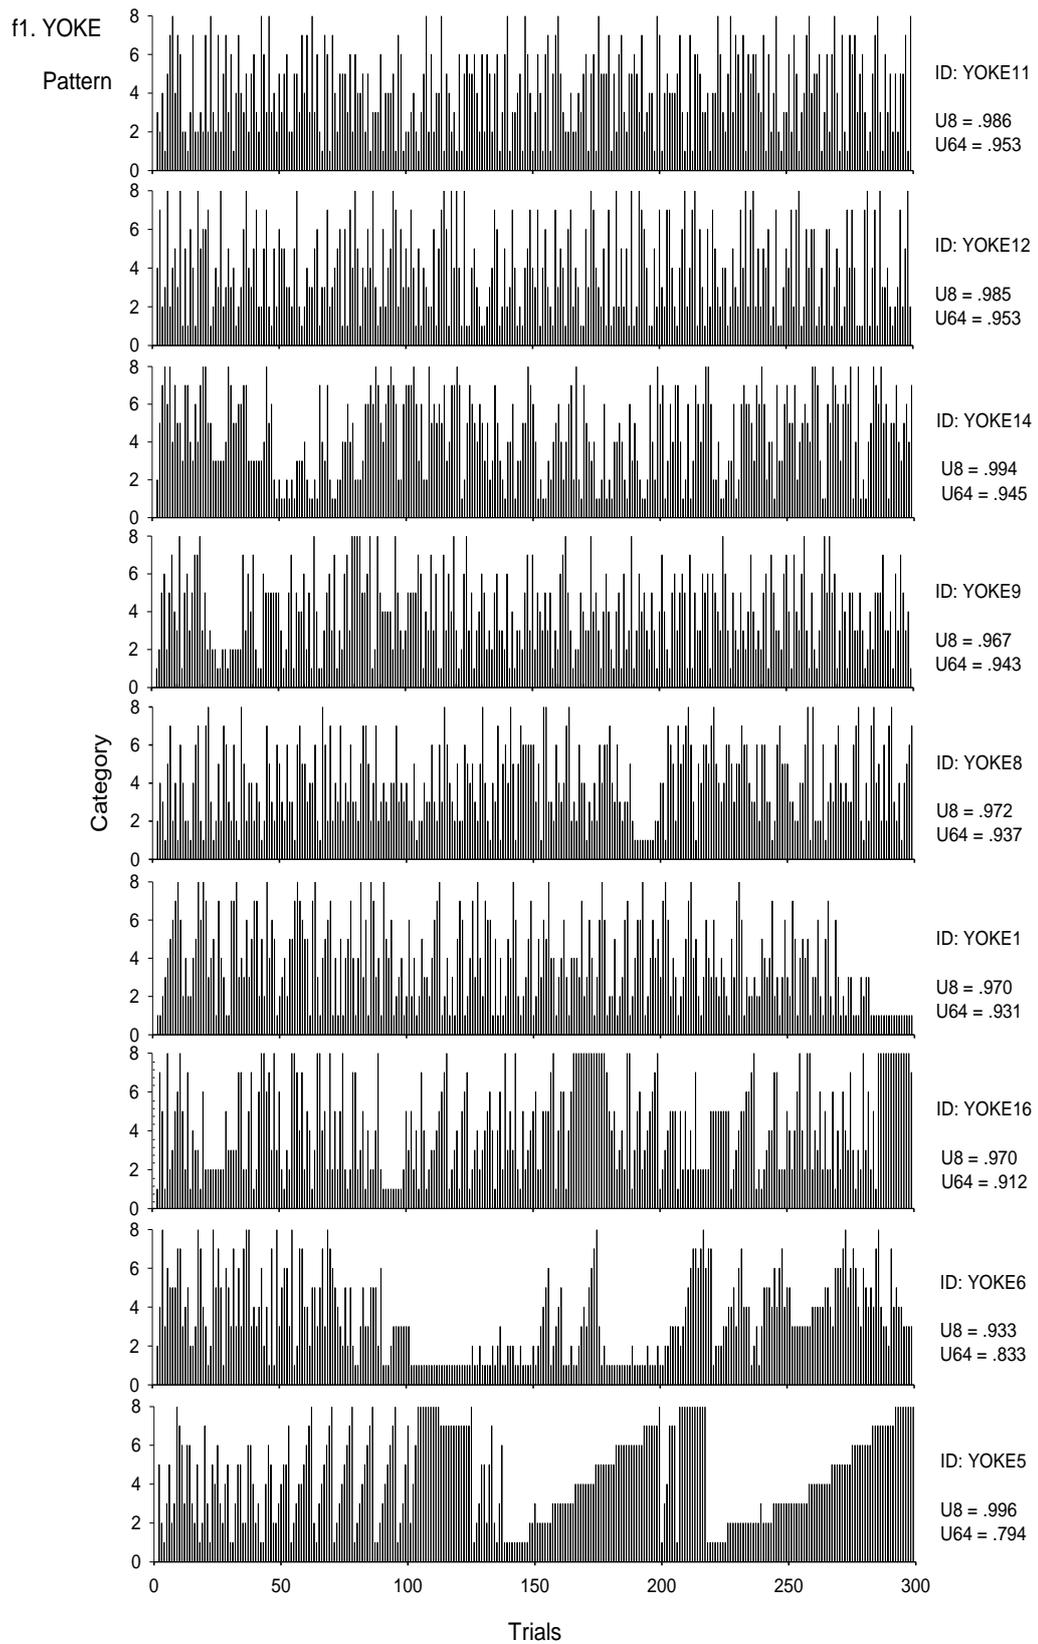
Appendix II. d. 2. Response graphs for YOKE group for Shape dimension and U-values calculated based on these responses.



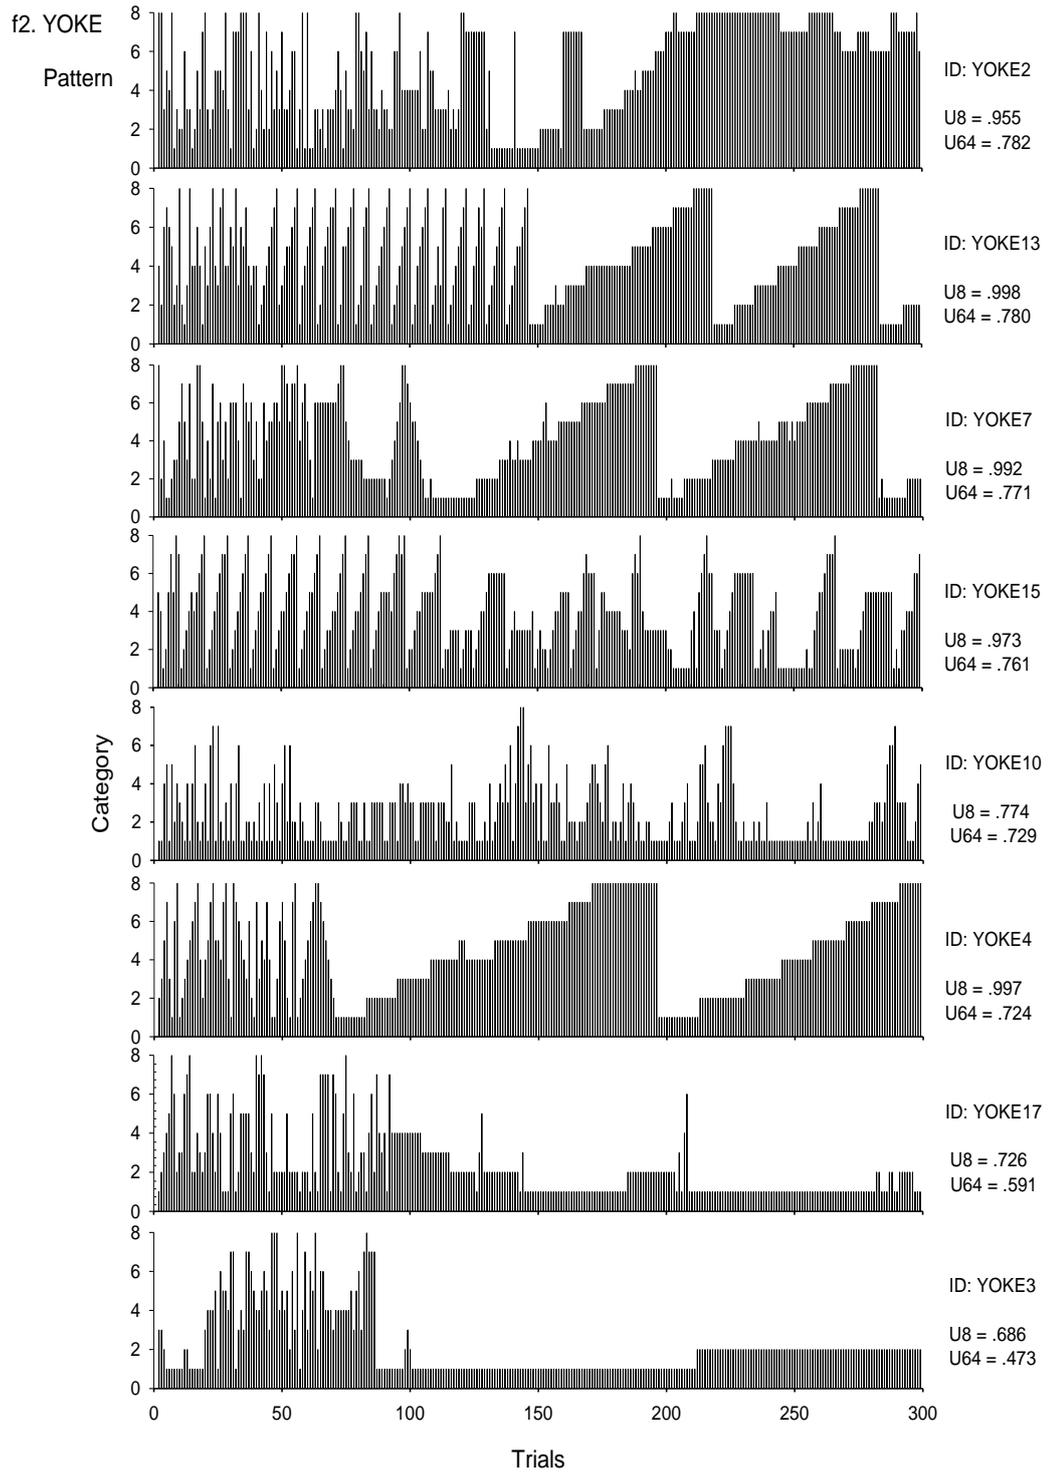
Appendix II. e. 1. Response graphs for YOKE group for Colour dimension and U-values calculated based on these responses.



Appendix II. e. 2. Response graphs for YOKE group for Colour dimension and U-values calculated based on these responses.

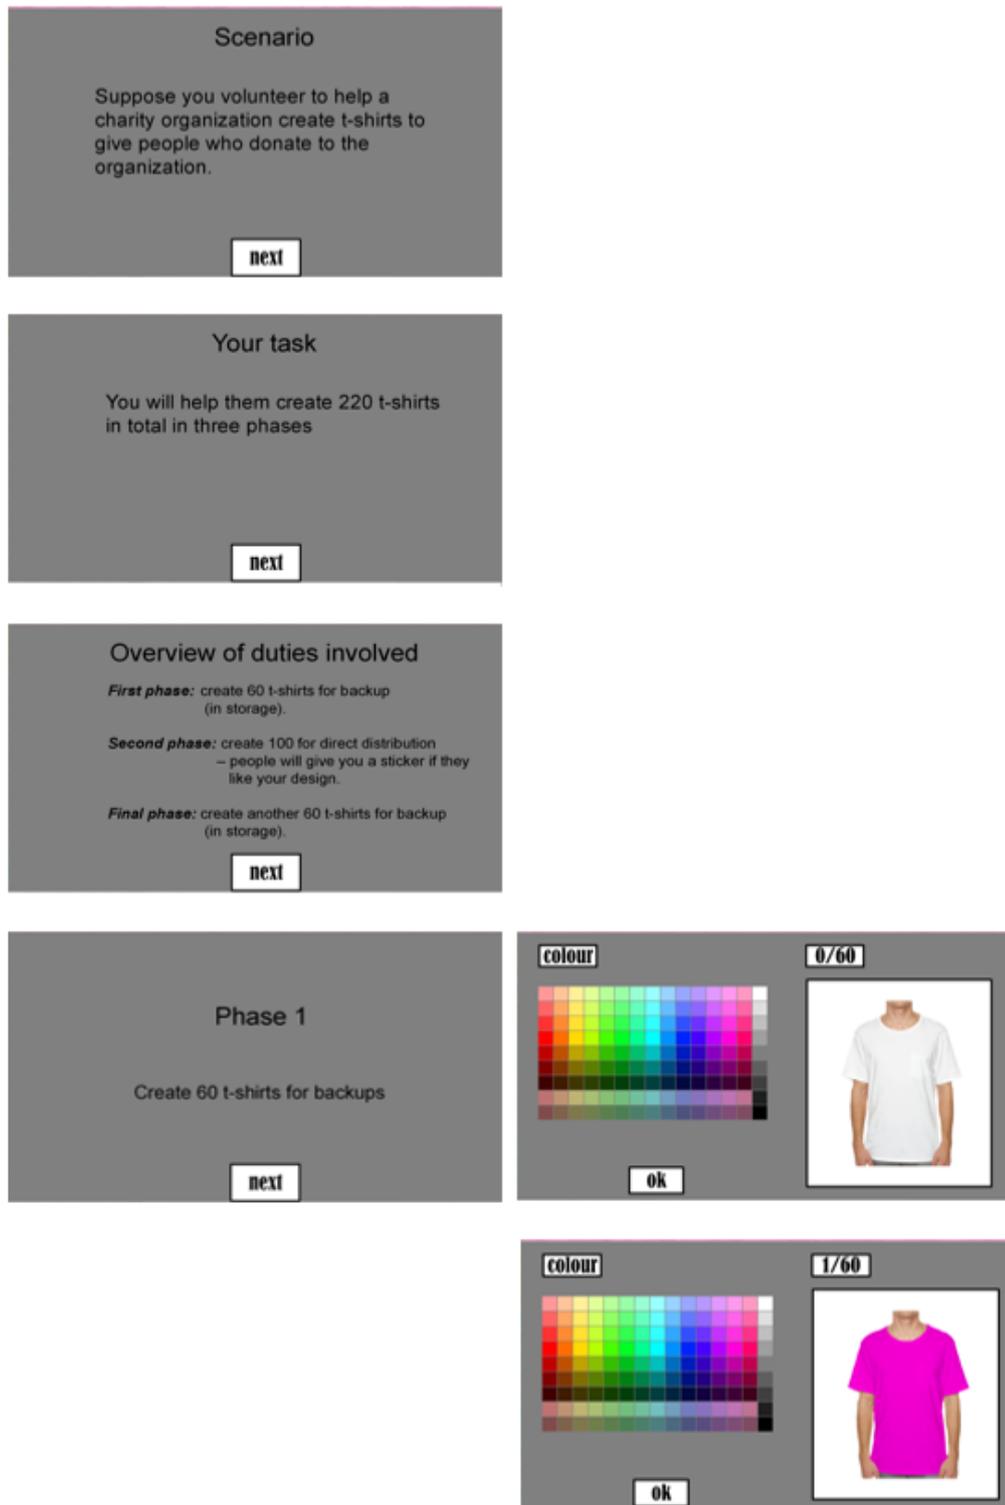


Appendix II. f. 1. Response graphs for YOKE group for Colour dimension and U-values calculated based on these responses.

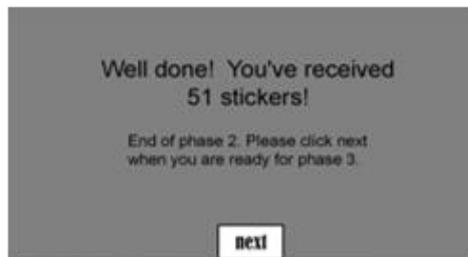
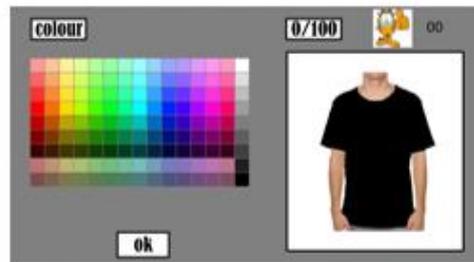


Appendix II. f. 2. Response graphs for YOKE group for Pattern dimension and U-values calculated based on these responses.

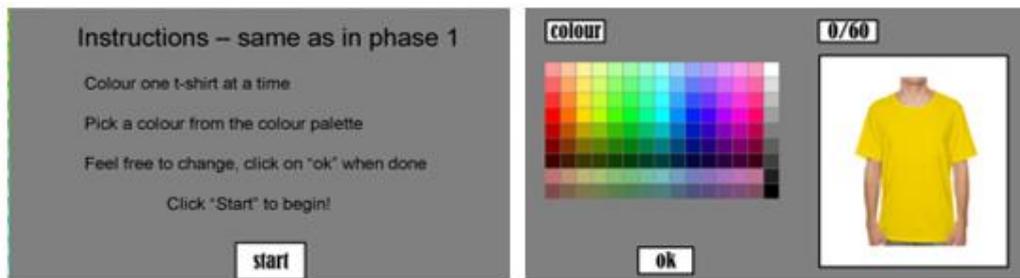
Appendix III



Appendix III. 1. Screen captures of computer program for Experiment 4, Overall instructions and Phase 1.



Appendix III. 2. Screen captures of computer program for Experiment 4, Phase 2.



Appendix III. 3. Screen captures of computer program for Experiment 4, Phase 3.